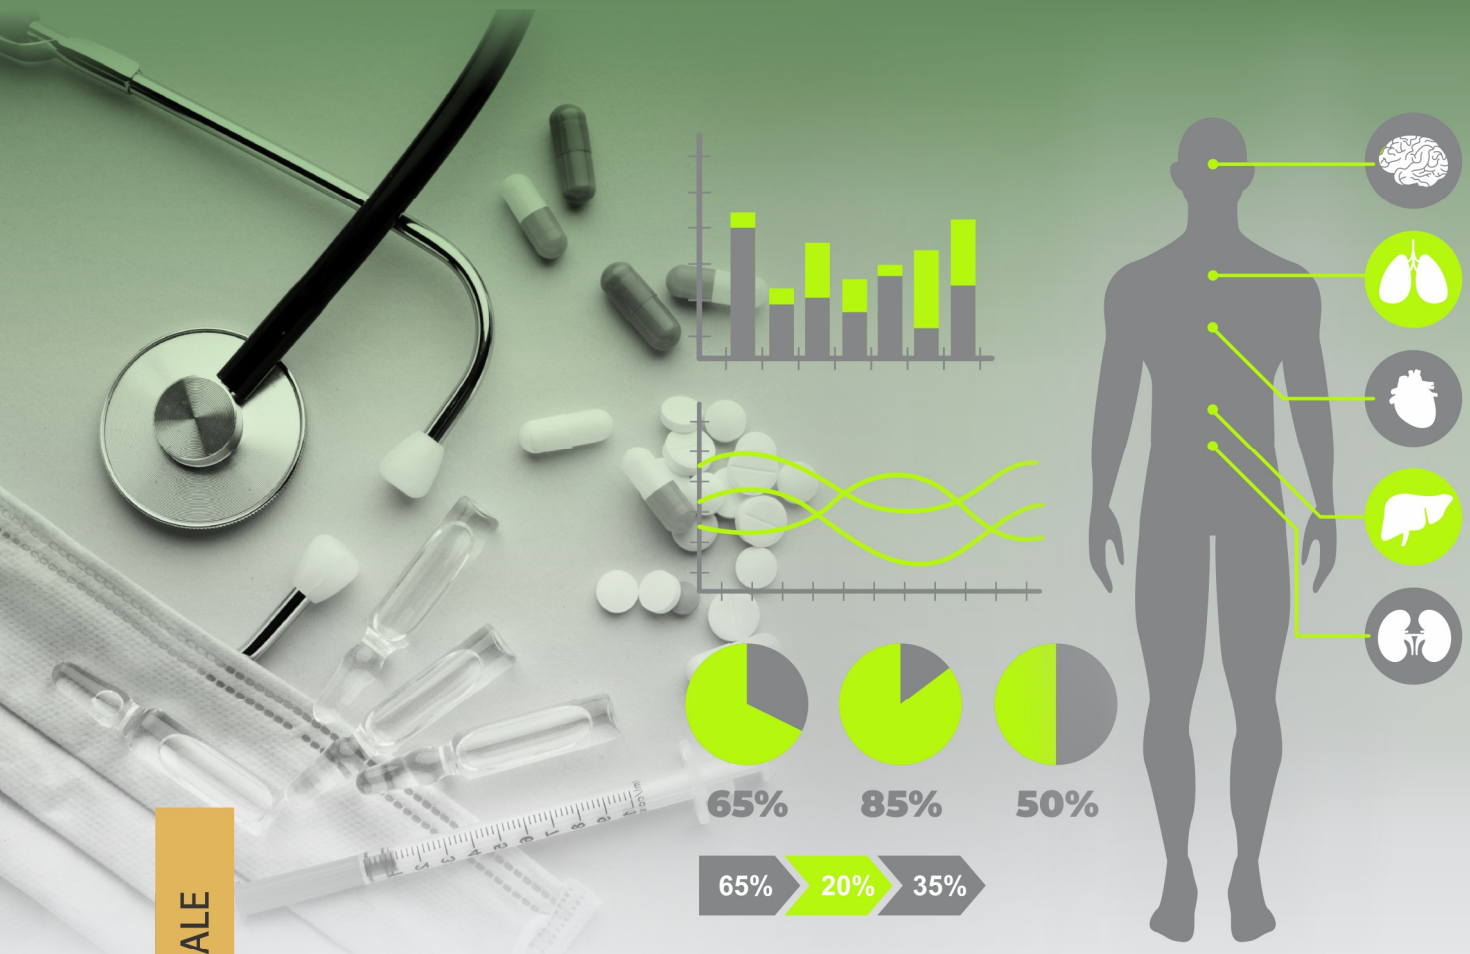
„VICTOR BABEŞ" UNIVERSITY
OF MEDICINE AND PHARMACY
FROM TIMIŞOARA

CĂLIN MUNTEAN

MEDICAL INFORMATICS

Methods, Systems and Practice. Horizont 2026

Lectures and practical sessions for medical students



MANUALE

„Victor Babeș" Publishing House
Timișoara, 2026

Editura „Victor Babeş”
Piața Eftimie Murgu nr. 2, cam. 316, 300041 Timişoara
Tel./Fax 0256 495 210
e-mail: evb@umft.ro
https://www.umft.ro/ro/organizare-evb/

Director: Prof. univ. dr. Sorin Ursoniu

Colecția: MANUALE
Coordonatori colecție: **Prof. univ. dr. Codruța Şoica**
Prof. univ. dr. Daniel Lighezan

Referent științific: Prof. univ. dr. Daliborca Vlad

Indicativ CNCSIS: 324

© 2026
Toate drepturile asupra acestei ediții sunt rezervate. Reproducerea parțială sau integrală a textului, pe orice suport, fără acordul scris al autorilor este interzisă și se va sancționa conform legilor în vigoare.

ISBN 978-606-786-642-1

The patient examples in this book are fictitious. The information on legislation, national platforms, and standards reflects the situation as of September 27, 2026. This book does not replace the clinical guidelines and regulations in force, which must be consulted in their current versions.

Contents

Preface ix
How to Read This Book x
List of Abbreviations xii
Part I. Lectures 1
Chapter 1. Medical Informatics: Subject, Domains, Competencies 2
1.1 Definitions and Subject 2
1.2 Historical Milestones: From the First Hospital Computers to Generative Artificial Intelligence 4
1.3 Domains of Biomedical Informatics: From Molecules to Populations 6
1.4 Digital Competencies of the Physician 7
1.5 Searching the Scientific Literature 9
1.6 Medical Informatics in Romania and in Timişoara 11
1.7 How to Use This Book 12
Chapter 2. Data, Information, and Knowledge in Medicine 15
2.1 Medical Data 15
2.2 Information Theory Applied in Medicine 17
2.3 Data, Information, Knowledge, Wisdom 21
2.4 Communicating Information in Medical Practice 22
2.5 Evidence-Based Medicine in Digital Format 24
2.6 Information in Biology 25
Chapter 3. Structured Data, Files, and Databases 28
3.1 Files and Formats 28
3.2 Data Modeling 29
3.3 Relational Database Management Systems and SQL 31
3.4 NoSQL Databases, Data Warehouses and Data Lakes, Big Data 32
3.5 Common Data Models for Research 34
3.6 Data Quality and Medical Registries 36
3.7 Electronic Data Capture Tools 38
Chapter 4. Medical Classifications, Nomenclatures, and Terminologies 42
4.1 Why We Code 42
4.2 ICD-10 and National Modifications 43
4.3 ICD-11 45
4.4 Other WHO Classifications 49

4.5 Clinical Terminologies 49
4.6 Diagnosis-Related Groups (DRG) 51
4.7 Computer-Assisted Coding and Artificial Intelligence 53
Chapter 5. Ethics, Data Protection, and Information Security in Health Care 56
5.1 Ethical Principles Applied to Health Data 56
5.2 The GDPR for Physicians 57
5.3 The European Health Data Space 61
5.4 Information Security 61
5.5 Equity and Algorithmic Bias 64
5.6 Academic Integrity 65
Chapter 6. Biomedical Signals: Acquisition, Processing, and Assisted Interpretation 68
6.1 The Biomedical Signal 68
6.2 The Acquisition Chain 70
6.3 Digital Processing 72
6.4 Artifacts and Signal Quality 75
6.5 Formats and Open Databases 76
6.6 Interpretation Assisted by Artificial Intelligence 77
6.7 Romania in 2026 79
Chapter 7. Digital Medical Imaging: DICOM, PACS, and Computer-Assisted Analysis 82
7.1 The Digital Image 82
7.2 Imaging Modalities from an Informatics Perspective 84
7.3 Image Compression 85
7.4 The DICOM Standard 86
7.5 PACS, RIS, VNA, and the Radiology Workflow 88
7.6 Image Processing and Analysis 90
7.7 Artificial Intelligence in Imaging, 2026 90
Chapter 8. The Electronic Health Record and Interoperability 95
8.1 The Electronic Record: Concepts and Effects 95
8.2 Interoperability 98
8.3 The International Patient Summary and Cross-Border Services 102
8.4 The European Health Data Space 103
8.5 Romania in 2026 104
8.6 Usability, Safety, and Human Factors 106
Chapter 9. Telemedicine: Forms, Evidence, Limits, and Legal Framework 109
9.1 Definitions and Taxonomy 109

9.2 Technical Infrastructure 111
9.3 Evidence of Effectiveness, 2020-2026 113
9.4 Limits and Risks 116
9.5 Legal and Reimbursement Framework 117
9.6 Documenting the Teleconsultation 118
Chapter 10. Mobile Health, Wearable Devices, and Language Models in the Consultation 122
10.1 Mobile Health 122
10.2 Wearable Devices and Sensors 124
10.3 Software as a Medical Device 127
10.4 Device Interoperability 129
10.5 Large Language Models in the Consultation 130
10.6 Romania in 2026 133
Chapter 11. The Information System, Information Security, and the Electronic Signature 136
11.1 The Information System and the Computer-Based Information System 136
11.2 Information Security Management 139
11.3 Cryptography Basics for Physicians 141
11.4 The Electronic Signature 143
11.5 Standardization in Medical Informatics 144
11.6 The Human Factor 146
Chapter 12. Hospital Information Systems and the Single Integrated Information System of Romanian Health Insurance 150
12.1 The Hospital's Information System: Components and Modules 150
12.2 Electronic Prescribing in the Hospital and Clinical Decision Support 153
12.3 Case Study: Computerizing Intensive Care 154
12.4 The Romanian Health Insurance Information System 156
12.5 The Digital Transformation of Hospitals and Digital Maturity 158
12.6 Indicators and Dashboards 159
12.7 Resilience: When the System Stops 160
Chapter 13. Artificial Intelligence in Medicine: Methods, Evaluation, Regulation (2026) 163
13.1 What Artificial Intelligence in Medicine Is 163
13.2 The Life Cycle of a Model 165
13.3 Performance Evaluation 167
13.4 Clinical Evidence 2020-2026 170
13.5 Explainable AI 172
13.6 Risks 172

13.7 Regulation in 2026 173
13.8 Further Reading 175
Chapter 14. Synthesis: The Physician's Competencies in the Digital Environment and the Ethics of Using Artificial Intelligence 178
14.1 The Map of the Discipline 178
14.2 Competencies at the End of the First Year 179
14.3 The 2024-2031 Regulatory Calendar That Will Change Practice 180
14.4 The Ethics of AI Use by Students and Physicians 183
14.5 Directions Beyond 2026 185
14.6 Review Questions 186
Part II. Practical Sessions 192
Practical Session 1. The Computer, the Operating System, and the Collaborative Workspace (Google Workspace) 193
L1.1 The Computer and Files: Basic Notions 193
L1.2 Sharing, Versions, and Account Security 195
Practical Session 2. Medical Information on the Web: Searching and Evaluating Quality 198
L2.1 From the Clinical Question to the Query 198
L2.2 Wikipedia: A Starting Point, Not a Source to Cite 200
L2.3 Tools for Evaluating Information for Patients 200
L2.4 A Chatbot's Answer Compared with the Guideline 200
Practical Session 3. Electronic Data Collection Forms: Google Forms and Epi Info 203
L3.1 The Protocol and the Data Dictionary 204
L3.2 Building the Form in Google Forms 205
L3.3 The Pilot Test 206
L3.4 Collection and Export 206
L3.5 Data Cleaning in Google Sheets 206
L3.6 The Descriptive Table 207
L3.7 The Epi Info 7 Variant 207
L3.8 Data Protection 208
Practical Session 4. Coding Diagnoses with the WHO Online Tools (ICD-10 and ICD-11) 209
L4.1 Preparation (5 minutes) 210
L4.2 The ICD-10 Tabular List (15 minutes) 210
L4.3 The ICD-11 Browser and the Coding Tool (15 minutes) 211
L4.4 Coding Ten Cases (35 minutes) 211
L4.5 Postcoordination Exercise (10 minutes) 212

L4.6 ICD-10-ICD-11 Mapping (10 minutes) 212
L4.7 Underlying Cause of Death and Order within the Cluster (10 minutes) 212
Practical Session 5. Correct Citation, Plagiarism Prevention, and Writing a Scientific Essay 215
L5.1 Anatomy of a Vancouver Reference 216
L5.2 The Reference Manager 217
L5.3 Paraphrase, Quotation, Copying 218
L5.4 The Similarity Report 219
L5.5 Declaring the Use of Generative AI 219
L5.6 Structure of the Essay 219
Practical Session 6. Analysis of ECG and EEG Signals with Free Software 221
Practical Session 7. Medical Image Processing with ImageJ/Fiji 227
L7.1 Installation and Setup 228
L7.2 Opening the Image and Reading the DICOM Header 229
L7.3 The Histogram and the Display Window 229
L7.4 Spatial Calibration and Measurements 229
L7.5 Noise Filtering 230
L7.6 Threshold Segmentation and Object Counting 230
L7.7 Recording a Macro and Exporting the Results 231
L7.8 Discussion: Limitations and Reproducibility 231
Practical Session 8. Working with an Electronic Health Record: The OpenEMR Demo and FHIR Resources 233
L8.1 Part A: The Fictitious Patient in OpenEMR 234
L8.2 Part B: The Same Concepts as FHIR Resources 235
L8.3 Part C: Comparison with the Romanian E-Prescription 236
Practical Session 9. Simulating a Teleconsultation: Preparation, Consent, Documentation, and Limits 239
Practical Session 10. Simulated Consultation with a Language Model: Verification, Errors, and Disclosure 245
Practical Session 11. Practical Security: Passwords, Multi-Factor Authentication, Verifying an Electronic Signature, and Recognizing Phishing 251
L11.1 How Strong Is a Password? (20 Minutes) 252
L11.2 2-Step Verification on the Google Account (15 Minutes) 252
L11.3 Signing and Verifying a Fictitious PDF in Acrobat Reader (35 Minutes) 253
L11.4 The TLS Certificate of a Website (10 Minutes) 254
L11.5 Recognizing Phishing: A Test with Fictitious E-mails (25 Minutes) 254
L11.6 The Incident Reporting Procedure (10 Minutes) 255

Practical Session 12. Information Flow in the Physician's Office, Outpatient Clinic, and Hospital: Modeling and Indicators 257
L12.1 The Scenario 258
L12.2 The Swimlane Diagram in draw.io 259
L12.3 Documents and Data Elements 259
L12.4 The Data Set and the Spreadsheet 260
L12.5 Calculating the Indicators 261
L12.6 Interpretation 262
Practical Session 13. Critical Appraisal of an AI Tool: From the Confusion Matrix to Debate 263
Practical Session 14. The Practical Examination: Sample Paper, Grading Scale, and Review 269
L14.1 The Sample Paper 270
L14.2 Solutions and Grading Scale 272
L14.3 Frequent Mistakes 273
L14.4 Preparation Plan 273
Appendix A. Answers to the Self-Assessment Questions 275
Appendix B. Glossary 276
References 281

Preface

Why does a medical informatics book carry a date as precise as "Horizon 2026"? Because, unlike anatomy, the subject matter of this discipline shifts before our eyes. In the years in which I wrote these pages, some of the landmarks taught to students changed twice. The eleventh revision of the WHO classification of diseases, ICD-11, has been in effect since 2022, and the WHO fact sheet of May 2026 counts 132 countries at various stages of implementation. Yet, as I complete this book, Romanian hospitals still group cases by diagnoses coded in ICD-10-AM, third edition. The HL7 FHIR standard is taught in version R4, on which most implementations rely, even though R5 was published in 2023 and R6 is being balloted. The European regulation on artificial intelligence was adopted in 2024, and in July 2026 the timetable of obligations for high-risk systems was shifted by Regulation (EU) 2026/1744. The European Health Data Space has been in force since March 2025, but its obligations apply in stages until 2031. In Romania, the eSănătateaMea portal was launched on September 1, 2026, and the IT platform of the health insurance system was modernized through the National Recovery and Resilience Plan (PNRR).

A first-year student will practice in a system in which all these deadlines will already have passed. It did not seem honest to me to offer such a student an undated picture. I chose instead to write every Romanian and European fact as a fact in motion, with the date on which it was checked and with its official source, so that readers can repeat the check a year or five years from now.

Artificial intelligence takes up more space in this book than in the textbooks of a decade ago, but it does not take center stage. I have tried to present it as a young physician encounters it: as a tool that must be assessed with the same questions as a diagnostic test, that can be convincingly wrong, and for which responsibility remains with the human. For this reason, the calculation of sensitivity and predictive value, the verification of references, and the protection of patient data appear in several chapters, not only in the one devoted to artificial intelligence. A student who masters them will also be able to judge the tools that appear after 2026.

This book does not replace the textbooks of the discipline; it cites them and builds on them. The textbook of medical informatics and biostatistics by Mihalaș and Lungeanu [1] remains the reference for biostatistics and for the foundations that this book does not revisit in detail. Focșa's volume on medical informatics, from data to artificial intelligence [2], and Mihalaș's guide to artificial intelligence in medicine [3] cover many of the topics discussed here from other angles. I recommend reading them in parallel. What the present book adds is a close link with the 2025 curriculum, with each week's practical sessions, and with the situation in Romania at a precise date.

The structure follows the 2025 syllabus of the Discipline of Medical Informatics and Biostatistics of the "Victor Babeș" University of Medicine and Pharmacy of Timișoara: 14 lectures and 14 practical sessions, one pair per week. The same content is taught to students in the Romanian-, English-, and French-language study lines, which is why this English edition keeps the Romanian names of national institutions, laws, and platforms at first mention, alongside their English equivalents. The first chapter defines the discipline and the digital competencies of the physician. Next come data and information theory, databases, classifications and terminologies, and ethics and data protection. Chapters 6 and 7 deal with signals and images, Chapters 8-10 with the electronic health record, telemedicine, and mobile health, and Chapters 11-13 with the

information system, hospital and national systems, and artificial intelligence. The last chapter is a synthesis, with a map of the discipline, a timeline of regulations, and rules for the use of artificial intelligence.

The practical sessions have a longer history than this book. Some protocols come from the discipline's volumes of practical applications, which I coauthored: the 2010 guide to practical sessions [4], the collection of practical applications in medical informatics and biostatistics published the same year [5], and the 2018 volume on biomedical data processing [6]. They have been revised and updated for the free tools available in 2026: Google Forms and Google Sheets, the WHO ICD browsers, Fiji, OpenEMR, test FHIR servers, and public signal databases. Others are new, because their topics were not part of the curriculum a decade ago. The simulated consultation with a language model and the critical evaluation of an artificial intelligence tool are two examples.

Every lecture chapter is built the same way. It opens with operational objectives, phrased as actions that students should be able to perform by the end. The content moves from concepts to methods and, where relevant, to the situation in Romania, presented in a separate box. A box of common pitfalls follows, that is, the mistakes I have seen repeated in seminars and examinations, and then a box of key points. The chapter closes with five multiple-choice questions in the format of the residency entrance examination, with the answers given in Appendix A. Each practical session includes an objective, the tools with their version and access date, the working steps, what to submit with an indicative grading scheme, and a variant for rooms without a network connection.

I have held to a few working rules, which I also ask of my students. Every reference was actually located in PubMed, in CrossRef, or on the website of the issuing institution before being cited, and numbers are reproduced as they appear in the source. Charts based on real data name their source. Those constructed for teaching purposes are labeled as synthetic data. The patients in the examples are fictitious.

The information on Romanian and European legislation, national platforms, and versions of standards reflects the situation as of September 27, 2026. Some of it will change before the book reaches its readers. Where an official source could not be found, I preferred to say so rather than fill the gap with an assumption.

How to Read This Book

The book can be read in order, week by week, but this is not required. The chapters refer to one another, and the map of the discipline in Chapter 14 shows how they are connected. Readers who need a topic quickly (for example, coding a diagnosis or calculating the positive predictive value) can go straight to the relevant chapter and its paired practical session.

For regular study I recommend three steps. First, read the objectives, so that you know what you should be able to do by the end. Next, work through the text together with the figures and tables, which are not illustrations but part of the content. Finally, answer the questions without looking at the answers in Appendix A, and check them only afterward. The common pitfalls are worth reading twice, once before the practical session and once before the examination.

The practical sessions are done at a computer. They require a Google account and a few free programs, but each also has an offline variant. The files submitted become the material for the practical examination, whose complete model is given in Practical Session 14.

Citations are numbered in Vancouver style, in order of first appearance, and the complete list is at the end of the book. Abbreviations are collected in a separate list, and the main terms are briefly defined in the glossary in Appendix B, with a reference to the chapter in which they are explained.

A final recommendation concerns data. Every legal deadline, every version of a standard, and every number about Romanian platforms in this book must be checked at the source before being used in a paper or in a decision. The book always indicates the source. I have written nothing about real patients, and readers are asked to follow the same rule in their own exercises, including when they use an artificial intelligence tool.

Timișoara, September 2026 Călin Muntean

List of Abbreviations

List of abbreviations used in this book, in alphabetical order. For Romanian institutions, documents, and platforms, the Romanian name is given first, followed by an English translation in parentheses.

Abbreviation	Full form
ACHI	Australian Classification of Health Interventions
ACID	atomicity, consistency, isolation, durability (properties of transactions)
ACOM	content model for health observations defined in IEEE 11073-10206
ACR	American College of Radiology
ADC	analog-to-digital converter (Romanian: CAN)
ADR	Autoritatea pentru Digitalizarea României (Romanian Authority for Digitalization)
ADT	patient administration module: admission, discharge, transfer
AES	Advanced Encryption Standard, a symmetric encryption standard
AF	atrial fibrillation
AI	artificial intelligence
AI Act	Regulation (EU) 2024/1689 on artificial intelligence
ALOS	average length of stay
AMIA	American Medical Informatics Association
ANMDMR	Agenția Națională a Medicamentului și a Dispozitivelor Medicale din România (National Agency for Medicines and Medical Devices of Romania)
ANSPDCP	Autoritatea Națională de Supraveghere a Prelucrării Datelor cu Caracter Personal (Romanian National Supervisory Authority for Personal Data Processing)
API	application programming interface
ASCII	American Standard Code for Information Interchange
ASRO	Asociația de Standardizare din România (Romanian Standards Association)
ATC	Anatomical Therapeutic Chemical classification of medicines
ATI	anestezie și terapie intensivă (anesthesia and intensive care; the ICU department in Romanian hospitals)
AUC	area under the ROC curve
BI-RADS	Breast Imaging Reporting and Data System
BMI	biomedical informatics
BP	blood pressure
BPMN	Business Process Model and Notation
CA	certification authority
CAN	convertor analog-numeric, Romanian abbreviation for the analog-to-digital converter; see ADC
CAS	casa de asigurări de sănătate (county health insurance house)
CDA	Clinical Document Architecture
CDC	Centers for Disease Control and Prevention, US
CDS	clinical decision support
CE	European conformity marking
CEAS	cardul electronic de asigurări de sănătate (electronic health insurance card)
CEI	cartea electronică de identitate (electronic identity card)
CEN	European Committee for Standardization (CEN/TC 251: technical committee for health informatics)
CGM	continuous glucose monitoring
CI	confidence interval
CIM	Clasificarea internațională a maladiilor, Romanian name of the International Classification of

Abbreviation	Full form
	Diseases; see ICD
CLAIM	reporting guideline for studies of AI in medical imaging
CMI	case-mix index
CNAS	Casa Națională de Asigurări de Sănătate (National Health Insurance House)
CNP	cod numeric personal (personal numeric code, the Romanian national identification number)
CONSORT-AI	reporting guideline for randomized trials of AI interventions
COPD	chronic obstructive pulmonary disease
COPE	Committee on Publication Ethics
CPOE	computerized provider order entry
CR	computed radiography (phosphor plates read after exposure)
CRF	case report form
CSIRT	Computer Security Incident Response Team
CSV	comma-separated values
CT	computed tomography
CVA	cerebrovascular accident (stroke)
DBMS	database management system
DBT	digital breast tomosynthesis
DDD	defined daily dose
DECIDE-AI	reporting guideline for the early-stage clinical evaluation of decision support systems driven by AI
DECODE	Digital Health Competencies in Medical Education framework
DES	Dosarul electronic de sănătate (the Romanian national electronic health record of the patient)
DICOM	Digital Imaging and Communications in Medicine
DiGA	digital health applications reimbursed in Germany (Ger. *Digitale Gesundheitsanwendungen*)
DigComp	European Digital Competence Framework for Citizens
DIKW	data, information, knowledge, wisdom
DNA	deoxyribonucleic acid
DNSC	Directoratul Național de Securitate Cibernetică (Romanian National Cyber Security Directorate)
DOI	Digital Object Identifier
DPIA	data protection impact assessment
DPO	data protection officer
DR	direct digital radiography
DRG	diagnosis-related groups
E2EE	end-to-end encryption
EBM	evidence-based medicine
ECG	electrocardiogram
EDC	electronic data capture
EDF, EDF+	European Data Format for signals
EEG	electroencephalogram
EEHRxF	European Electronic Health Record Exchange Format
EFMI	European Federation for Medical Informatics
EHDEN	European Health Data & Evidence Network
EHDS	European Health Data Space, Regulation (EU) 2025/327
EHR	electronic health record
eIDAS	Regulation (EU) No 910/2014 on electronic identification and trust services (eIDAS 2: Regulation

Abbreviation	Full form
	(EU) 2024/1183)
EMG	electromyogram
EMR	electronic medical record
EMRAM	Electronic Medical Record Adoption Model
ENISA	European Union Agency for Cybersecurity
ERAS	Enhanced Recovery After Surgery
ERP	enterprise resource planning
ETL	extract, transform, load
EU	European Union
EUDI	European Digital Identity Wallet
FAIR	Findable, Accessible, Interoperable, Reusable
FDA	US Food and Drug Administration
FFT	fast Fourier transform
FHIR	Fast Healthcare Interoperability Resources, an HL7 standard
FIPS	US Federal Information Processing Standards
FN	false negative
FOCG	foaia de observație clinică generală (general clinical observation chart, the Romanian inpatient medical record)
FP	false positive
GDPR	General Data Protection Regulation, Regulation (EU) 2016/679
GHS, GHSP	Bluetooth Generic Health Sensor Profile
GRADE	Grading of Recommendations Assessment, Development and Evaluation (system for rating the certainty of evidence and the strength of recommendations)
HbA1c	glycated hemoglobin
HG	Hotărârea Guvernului (Government Decision)
HIS	hospital information system
HL7	Health Level Seven, the organization and the family of standards (HL7 v2: the messaging standard)
HR	heart rate
HRV	heart rate variability
HTN	hypertension
HTTPS	HTTP secured by TLS
HU	Hounsfield units
IBSI	Image Biomarker Standardisation Initiative
ICD	International Classification of Diseases; ICD-10, ICD-11
ICD-10-AM	Australian modification of ICD-10, used in Romania in its third edition
ICD-10-CM	clinical modification of ICD-10 used in the US
ICF	International Classification of Functioning, Disability and Health
ICHI	International Classification of Health Interventions
ICMJE	International Committee of Medical Journal Editors
ICT	information and communication technology
IDMP	ISO standards for the Identification of Medicinal Products
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
IHE	Integrating the Healthcare Enterprise
IMIA	International Medical Informatics Association

Abbreviation	Full form
INMSS	Institutul Național de Management al Serviciilor de Sănătate (National Institute for Health Services Management)
IOD	DICOM Information Object Definition
I-PASS	patient handover protocol: illness severity, patient summary, action list, situation awareness and contingency planning, synthesis by receiver
IPS	International Patient Summary
ISMS	information security management system
ISO	International Organization for Standardization (ISO/TC 215: technical committee for health informatics)
ITIL	IT Infrastructure Library, a set of good practices for IT service management
IVDR	Regulation (EU) 2017/746 on in vitro diagnostic medical devices
JSON	JavaScript Object Notation
LIS	laboratory information system
LLM	large language model
LOINC	Logical Observation Identifiers Names and Codes
MA	multiple answer (type of multiple-choice question)
MARS	Mobile App Rating Scale
MDC	major diagnostic category
MDCG	Medical Device Coordination Group
MDR	Regulation (EU) 2017/745 on medical devices
MeSH	Medical Subject Headings, the NLM controlled vocabulary
MFA	multi-factor authentication
ML	machine learning
MMS	ICD-11 linearization for Mortality and Morbidity Statistics
MPPS	Modality Performed Procedure Step
MR, NMR	magnetic resonance, nuclear magnetic resonance
MRI	magnetic resonance imaging
MS	Ministerul Sănătății (Romanian Ministry of Health)
NAM	US National Academy of Medicine
NIS2	Directive (EU) 2022/2555 on the security of network and information systems
NIST	National Institute of Standards and Technology, US
NLM	US National Library of Medicine
NNT	number needed to treat
NoSQL	non-relational databases (*not only SQL*)
NPV	negative predictive value
OHDSI	Observational Health Data Sciences and Informatics
OMOP CDM	Observational Medical Outcomes Partnership Common Data Model
OR	odds ratio
OUG	ordonanță de urgență a Guvernului (Government Emergency Ordinance)
PACS	Picture Archiving and Communication System
PDMS	patient data management system in intensive care
PET	positron emission tomography
PGHD	patient-generated health data
PHR	personal health record
PIAS	Platforma informatică din asigurările de sănătate (health insurance IT platform)

Abbreviation	Full form
PICO	Population, Intervention, Comparator, Outcome
PKI	public key infrastructure
PMID	PubMed identifier of an article
PNRR	Planul Național de Redresare și Reziliență (National Recovery and Resilience Plan)
PPG	photoplethysmography
PPV	positive predictive value
PS	practical session (Romanian: LP, lucrare practică)
RAG	retrieval-augmented generation
RAM	random access memory, the working memory of a computer
RBAC	role-based access control
RCT	randomized controlled trial
REST	representational state transfer, the web interface style used by FHIR and DICOMweb
RIM	HL7 Reference Information Model
RIS	Radiology Information System
RMSSD	root mean square of successive differences between NN intervals
ROC	receiver operating characteristic curve
ROI	region of interest
RR	interval between two successive R waves; in the context of studies, relative risk
RSA	Rivest-Shamir-Adleman asymmetric encryption algorithm
SA	single answer (type of multiple-choice question)
SBAR	Situation, Background, Assessment, Recommendation (communication protocol)
SDNN	standard deviation of NN intervals
SHAP	SHapley Additive exPlanations
SIEC-ATI	sistemul informatic pentru evidența clinică a secțiilor ATI (information system for the clinical records of Romanian ICU departments)
SIPE	Sistemul informatic de prescripție electronică (Romanian electronic prescription system)
SIUI	Sistemul informatic unic integrat al asigurărilor de sănătate (Single Integrated Information System of the Romanian health insurance system)
SMART	authorization framework for applications running on FHIR data (SMART App Launch, SMART on FHIR)
SMD	standardized mean difference
SNOMED CT	Systematized Nomenclature of Medicine Clinical Terms
SNR	signal-to-noise ratio
SOP	DICOM Service-Object Pair
SpO2	peripheral oxygen saturation, measured by pulse oximetry
SQL	Structured Query Language
SRIM	Societatea Română de Informatică Medicală (Romanian Society of Medical Informatics)
STARD-AI	reporting guideline for diagnostic accuracy studies of AI tests
SWF	IHE Scheduled Workflow profile for radiology
TLS	Transport Layer Security
TN	true negative
TP	true positive
TRIPOD+AI	reporting guideline for the development and validation of prediction models, including those built with machine learning
UID	unique identifier of a DICOM object, unique worldwide

Abbreviation	Full form
UMLS	Unified Medical Language System
US	United States of America
USMLE	United States Medical Licensing Examination
UTF-8	Unicode Transformation Format, 8-bit: the Unicode encoding form using 1-4 bytes
VNA	vendor neutral archive
WFDB	Waveform Database format and software
WHA	World Health Assembly
WHO	World Health Organization
WSI	whole slide imaging in digital pathology
XAI	explainable AI
XML	eXtensible Markup Language

PART I

Part I. Lectures

CHAPTER 1

Chapter 1. Medical Informatics: Subject, Domains, Competencies

Objectives

After studying this chapter, the student will be able to:

1. define biomedical informatics, medical informatics, health informatics, and clinical informatics, and explain how they differ.
2. distinguish, using a clinical example, a datum from a piece of information and from a piece of knowledge.
3. arrange at least eight milestones of the discipline in chronological order, from punched cards to large language models.
4. describe the domains of biomedical informatics along the spectrum of molecules, tissues, patient, and population.
5. list the domains of digital competence required of physicians and estimate his or her own level.
6. build a PubMed query with MeSH terms and Boolean operators and verify a reference proposed by an artificial intelligence tool.

1.1 Definitions and Subject

1.1.1 From "Medical Computing" to Biomedical Informatics

The name of the discipline has changed several times over the past sixty years, and each change says something about the way its subject was understood. In the 1970s, the English-speaking world spoke of *medical computer science*, that is, the subdivision of computer science that applies the general methods of the field to medical problems. Also in the 1970s, modeled on the French *informatique*, the term *medical informatics* appeared, and it gained ground mainly in Europe in the 1980s [7]. The new term shifted the emphasis from the computer to information. It encompassed medical statistics, record keeping, and the study of the nature of medical information, not just programming [7].

The adjective "medical" was later criticized as being too closely tied to the physician and to disease. This gave rise to *health informatics* and, with the Human Genome Project of the 1990s, to *bioinformatics*. To encompass both directions, part of the community adopted the name biomedical informatics (BMI). One sign of this change is the journal *Computers and Biomedical Research*, launched in 1967, which became the *Journal of Biomedical Informatics* in 2001 [7].

The definition adopted by the American Medical Informatics Association (AMIA) and taken up by the reference textbook of the field describes biomedical informatics as the interdisciplinary field that studies and pursues the effective use of biomedical data, information, and knowledge for scientific research, problem solving, and decision making, with the aim of improving human

health [7]. The definition has four corollaries. The first sets the scope of the field, from molecules to individuals and populations, and from biological systems to social systems. The second concerns the theory and methods for generating, storing, retrieving, using, and sharing data and knowledge. The third states that the discipline builds on computer science, telecommunications, and the information sciences and contributes to them. The fourth recalls that people are the end users, which is why the design and evaluation of solutions also draw on the social and behavioral sciences [7].

1.1.2 Related Terms

The terms in Table 1.1 appear frequently in the literature, and they are not synonyms. Shortliffe and Chiang propose a simple convention: biomedical informatics is the core discipline, and the other names designate its fields of application [7]. Under this convention, "medical informatics" denotes applied research and practice focused on disease and on the role of the physician, while "health informatics" covers clinical informatics and public health informatics taken together.

Table 1.1. Terms used for the discipline and its fields. Source: [7-9].

Term	Romanian term	What it covers
Biomedical informatics	informatică biomedicală	the core discipline: methods, theories, tools
Medical informatics	informatică medicală	applications focused on disease and on the physician's work
Health informatics	informatica sănătății	clinical and public health informatics together
Clinical informatics	informatică clinică	applications for the care of the individual patient
Bioinformatics	bioinformatică	biological information, especially at the molecular level
Public health informatics	informatica sănătății publice	applications for populations, prevention, surveillance

In Romanian medical schools, the discipline taught to students is called medical informatics. The first course, introduced as an elective in 1984, became compulsory for all medical students in 1990 [10]. In this book, "medical informatics" refers to the discipline as the medical student encounters it, and the term "biomedical informatics" is used where the underlying science is meant.

1.1.3 Data, Information, Knowledge

A clinical datum is a single observation about a patient: a temperature, a red blood cell count, a history of rubella. Shortliffe and Chiang show that a datum is defined by five elements: the patient, the parameter observed, the value, the time of the observation, and the method by which it was obtained [11]. Information consists of data that have been analyzed and organized so that they acquire meaning. Knowledge results from the formal or informal analysis of information and includes study results as well as heuristics and models [11].

An example helps. A blood pressure of 180/110 mmHg measured in a patient is a datum. When such data are brought together for many patients, information emerges, and its analysis may show that patients with high blood pressure have myocardial infarctions more often, which is a piece of knowledge [11]. The boundary depends on context. "The patient has hypertension" is an

interpretation of several measurements, but it becomes a simple input datum for a decision support program. A database stores individual observations, without analysis. A knowledge base contains facts, heuristics, and models used for problem solving [11]. Chapter 2 returns to these notions together with the measurement of information.

1.1.4 Scientific Discipline and Professional Practice

Biomedical informatics can be seen as a basic biomedical science with numerous fields of application. It is an experimental science: it formulates questions, designs experiments, analyzes the results, and uses what it has learned for new experiments [7]. Basic and applied research build on each other. Problems originate in medical practice, and generalizable solutions return to it. The discipline draws on several component sciences, not only on computer science: the decision sciences, statistics, the cognitive sciences, the information sciences, and the management sciences [7].

The relationship with biomedical engineering deserves clarification. Biomedical engineering departments emerged in the late 1960s and dealt mainly with instrumentation, monitoring, and imaging techniques. Biomedical informatics, a younger field, places biomedical information and knowledge, and their management with the aid of computers, at its center [7]. The two often meet, for example in signal and image processing (Chapters 6 and 7).

Informatics is also a profession. In the United States, clinical informatics was approved as a medical subspecialty in 2013, with a certification examination [7]. For physicians who will not become informatics specialists, the stakes are different: to use digital tools correctly, to recognize their limits, and to take part in selecting and evaluating them. The same source notes that the uptake of an innovation into routine practice can take 17 years on average [7]. That is precisely why it matters that physicians be able to judge the evidence.

1.2 Historical Milestones: From the First Hospital Computers to Generative Artificial Intelligence

The history of medical informatics is also the history of computers. The first application of automated computing relevant to medicine was Herman Hollerith's punched-card system, used in the 1890 US census and later adapted for epidemiological surveys [7]. Stored-program electronic computers appeared in the late 1940s, and general-purpose computers were sold commercially from the mid-1950s onward [7]. Artificial intelligence (AI) took shape at the same time. The 1956 Dartmouth conference marked its establishment as a formal discipline [12].

In the 1960s, several American centers built the first hospital information systems. At Massachusetts General Hospital, Barnett and colleagues developed applications over three decades, and the MUMPS language, created specifically for medical applications, was for several years the most widely used language for processing medical records [7]. Scientific documentation was computerized in parallel. The printed Index Medicus remained the main source for bibliographic retrieval until 1971, when the US National Library of Medicine made available the electronic version, MEDLARS, with records cataloged since 1966, followed by MEDLINE (*MEDLARS Online*) [13].

The 1970s brought the first decision support systems. In 1972, de Dombal's group used a Bayesian model for the diagnosis of acute abdominal pain. MYCIN, a rule-based system for infections, was never used in the clinic, but it served as a prototype for many later systems [14]. The minicomputers of the early 1970s, and then the microprocessor, followed by the personal

computer, in the late 1970s and early 1980s, brought computing into every department and then onto every desk [7].

As systems multiplied, so did the need for them to communicate. The HL7 organization was founded in 1987, and the 1990s saw the creation of the European standardization committee for health informatics and of the technical committee ISO/TC 215 [15]. The HL7 FHIR standard, introduced in 2011, simplified data exchange through resources described in XML or JSON [15]. In 2009, the US HITECH Act offered financial incentives to hospitals and physicians who adopted certified electronic health records and actually used them [7].

In the early 2000s, deep learning overcame many of the limitations of early AI models [16]. In 2021, AlphaFold predicted the three-dimensional structure of proteins with an accuracy comparable, in most of the cases tested, to that of experimentally determined structures. The problem had been open for more than 50 years [17]. Since late 2022, ChatGPT has made large language models (LLMs) accessible to anyone [18]. In 2023, ChatGPT achieved scores at or near the passing threshold on all three steps of the United States Medical Licensing Examination (USMLE) [19]. In the same year, the Flan-PaLM model reached 67.6% accuracy on USMLE-style questions, and its version adapted to the medical domain, Med-PaLM, produced encouraging results but remained inferior to clinicians in human evaluation [20].

The COVID-19 pandemic accelerated telemedicine, data exchange, and the use of AI. A 2023 analysis concludes that digital health matured during the pandemic, but that sociotechnical problems, as well as problems of evaluation and trust, remained [21]. The European Union responded with regulation: the Artificial Intelligence Regulation was adopted in 2024 [22], and the Regulation on the European Health Data Space in 2025 [23]. Figure 1.1 summarizes these milestones.

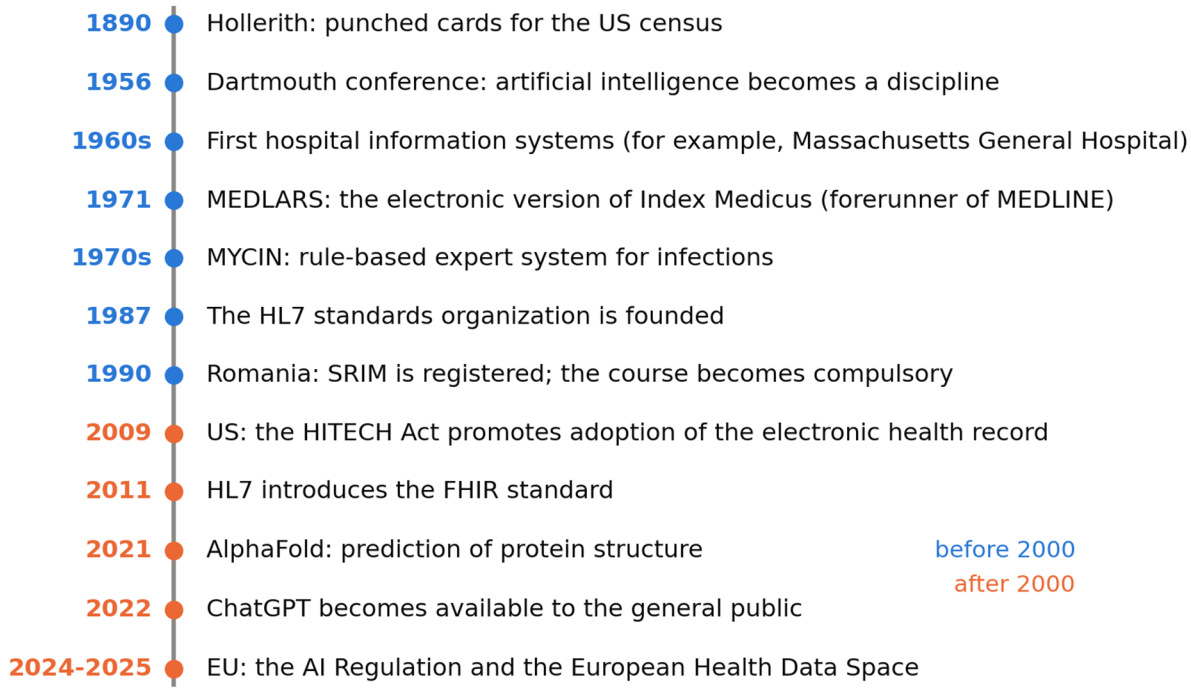


Figure 1.1. Milestones in the history of medical informatics. The axis is not proportional to time. Sources: [7,10,12-15,17,18,22,23].

1.3 Domains of Biomedical Informatics: From Molecules to Populations

The same methods (data representation, standards, reasoning, evaluation) apply at different levels of the organization of life. Shortliffe and Chiang describe a four-step spectrum: molecular and cellular processes, tissues and organs, the individual (the patient), and populations and society [7]. Each step corresponds to a field of application, as shown in Figure 1.2. The boundaries are blurred. Pharmacogenomics, for example, links genotypic and phenotypic databases and lies at the intersection of bioinformatics and clinical informatics [7].

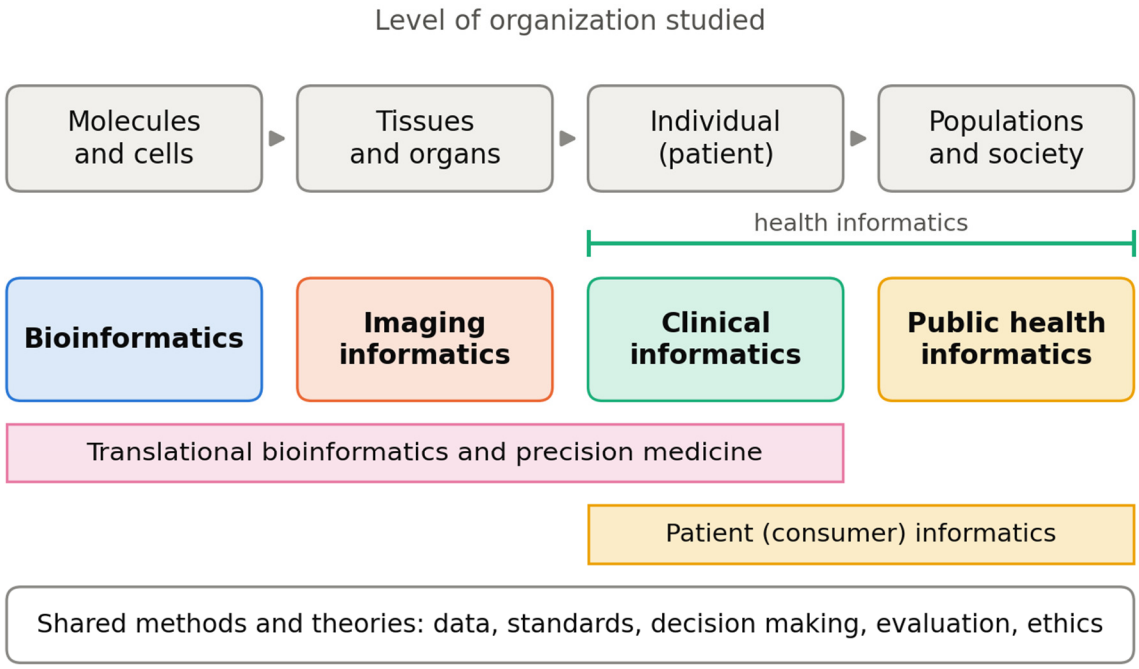


Figure 1.2. The domains of biomedical informatics along the spectrum from molecules to populations. Adapted from [7].

1.3.1 Bioinformatics

Bioinformatics studies how information is represented and analyzed in biological systems, especially at the molecular level. Whereas clinical informatics manages information related to patient care, bioinformatics deals with information from the basic biological sciences [8]. Sequencing, gene expression, and metabolomic data are now produced at a pace that only computational methods can keep up with. AlphaFold is an example of a bioinformatics result with a direct impact on medical research [17].

1.3.2 Imaging Informatics

Medical images are either acquired directly in digital format or converted into it. Imaging informatics deals with image acquisition, the representation of image content, storage and management, processing, and interpretation, possibly with computer assistance [24]. The number of images produced in a single examination can reach hundreds or thousands [24]. Chapter 7 covers the DICOM standard and image processing in detail.

1.3.3 Clinical Informatics

Clinical informatics covers applications oriented toward the individual patient, in medicine, nursing, or dentistry [7]. The electronic health record (EHR), e-prescription, clinical decision support systems, and telemedicine belong to this domain. They are the subject of Chapters 8-12.

1.3.4 Public Health Informatics

Public health informatics is the systematic application of informatics methods and tools in support of public health goals, regardless of the setting in which the work is done [9]. Public health has three core functions: assessment of the health status of the population, policy development, and assurance of services [9]. Each requires different data and different systems. The pandemic showed how closely electronic health records are linked to public health informatics in the management of an epidemic [7]. In the same context, the concept of the learning health system is discussed, defined as a system that continuously uses a data-driven infrastructure for policy, public health, and the personalization of care [25].

1.3.5 Patient (Consumer) Informatics

Patient-centered care has become a core component of modern medicine, and patient informatics (*personal* or *consumer health informatics*) studies the tools through which people manage their own health: patient portals, personal health records, mobile applications, and home sensors [26]. Health literacy and the digital divide also belong here [26].

1.3.6 Translational Informatics

Translational bioinformatics applies the methods of bioinformatics to human health and disease. It helps laboratory discoveries reach the bedside and, ultimately, population health, through precision medicine [27]. Unlike "classical" bioinformatics, it works mainly with human data, which is why it faces regulatory and confidentiality issues that do not arise with data from model organisms [27].

Data protection cuts across all these domains. A bibliometric analysis found 7,165 articles on data privacy in medical informatics published between 1975 and 2023, with an annual number that kept increasing. The most studied topics were mobile health and telemedicine, security and authentication, and cloud computing [28]. Chapter 5 deals with these issues.

1.4 Digital Competencies of the Physician

What does a physician need to know about informatics? The answer has changed as digital tools have entered every stage of medical care. Several organizations have published competency frameworks, and these overlap to a large extent.

1.4.1 IMIA Recommendations

The International Medical Informatics Association (IMIA) published recommendations on education in biomedical and health informatics in 2000 and 2010, and the second revision appeared in 2023 [29]. The revision was prepared by a working group with representatives from all IMIA regions, on the basis of an international Delphi study. It describes the educational needs of three roles (user, generalist, and specialist) in six domains: the core principles of the discipline; the health sciences and health services; computer, data, and information sciences; the social and behavioral sciences; the management sciences; and, finally, informatics specialization [29]. For medical students, the relevant role is first and foremost that of the competent user.

1.4.2 Frameworks for Medical Education

The DECODE framework (*Digital Health Competencies in Medical Education*), published in 2025, resulted from a Delphi study with 211 experts from 79 countries and territories. It groups competencies into four domains: professionalism in digital health, patient and population digital health, health information systems, and health data science. The 19 competencies have 33 mandatory and 145 optional learning outcomes [30]. An earlier study, based on a literature review and a Delphi panel of 18 experts, selected 40 of 65 digital health topics as necessary in the curriculum for medical students [31].

Students' attitudes are generally favorable. A 2026 meta-analysis of 96 cross-sectional studies from 37 countries, with more than 45,000 students, estimated a positive attitude toward AI in 76.9% of students and 76.6% support for including it in the curriculum. At the same time, 39.9% feared that physicians could be replaced, and the variation between centers was very large [32].

1.4.3 The European Framework and the WHO Strategy

The European Digital Competence Framework for Citizens, DigComp, was published in version 2.2 in 2022 [33]. Version 3.0, published on November 27, 2025, keeps the five areas and 21 competences but integrates AI-related competences into all areas [34]. The five areas are information search, evaluation, and management; communication and collaboration; content creation; safety (together with well-being and responsible use); and, last, problem identification and solving [34].

In 2025, in Policy Brief 71, the European Observatory on Health Systems and Policies analyzed the digital competencies of health professionals in the EU. The competency frameworks analyzed emphasize information and data literacy, communication and collaboration, safety and security, problem solving, and digital content creation [35]. The available training is considered insufficient, with gaps in advanced areas (robotics, AI, genomics) and an excessive reliance on courses in English. Participants from Italy and Romania in the focus group discussions explicitly reported that a lack of English language skills limits their access to courses [35]. The document also presents an example of regulation. France adopted a common digital competency framework for the health professions, with five domains, and a law required its inclusion in the higher education training of 12 professions by 2024-2025 [35].

At the global level, the WHO Global Strategy on Digital Health 2020-2025 was endorsed by the World Health Assembly in 2020. Among the actions proposed to Member States is the inclusion of digital health competencies in the training programs of all health professionals [36]. In May 2025, the World Health Assembly approved the extension of the strategy until 2027 [37,38]. At that time, 129 countries had a national digital health strategy [37].

Table 1.2 translates these frameworks into concrete tasks for the first year and indicates where they are practiced in this book.

Table 1.2. Digital competence areas relevant in the first year. Source: adapted from [30,34,35].

Area (DigComp)	What it means for the medical student	Where it is practiced
Information and data	PubMed searching with MeSH, evaluation of sources, reference management	Ch. 1, 2; PS 1, PS 2
Communication and collaboration	controlled sharing, teamwork on documents, teleconsultation	PS 1, Ch. 9
Content creation	documents, tables, charts, correct citation	PS 3, PS 5
Safety and responsible use	passwords, two-step verification, GDPR, AI ethics	PS 1, Ch. 5, 11
Problem solving	choosing the right tool, verifying AI output	Ch. 10, 13

1.5 Searching the Scientific Literature

1.5.1 PubMed, MEDLINE, and MeSH

PubMed is the free search engine of the US National Library of Medicine and contains more than 40 million citations from MEDLINE, life science journals, and online books [39]. On September 27, 2026, a query of the entire database returned 41,213,139 records [39]. MEDLINE is the component indexed with the controlled vocabulary MeSH (Medical Subject Headings). In the edition described by the cited textbook, MeSH contained more than 28,000 descriptors and more than 90,000 synonyms (entry terms), organized hierarchically in 16 trees, such as Diseases, Organisms, or Chemicals and Drugs [13]. The hierarchy makes it possible to "explode" a term, that is, to include the more specific terms automatically. Subheadings narrow the meaning, for example to diagnosis or therapy, and the publication type makes it possible to select reviews, randomized trials, or meta-analyses [13].

Indexing has gradually been automated. The NLM automatic indexing tool evolved from a rule-based system in 2002 to MTIX, based on neural networks, in 2024. A 2026 review shows that about one third of records still require human correction, which is why searchers should watch out for algorithmic ambiguities [40].

Figure 1.3 uses PubMed as a data source. The annual number of records indexed with the descriptor "Medical Informatics" (which, through explosion, includes all subordinate terms) increased until 2019, then fell to a low in 2022-2023 and partially recovered in 2024-2025. The number of records indexed with "Artificial Intelligence" or "Machine Learning" increased about 14-fold between 2005 and 2025, from 3,579 to 51,707. The decline in the medical informatics series should not automatically be read as a decline in interest. It may also reflect changes in indexing, and recent records may receive their MeSH descriptors with a delay. This is a hypothesis, not a result verified in this chapter.

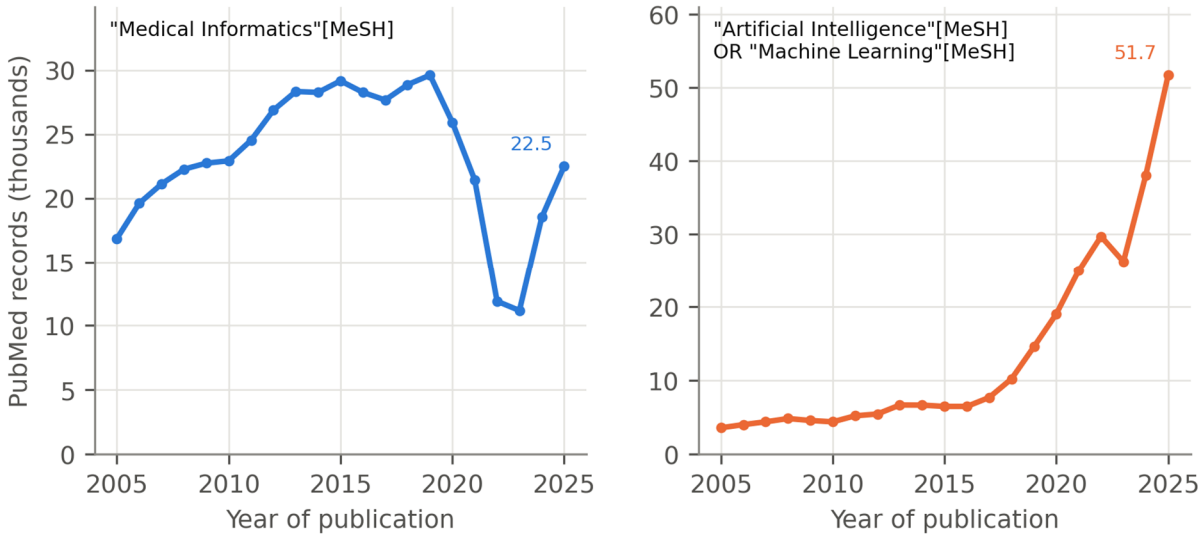


Figure 1.3. Annual number of PubMed records indexed with the descriptor "Medical Informatics" (left) and with "Artificial Intelligence" or "Machine Learning" (right), 2005-2025. The vertical axes have different scales. Source: PubMed (NCBI), query run on September 27, 2026, via E-utilities; author's calculations [39]. The data are in note/date_pubmed_C01.csv.

1.5.2 Boolean Operators, Fields, and Filters

A PubMed search combines sets of results by means of Boolean operators. AND restricts the results to documents that contain all the concepts. OR brings together the different ways of expressing the same concept, and NOT removes one set from another [13]. Operators are written in capital letters, the search is processed from left to right, and parentheses group concepts [41]. Field tags specify where to search: [mh] for MeSH terms, [majr] for major topic, [tiab] for title and abstract, [pt] for publication type, and [dp] for date of publication. The asterisk truncates a term (at least four characters must precede it), and quotation marks search for an exact phrase [41].

When the user types a few words in the search box, PubMed tries to match them successively with MeSH terms, journal titles, common phrases, and authors. Whatever remains unmatched is searched as free text [13]. The filters on the left side of the page limit the results by article type, publication date, species, language, or age, and the Clinical Queries feature applies validated strategies for questions of therapy, diagnosis, harm, and prognosis [13,41]. The default sort order, Best Match, is based on a relevance algorithm [41]. In Practical Session 2, these steps are practiced on a clinical question.

1.5.3 The Evidence Pyramid and Secondary Sources

Not all articles carry the same weight. The classic hierarchy of evidence starts from animal studies, translational research, and expert opinion, rises through case reports and case series, then through analytical observational studies, reaches randomized controlled trials, and has systematic reviews and meta-analyses at the top [42]. The hierarchy helps prioritize reading, but it does not replace the critical appraisal of each article [42]. Over time, evidence-based practice has shifted from the critical appraisal of individual studies toward the use of evidence that has already been synthesized, that is, systematic reviews, guidelines, and syntheses integrated into clinical systems [13]. For first-year students, the practical rule is simple: look first for a recent synthesis and a guideline, and only then for primary studies.

1.5.4 AI-Assisted Search Tools and Their Limits

General-purpose chatbots produce fluent texts with references that look credible. Several studies have checked these references one by one. In a 2023 study, ChatGPT-3.5 generated 30 short medical texts with 115 references. Of these, 47% were fabricated, 46% existed but contained errors, and only 7% were authentic and accurate. An incorrect PMID was the most frequent error [43]. When asked to reproduce 11 systematic reviews, GPT-3.5, GPT-4, and Bard had hallucination rates of 39.6%, 28.6%, and 91.4%, respectively, with a precision of article selection of at most 13.4% [44]. In an interdisciplinary study with ChatGPT (GPT-3.5), about three quarters of the citations existed, but the DOI was correct in only 32.7% of cases in the natural sciences and in 8.5% in the humanities [18].

Technical solutions reduce the problem but do not eliminate it. OpenScholar, a specialized model that first retrieves passages from real articles and only then drafts the answer (retrieval-augmented generation, RAG), achieved a citation accuracy comparable to that of experts, whereas GPT-4o fabricated citations in 78-90% of cases [45]. In a randomized study with 30 urologists, a chatbot with RAG and in-text references gave correct recommendations in 73% of cases, compared with 50% for ChatGPT, and 28% of the ChatGPT citations were nonexistent or outdated [46]. A systematic review of 519 studies evaluating LLMs in healthcare found that only 5% had used real patient care data [47].

The practical conclusion for students is this: an AI tool can help formulate the question or find synonyms, but every reference must be checked in PubMed or through its DOI before it is cited. The same principle was applied, incidentally, in writing this book.

Table 1.3. Accuracy of references produced by language models in studies published in 2023-2026. Source: [43-46].

Study	Task	Main result
Bhattacharyya 2023	115 references in medical texts (GPT-3.5)	47% fabricated, 7% accurate
Chelli 2024	reproducing 11 systematic reviews	hallucination 28.6-91.4%
Asai 2026	literature synthesis, without and with RAG	GPT-4o: 78-90% fabricated citations
Carl 2026	uro-oncology cases, RCT with 30 urologists	28% of ChatGPT citations nonexistent or outdated

1.6 Medical Informatics in Romania and in Timișoara

The first Romanian event dedicated to medical informatics took place in 1977. The Romanian Society of Medical Informatics (Societatea Română de Informatică Medicală, SRIM) could be officially registered only in 1990, after the fall of the communist regime [10]. Romania was a corresponding member of IMIA from 1986 and a full member from 1994, the year in which SRIM also became a member of the European Federation for Medical Informatics (EFMI) [48]. Two SRIM representatives later led EFMI: George Mihalaș (2006-2008) and Lăcrămioara Stoicu-Tivadar (2018-2020). SRIM organized EFMI special topic conferences in Bucharest, in 2001, and in Timișoara [10]. The society has its headquarters in Timișoara [49]. In May 2026, SRIM held the second edition of the jRoMedINF conference in Bucharest, on the theme of the life cycle of medical data [49].

Timișoara also hosted the EFMI special topic conference of November 27-29, 2024, organized by SRIM at the Politehnica University of Timișoara, with the West University of Timișoara and the "Victor Babeș" University of Medicine and Pharmacy as co-organizers. The extended

abstracts of the posters were published in a supplement of the journal *Applied Medical Informatics* [50]. The journal is published at the "Iuliu Haţieganu" University of Medicine and Pharmacy in Cluj-Napoca; it is open access and charges no publication fees [51]. EFMI also mentions the collaboration of SRIM with the Romanian Society of Applied Medical Informatics [10].

In Timişoara, the discipline has a teaching tradition of its own. The textbook of medical informatics and biostatistics by Mihalaş and Lungeanu [1] and Focşa's recent volume on medical informatics, from data to artificial intelligence [2], are core readings for this course.

Standardization is the third component of professional life. In Romania, medical informatics standards are monitored by the technical committee ASRO/CT 319 "Informatică medicală" (Medical Informatics), the national counterpart of the ISO/TC 215 and CEN/TC 251 committees [52]. The author of this book is active in the ASRO technical committee for medical informatics.

Romania in 2026

- The medical informatics course has been compulsory for all medical students since 1990 [10]. At the "Victor Babeş" University of Medicine and Pharmacy of Timişoara, it is taught in the first year, according to the 2025 syllabus.
- SRIM, headquartered in Timişoara, has been a member of IMIA and EFMI since 1994 [48,49].
- The ISO and CEN medical informatics standards are monitored by ASRO/CT 319 [52].
- European documents point to the English-language barrier that limits the access of Romanian professionals to digital training [35].

1.7 How to Use This Book

The book follows the 2025 syllabus of the discipline: 14 lectures and 14 practical sessions, one pair per week. Each lecture chapter ends with multiple-choice questions, and each practical session describes what is to be submitted for the practical examination. Table 1.4 shows the correspondence.

Table 1.4. Correspondence between lectures, practical sessions, and the chapters of this book (2025 syllabus).

Week	Lecture	Practical session	Chapter
1	Subject of medical informatics; literature searching	Hardware and software; Google Workspace	Ch. 1
2	Information theory; data and knowledge	Web sources, quality assessment	Ch. 2
3	Files, DBMS, portability, NoSQL	Google Forms, Epi Info	Ch. 3
4	ICD-10, ATC, DRG in Romania	Online ICD coding	Ch. 4
5	Ethics, GDPR, data security	Citation, plagiarism, essay	Ch. 5
6	Biosignals	ECG/EEG, free software	Ch. 6
7	Biomedical images	ImageJ/Fiji	Ch. 7
8	eHealth I: the electronic health record	OpenEMR, demo version	Ch. 8
9	eHealth II: telemedicine	Simulated teleconsultation	Ch. 9
10	mHealth, interoperability	Consultation with an LLM	Ch. 10
11	Information system; electronic signature	Security, signature	Ch. 11
12	The health information system, SIUI, HIS	Medical practices, polyclinics, hospitals	Ch. 12
13	AI in medicine	AI projects and debates	Ch. 13
14	Review	Practical examination	Ch. 14

The book cites the textbooks of the discipline but does not replace them [1,2]. Students use them in parallel, especially for biostatistics and for the topics covered there in detail. Advanced research methods, such as digital twins, federated learning, or the reporting of prediction models according to TRIPOD+AI, are covered in the author's doctoral monograph.

Common Pitfalls

- Believing that medical informatics means "knowing how to use a computer." Its subject is information and decision making, and the computer is a tool.
- Using "bioinformatics" as a synonym for the entire discipline. Bioinformatics is only the domain that works at the molecular and cellular level.
- Searching PubMed as if it were Google, with a long sentence, without MeSH terms and without operators. The result is either too broad or incomplete.
- Writing "and" instead of AND, or forgetting the parentheses when combining OR with AND.
- Citing a reference proposed by a chatbot without finding it in PubMed or by its DOI. Studies show high rates of fabricated references.
- Confusing the hierarchy of evidence with quality: a weak meta-analysis does not become good because it sits at the top of the pyramid.

Key Points

- Biomedical informatics studies the effective use of data, information, and knowledge for research, decision making, and health.
- A datum is an observation, information is an organized set of data, and knowledge results from the analysis of information.
- The fields of application range from molecules (bioinformatics) to populations (public health informatics).
- The IMIA, DECODE, and DigComp frameworks require physicians to have competencies in information, communication, safety, and problem solving.
- PubMed is searched efficiently with MeSH terms, Boolean operators written in capital letters, parentheses, and filters.
- Every reference produced by an AI tool must be checked at the source.
- In Romania, the discipline has been compulsory since 1990, and SRIM has its headquarters in Timișoara.

Self-Assessment Questions

1. (SA) According to the definition adopted by AMIA, biomedical informatics is:

a) a synonym for bioinformatics
b) the branch of engineering that designs medical equipment
c) the interdisciplinary field that studies the effective use of biomedical data, information, and knowledge for research, problem solving, and decision making, with the aim of improving health
d) the discipline concerned exclusively with the electronic health record
e) a subspecialty of medical statistics

2. (SA) The entry "patient X, blood pressure 180/110 mmHg, time 14:30, automatic blood pressure monitor" represents:

a) a piece of knowledge
b) a clinical datum
c) a knowledge base
d) a heuristic
e) a practice guideline

3. (SA) In PubMed, the operator that keeps only the articles containing both concepts is:

a) OR
b) NOT
c) AND
d) the asterisk (*)
e) quotation marks

4. (MA) Which of the following statements about references generated by large language models are supported by the studies presented in this chapter?

a) in a study of 115 references produced by ChatGPT-3.5, 47% were fabricated
b) all the models tested had hallucination rates below 5%
c) when reproducing systematic reviews, Bard had a hallucination rate of 91.4%
d) a chatbot with retrieval from sources (RAG) and in-text references increased the proportion of correct recommendations in a randomized study
e) the PMID proposed by the model is correct in all cases

5. (MA) Which of the following are competence areas of the European DigComp framework?

a) information search, evaluation, and management
b) communication and collaboration
c) content creation
d) digital clinical pharmacology
e) problem identification and solving

CHAPTER 2

Chapter 2. Data, Information, and Knowledge in Medicine

Objectives

By the end of this chapter, the student will be able to:

- classify a medical datum by structure (structured, semistructured, unstructured) and by measurement scale (nominal, ordinal, numeric)
- calculate the amount of information of an outcome and the entropy of a binary variable, using the base-2 logarithm
- explain, with a numerical example, the relationship between pre-test probability, post-test probability, and the information gained from a diagnostic test
- estimate the order of magnitude of an ECG file, a CT examination, and a sequenced genome
- describe the data-information-knowledge-wisdom hierarchy and at least two of its limitations
- identify the sources of noise in a patient handover and argue for the use of a structured protocol (I-PASS, SBAR).

2.1 Medical Data

2.1.1 What Is a Clinical Datum

In informatics terms, a clinical datum is a single observation about a patient: a temperature, a hemoglobin value, a history of rubella, or a blood pressure reading [11]. For such an observation to be processed by a computer, it must be described by five elements: the patient, the parameter observed, the value, the time of the observation, and the method by which it was obtained (patient report, thermometer, urine dipstick, laboratory analyzer) [11]. The blood pressure example shows why the data model matters. The physician reads "120/80" as a single value, but for analysis it represents two distinct pieces of information, the systolic and the diastolic pressure.

Time greatly complicates data management. In outpatient care, the calendar date of the visit is usually sufficient. In intensive care, the minute matters, and for a vital signs monitor even fractions of a second count.

2.1.2 Types of Data by Structure

Structured data have a fixed, predefined format: a numeric field for blood glucose, a diagnosis code, a calendar date. They can be sorted, counted, and compared directly. **Semistructured** data have tags, but the content of each tag may vary. A laboratory message in JSON or XML format is the typical example (Chapter 3). **Unstructured** data are free text, images, and audio recordings. Narrative text (the medical history, progress notes, the medical letter) represents a large share of the information collected in patient care [11]. It is rich in nuance but hard to query automatically without natural language processing methods.

2.1.3 Types of Data by Measurement Scale

Statistics uses a different classification, based on which operations are meaningful for the values:

- **nominal**: categories without order (blood group, sex, admitting ward)
- **ordinal**: ordered categories, but without equal distances between them (NYHA class, the Glasgow score grouped into categories, pain grade "mild/moderate/severe")
- **numeric**: values measured on a scale with units (blood pressure in mmHg, creatinine in mg/dL), which may be discrete (number of pregnancies) or continuous (weight).

Coding an ordinal variable with the digits 1, 2, 3 does not turn it into a numeric variable. The mean of pain grades, for example, has a questionable meaning, because "moderate" is not necessarily twice as much as "mild."

2.1.4 Special Forms of Data

Besides isolated values, medicine also produces data with a structure of their own:

- **time series**: the same variable measured repeatedly (hourly blood glucose, the temperature on the observation chart)
- **signals**: continuous variations sampled by a device (ECG, EEG, pulse oximetry)
- **images**: radiographs, CT, MRI, scanned pathology images
- **omics data**: genomics, transcriptomics, proteomics, metabolomics. The field started from genomics and expanded to numerous high-throughput measurement technologies, and current trends are single-cell analysis and the integration of several omics layers in the same analysis [53]
- **patient-generated health data** (PGHD): clinically relevant data captured by the patient outside the traditional care setting. In a scoping review that included 19 studies, the data integrated into the electronic record were biometric and activity data in 57.9% of the studies, and integration into the electronic health record (EHR) was only in its early stages [54].

Table 2.1. Types of medical data, with examples and their usual form of storage. Source: synthesis based on [11,53,54].

Type	Clinical example	Structure	Typical storage
Numeric value	Serum potassium 4.1 mmol/L	Structured	Numeric field with unit
Code	ICD-coded diagnosis	Structured	Controlled text field
Narrative text	Medical history, discharge summary	Unstructured	Free text
Signal	12-lead ECG	Series of samples	Binary file
Image	CT examination	Pixel matrix	DICOM files
Omics data	Genome sequence	Nucleotide string	FASTQ, BAM, VCF
Patient-generated data	Daily steps, home blood glucose	Variable	Mobile app, wearable device

2.1.5 Data Quality

A wrong datum leads to a wrong decision, even when the program that processes it is flawless. There are four classic dimensions of quality:

- **accuracy** (the value matches reality: the recorded weight is the actual weight)

- **completeness** (the required values are present)
- **timeliness** (the value is recent enough for the decision at hand)
- **consistency** (values do not contradict one another, either within a system or across systems: a male patient cannot have an ongoing pregnancy).

A systematic review of 103 papers published between 2013 and 2023 found that the dimension most often assessed in EHR data is completeness, followed by correctness, concordance, plausibility, and timeliness. The authors also added conformance and bias and found that there is still no standard assessment method [55]. Another review, based on 227 articles, proposed six dimensions (accessibility, accuracy, completeness, consistency, contextual validity, timeliness) and identified consistency as the dimension that influences all the others [56]. An example shows how the dimensions are linked. In the record of a patient with heart failure, weight is noted at admission but missing at discharge (a completeness problem). The admission value was copied from a record made two years earlier (a timeliness problem). On top of that, it is expressed in pounds in one system and in kilograms in another (a consistency problem). In this case, the automatic calculation of a diuretic dose or of the fluid balance starts from a wrong datum, even though each program works correctly.

Chapter 3 returns to the practical methods of quality control.

2.2 Information Theory Applied in Medicine

2.2.1 The Amount of Information and the Bit

The mathematical theory of information, formulated by Claude Shannon in the mid-twentieth century, starts from a simple idea: a message carries information to the extent that it reduces uncertainty. A certain event conveys nothing new. A rare event is highly surprising. The amount of information of an outcome with probability p is:

$$I = -\log_2 p \text{ (bits)}$$

The bit (*binary digit*) is the amount of information obtained when one learns the outcome of a choice between two equally probable alternatives. The functions of information theory (entropy, relative entropy, mutual information) apply directly to diagnostic tests, because a patient's disease state and the test result are random variables [57].

Numerical example. A patient is tested for a disease. If, before the test, the probability of disease is 50% (0.5), learning that the diagnosis is positive brings $I = -\log_2 0.5 = 1$ bit. If the disease has a prevalence of 1% in the tested population, learning that the patient is ill brings $-\log_2 0.01 \approx 6.64$ bits, whereas learning that the patient is healthy brings only $-\log_2 0.99 \approx 0.014$ bits. For a prevalence of 0.1%, a true positive result means almost 10 bits (Figure 2.2). Intuitively, the news that a young person without risk factors has a rare disease is surprising, while confirmation that he or she does not have it changes almost nothing.

2.2.2 Entropy

Entropy is the average information expected from a variable, that is, its uncertainty before observation:

$$H = -\Sigma\, p \cdot \log_2 p$$

For a binary variable (ill / healthy) with probabilities p and 1 - p, entropy is maximal, 1 bit, at p = 0.5. For p = 0.01, $H = -0.01 \cdot \log_2 0.01 - 0.99 \cdot \log_2 0.99 \approx 0.066 + 0.014 \approx 0.08$ bits (Figure

2.1). The difference has a direct clinical meaning. When the prevalence is 50%, the clinician knows almost nothing before the test. When the prevalence is 1%, the answer "healthy" is almost certain in advance, and the remaining uncertainty is small.

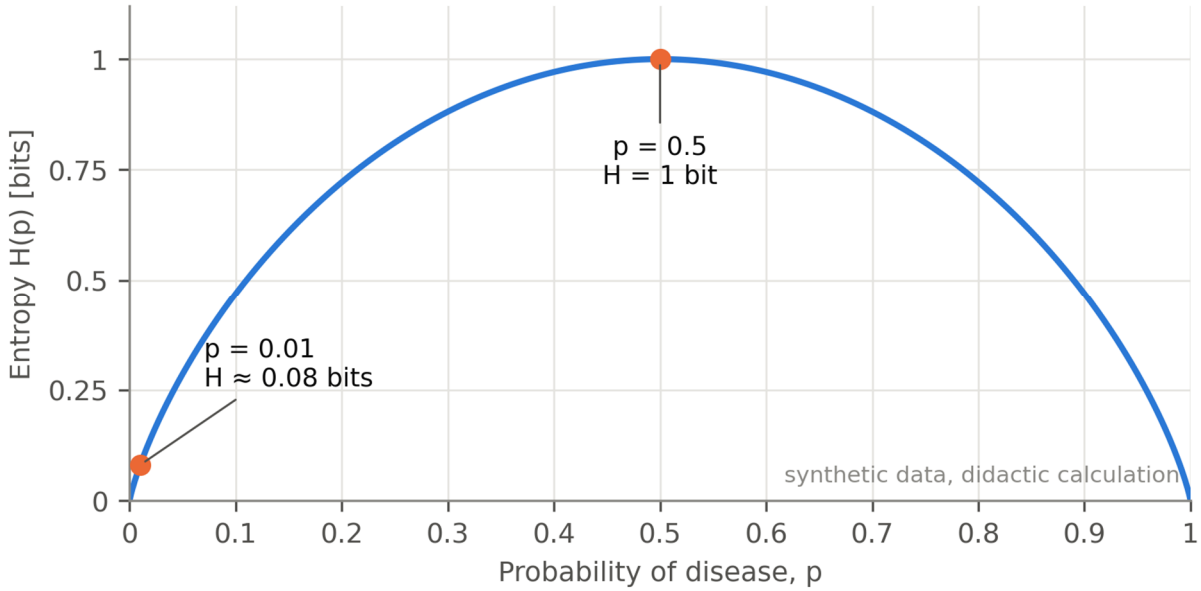


Figure 2.1. Entropy of a binary variable H(p) = -p·log2 p - (1-p)·log2(1-p), with the points p = 0.5 (1 bit) and p = 0.01 (≈ 0.08 bits). Didactic calculation, synthetic data.

Redundancy is the difference between the maximum capacity of a code and the information it actually transmits. In medicine it has a protective role. The patient's name, date of birth, and personal numeric code (cod numeric personal, CNP) partly repeat the same identity, so that an error in one field can be detected from the others. Noise is any disturbance that alters the message along the way: a motion artifact on the ECG, a digit transcribed incorrectly, a sentence not heard at shift change. Well-designed redundancy makes it possible to detect, and sometimes to correct, the errors caused by noise.

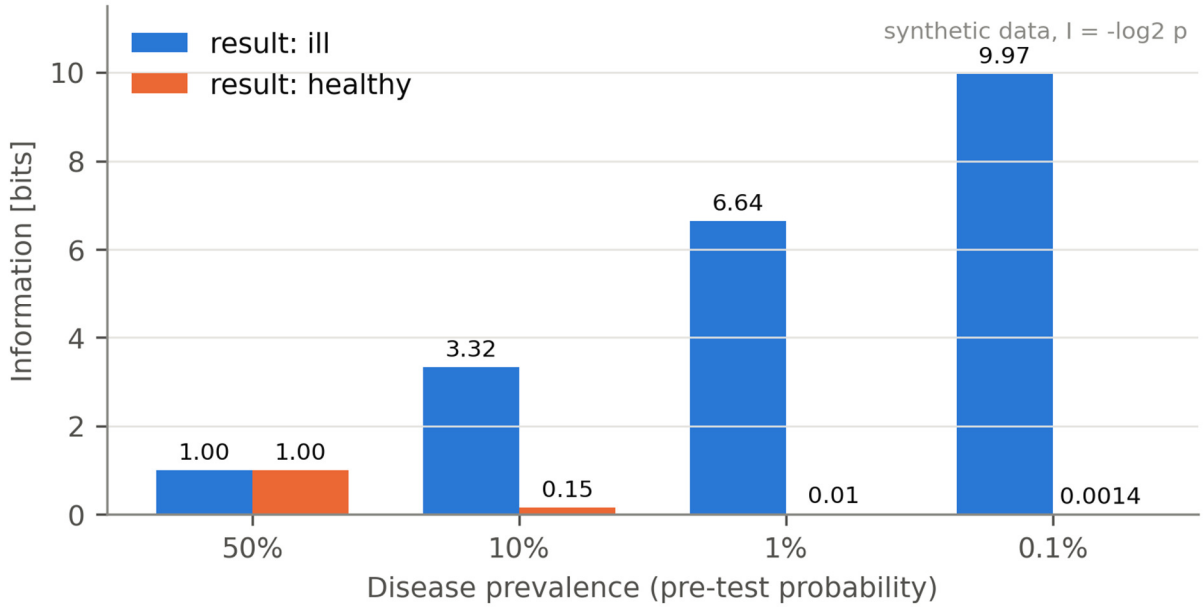


Figure 2.2. Amount of information (bits) obtained by learning the patient's true state, for four values of disease prevalence. An "ill" result brings more bits the rarer the disease, whereas a "healthy" result brings less and less. Didactic calculation, I = -log2 p, synthetic data.

2.2.3 Information Gained from a Diagnostic Test

Classic diagnostic reasoning starts from the pre-test probability (estimated from the history, the clinical examination, and the prevalence) and arrives, through Bayes' theorem, at the post-test probability, using the sensitivity and specificity of the test [58]. Likelihood ratios express how much a result changes the probability of disease and, combined with the pre-test probability, allow the post-test probability to be estimated [59].

Information theory adds a complementary question: by how much, on average, does uncertainty decrease after the test? The measure is mutual information, that is, the entropy before the test minus the average entropy after the test. Sensitivity and specificity, used separately, do not measure this reduction explicitly, whereas mutual information expresses it as a single number [60]. For example, an ultrasound examination for endometrial cancer with 94% sensitivity and 40% specificity reduces diagnostic uncertainty by only 10%, a result that the two percentages taken separately conceal [60].

Worked example. Consider a test with a sensitivity and a specificity of 90%, applied to patients with a pre-test probability of 10%:

- pre-test entropy: $H(0.10) \approx 0.47$ bits
- 18% of the patients have a positive test, and for them the post-test probability is 50%, hence $H = 1$ bit
- 82% have a negative test, with a post-test probability of 1.2% and $H \approx 0.10$ bits
- average entropy after the test: $0.18 \cdot 1 + 0.82 \cdot 0.10 \approx 0.26$ bits
- information gained: $0.47 - 0.26 \approx 0.21$ bits, that is, a reduction in uncertainty of about 45%.

The same test applied at a prevalence of 1% gains only 0.025 bits, a reduction of about 31%. The important teaching point concerns the patient with a positive test: the probability of disease rises from 10% to 50%, but that patient's individual uncertainty increases (from 0.47 to 1 bit). The test was useful, because it moved the patient into a category that requires further investigation, yet it did not resolve the diagnosis. A narrative review of 20 studies that used entropy in clinical decision making found that it serves mainly to quantify uncertainty and that no study has yet turned it into a validated operational decision tool [61].

2.2.4 Encoding: Bit, Byte, ASCII, Unicode

A computer stores all information as strings of bits. Eight bits form a byte (octet), which can take $2^8 = 256$ values. The first widespread standard for text, ASCII (American Standard Code for Information Interchange), used 7 bits, hence 128 characters, enough for the English alphabet [62]. Unicode was created to cover all the world's writing systems. Version 18.0, released on September 16, 2026, contains 172,808 characters [63]. The UTF-8 encoding form, widely used in data exchange, represents each character with 1-4 bytes and leaves ASCII characters unchanged [64].

For Romanian names, this detail has practical consequences. Unicode contains two different letters that look almost the same: "s with comma below" (ș, code point U+0219, added explicitly for Romanian) and "s with cedilla" (code point U+015F), and the same pair exists for t (U+021B versus U+0163) [65]. In UTF-8, ș is written as the bytes C8 99, and s with cedilla as C5 9F. The consequence is also visible in string length: the name Mureșan occupies 8 bytes in UTF-8, whereas the form without diacritics, Muresan, occupies 7, because ASCII letters are encoded in a single byte. To a computer, the name Mureșan written with ș (U+0219) and the same name written with s with cedilla (U+015F) are two different strings, although they look almost identical

on screen. The same patient may therefore appear as two distinct persons if one system (for example, an old laboratory program) uses the cedilla and another (the electronic record) uses the comma.

Example

- A hospital links admission records to laboratory results by name and date of birth (without the personal numeric code, for data protection reasons).
- The laboratory transmits the name Ştefănescu with the initial letter encoded as S with cedilla (U+015E), while the electronic record stores it with Ș with comma below (U+0218). Exact string comparison fails, and the results are not attached to the record.
- The solution: normalizing the text before comparison (all cedilla variants are converted to the comma variant, and the forms without diacritics may also be compared) and, where possible, a shared unique identifier.

The problem is not only a Romanian one. When the records of a clinical data warehouse from three French university hospitals had to be linked with the national death registry in the absence of a common identifier, an algorithm with advanced name cleaning and knowledge of the naming system increased linkage sensitivity from 82.7% to 93.3% [66].

2.2.5 The Size of Medical Files

Orders of magnitude help in understanding the costs of storage and transmission. A standard ECG from the public PTB-XL database lasts 10 seconds, has 12 leads, is sampled at 500 Hz, and is stored with 16-bit precision [67]. This gives 12 × 5,000 samples × 2 bytes = 120,000 bytes, or about 120 kB. A typical CT examination has images of 262,144 pixels (a 512 × 512 matrix) and around 500 images, hence about 250 MB, while a mammogram reaches about 153 MB [68]. Scanned pathology images are usually around 1 GB each [68].

The complete human genome (the T2T-CHM13 assembly) has 3.055 billion base pairs [69]. At 2 bits per base (Section 2.6), the sequence of a single copy would theoretically fit into about 764 MB. In practice, the raw data from whole-genome sequencing at a coverage of about 35× occupy roughly 65 GB in fastq.gz files, and the aligned BAM file takes up roughly another 55 GB [70]. The difference comes from the deliberate redundancy of sequencing (each position is read dozens of times) and from the quality scores attached to each base.

Table 2.2. Orders of magnitude of some medical files. The values for the ECG and the reference sequence are calculated from the published parameters. Source: [67-70].

Data type	Parameters	Order of magnitude
12-lead ECG, 10 s	500 Hz, 16 bits	~120 kB (calculated)
CT examination	512 × 512 pixels, ~500 images	~250 MB
Mammogram	4 images	~153 MB
Scanned slide (pathology)	one image	~1 GB
Genome, one copy of the sequence	3.055 billion bases × 2 bits	~764 MB (calculated)
Sequenced genome, raw data	~35× coverage, fastq.gz	~65 GB

2.3 Data, Information, Knowledge, Wisdom

2.3.1 The DIKW Hierarchy

The DIKW model (data, information, knowledge, wisdom) arranges the content of information systems on four levels. **Data** are raw observations. **Information** means data analyzed and organized so that they acquire meaning, for example for display or analysis. **Knowledge** results from the formal or informal interpretation of information and includes the results of studies, but also heuristics and models that may reflect the experience or the biases of those who formulate them [11]. **Wisdom** is the ability to apply this knowledge to the concrete situation, weighing the patient's values and context.

The classic example: a patient's blood pressure of 180/110 mmHg is a datum. When the data of many patients are brought together and analyzed, it turns out that hypertensive patients have myocardial infarctions more often, and this is knowledge [11]. Figure 2.3 illustrates the levels with an example from diabetes mellitus.

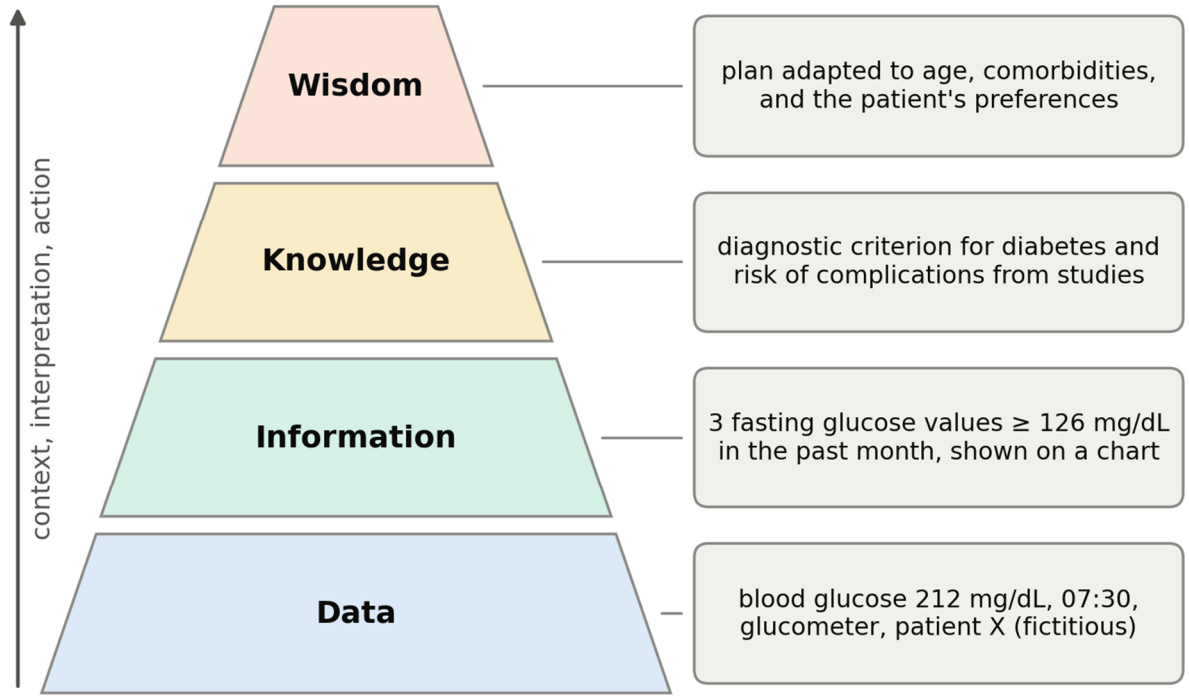


Figure 2.3. The data-information-knowledge-wisdom (DIKW) hierarchy, with clinical examples from diabetes mellitus at each level. Didactic diagram, adapted from the definitions in [11].

2.3.2 Limitations of the Model

The pyramid is useful for teaching, but it simplifies reality in at least three respects:

- **the boundaries depend on the level of abstraction**. The statement "the patient has hypertension" is an interpretation of measurements, but it becomes a mere datum for a system that calculates cardiovascular risk [11]
- **the flow is not one-directional**. Existing knowledge determines which data are collected: a physician who suspects a disease orders certain tests and not others

- **action is missing**. In public health surveillance, it has been proposed to insert a level of "intelligence" (knowledge put into action) between knowledge and wisdom, so that data reporting is followed by a response [71].

Knowledge graphs are a modern example of the passage from data to knowledge, achieved by linking biomedical entities into a network. The integration of real-world EHR data into these graphs has, however, remained limited [72].

2.3.3 Explicit Knowledge and Tacit Knowledge

Explicit knowledge can be written down, transmitted, and verified: a clinical guideline, a dose, a resuscitation algorithm. **Tacit** knowledge comes from experience and is hard to put into words. Examples are the clinical impression at the bedside or the surgeon's "hand." A systematic review accompanied by a hermeneutic analysis showed that, in critical situations such as the COVID-19 pandemic, guidelines and scientific knowledge were not sufficient to guide action, and tacit knowledge served as a way of teaching and learning skills related to values and ethics [73]. For informatics, the distinction has a practical consequence: only explicit knowledge can be encoded directly into a system.

2.3.4 Knowledge Representation

To be used by a computer, explicit knowledge must be written in a formalism:

- **production rules** have the form IF ... THEN ... The MYCIN system, the prototype of rule-based systems, represented knowledge about infections and antibiotics in this way [14]. The Arden Syntax standard organizes event-condition-action rules into medical logic modules (for example, an alert for an annual foot examination in a patient with diabetes) [14]
- **decision trees** represent the sequence of choices and uncertain events, each branch having a probability and a consequence [58]
- **ontologies** organize the types of entities in a domain into a hierarchy and make the relationships between them explicit, in the form of a graph [14]. The creation of ontologies and knowledge graphs has remained a major research direction [74].

The choice of formalism is not neutral. An analysis of several rule technologies (Arden Syntax, Drools, SWRL, and others), applied to antibiotic therapy guidelines, showed that they differ in expressiveness, interoperability, and industrial maturity [75]. Medical terminologies and ontologies are covered in Chapter 4, and decision support systems in Chapter 13.

2.4 Communicating Information in Medical Practice

2.4.1 The Sender-Channel-Receiver Model

Shannon described communication as a chain: a source produces a message, a transmitter encodes it, the message passes through a channel in which noise intervenes, and the receiver decodes it. In patient handover at shift change (also called handoff), the sender is the departing physician, the receiver is the physician taking over, and the channel is the verbal discussion, the written sheet, or the electronic record (Figure 2.4). Noise takes concrete forms. A meta-synthesis of qualitative studies on handover in intensive care identified, among other things, interruptions and distractions, lack of pertinent information, and one-way communication [76]. A systematic review of 60 studies on patient handover from prehospital emergency crews to emergency department staff described environmental factors, degradation and loss of information, lack of standardization, and poor recall [77].

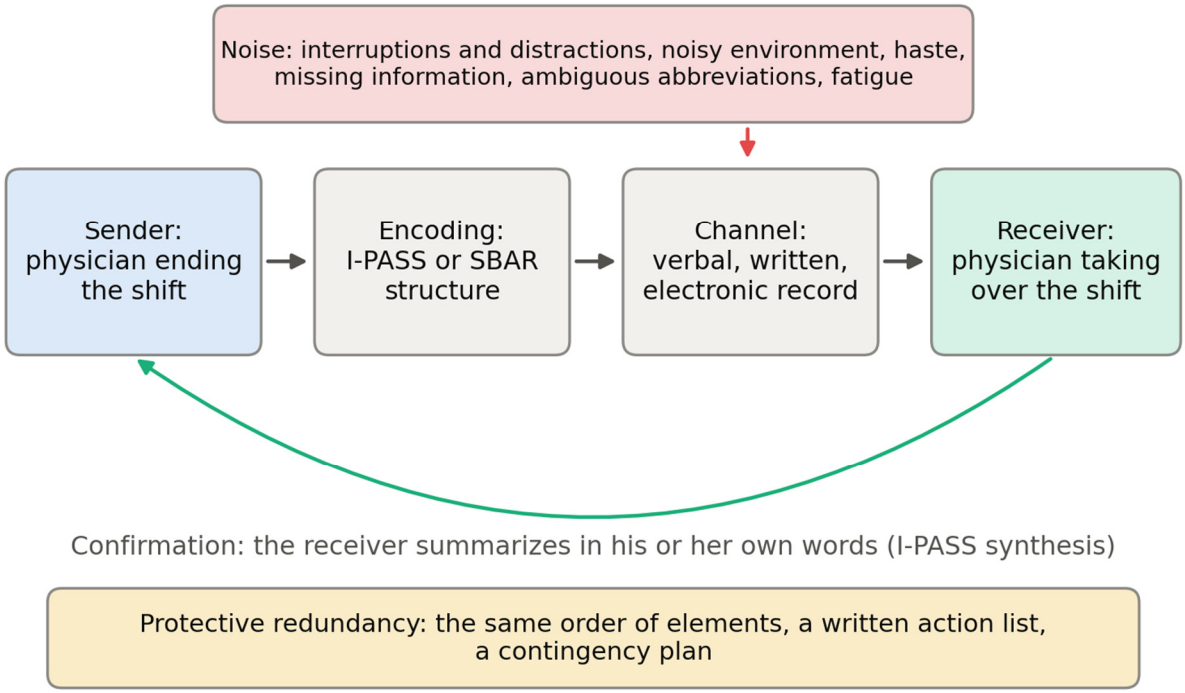


Figure 2.4. The sender-channel-receiver model applied to patient handover at shift change, with sources of noise and protective mechanisms (standard structure, confirmation by the receiver through a summary). Didactic diagram based on Shannon's model, with sources of noise from [76,77].

2.4.2 Communication Errors and Structured Protocols

The consequences are measurable. In a sample of US malpractice claims, communication failures were identified in 49% of cases, and 40% of these failures involved a failed handover. The information most often miscommunicated concerned contingency plans, the diagnosis, and the severity of illness [78].

Structured protocols in fact apply the principles of information theory: a fixed order of elements (a code shared by sender and receiver) and confirmation by the receiver (redundancy). **SBAR** (Situation, Background, Assessment, Recommendation) requires the presentation of the situation, the background, the assessment, and the recommendation. **I-PASS** (Illness severity, Patient summary, Action list, Situation awareness and contingency planning, Synthesis by receiver) explicitly adds the contingency plan and the synthesis performed by the receiver.

The evidence from 2020-2025 is encouraging but uneven. In a prospective study in 32 hospitals with 2,735 resident physicians, the implementation of I-PASS was associated with a 47% decrease in reported handover-related adverse events (major events fell from 1.7 to 0.9 per person-year). The proportion of verbal handovers that contained all five key elements rose from 20% to 66% [79]. A 2025 systematic review concluded that I-PASS reduces medical errors and adverse events with moderate certainty of evidence, while the evidence that SBAR improves safety outcomes is of low certainty [80]. For SBAR, another review showed that faithful application and improved clarity of communication appeared mainly in classroom-based studies. Effects on patients were observed only in complex, multicomponent interventions [81].

The channel may also be electronic. A systematic review of 9 publications on handover through electronic platforms suggested that this approach may improve the completeness of the information transmitted. It is not clear, however, whether this gain translates into better patient outcomes [82]. An electronic handover form therefore remains a support for structure, not a substitute for direct discussion and for confirmation by the receiver.

Example

I-PASS handover, fictitious patient, internal medicine ward, 20:00.

- I (severity): "Unstable patient, under close observation."
- P (summary): "68-year-old man, community-acquired pneumonia, second day of antibiotics."
- A (actions): "Reassess saturation at 22:00; check the blood cultures in the morning."
- S (contingencies): "If SpO2 falls below 90% on mask oxygen, call the on-call ICU physician."
- S (receiver synthesis): the physician taking over restates the information in his or her own words, and any discrepancies are clarified on the spot.

2.5 Evidence-Based Medicine in Digital Format

2.5.1 Levels of Evidence and Their Certainty

Evidence-based medicine (EBM) combines the best available evidence with the clinician's expertise and the patient's preferences. Traditionally, evidence is ranked by study design: systematic reviews and meta-analyses, randomized controlled trials, cohort studies, case-control studies, case series, expert opinion. The GRADE approach (Grading of Recommendations Assessment, Development and Evaluation), however, assesses the certainty of evidence separately for each outcome, starting from a question formulated according to the PICO structure (population, intervention, comparator, outcome), and it separates the judgment about the evidence from the move to a recommendation [83].

2.5.2 Evidence Syntheses and Living Systematic Reviews

Clinicians cannot read all primary studies. This is why syntheses are used: systematic reviews, scoping reviews, evidence maps, reviews of reviews. Rapid reviews and living reviews are modes of work that can be applied to several types of synthesis [84]. A living systematic review is updated continuously, as new studies appear. It is especially useful when evidence emerges quickly and is uncertain, and when new results may change practice [85]. The same principle applies to living guidelines. Their methodology includes deciding whether a guideline warrants the living approach, literature surveillance and the frequency of searches, as well as retiring recommendations from living mode once the topic stabilizes [86].

At the point of care, physicians usually consult synthesis sources that have already been appraised: the Cochrane Database of Systematic Reviews, BMJ Best Practice, UpToDate, and similar products [13].

Table 2.3. Types of evidence sources and their use. Source: [13,84-86].

Source type	What it offers	When it is used
Primary study (RCT, cohort)	The result of a single study	New questions, with no synthesis available
Systematic review	Synthesis of all studies on a question	Assessing the effect of an intervention
Living systematic review	Continuously updated synthesis	Fields where evidence changes quickly
Clinical guideline (GRADE)	Recommendations with certainty of evidence	Practice decisions
Point-of-care synthesis resource	Updated clinical summaries	Specific questions during the consultation

2.5.3 Computable Clinical Guidelines

A guideline published as text (PDF) must be read and interpreted by each physician. A **computable** guideline is translated into a form that the information system can execute on patient data. For example, a format based on FHIR resources has been proposed that represents the recommendation, its strength, the certainty of evidence, and the link to the systematic reviews on which it rests [87]. The US Centers for Disease Control and Prevention (CDC) translated the screening guidelines for cervical, breast, and colorectal cancer into computable guidelines and decision support tools. The cervical cancer tool was integrated into the EHR and tested clinically for 6 months. The other two were tested only with synthetic data [88]. The authors note that scaling up depends on structured data capture, that is, precisely on the data quality discussed in Section 2.1.

2.5.4 Automating Evidence Synthesis

A systematic review is expensive and slow. An AHRQ report from December 2024, organized as a living rapid review with an evidence map, included 56 studies that evaluated machine learning tools for different stages [89]. The results differ greatly from one stage to another:

- tools for identifying randomized trials had a sensitivity (recall) of 96-100% (median 98.5%)
- abstract screening tools had a median sensitivity of 93% and reduced the workload by a median of 50%
- automated data extraction was correct in proportions ranging from 0 to 99% (median 10%)
- risk-of-bias assessment showed modest agreement with human reviewers (median weighted kappa 0.16) [89].

The report concludes that the identification of randomized trials and the prioritization of abstracts are approaching routine use under human supervision, whereas searching and data extraction are not yet reliable enough [89]. A preregistered study with GPT-4 found moderate performance in data extraction and variable performance in screening, and the authors recommend substantial caution in using large language models for systematic reviews [90].

2.6 Information in Biology

2.6.1 DNA as a Code

The DNA molecule uses an alphabet of four bases (A, C, G, T). If Shannon's formula is applied to four equally probable symbols, each base can carry at most log2 4 = 2 bits. The complete human genome, with 3.055 billion base pairs [69], therefore contains, as an order of magnitude, about 6 billion bits for a single copy. This figure is a theoretical upper limit. The real sequence is not random (it contains repetitive regions), so the actual information content is lower, and compression programs exploit precisely this redundancy.

The density of DNA as a storage medium is very high. It is estimated at about 10^{1}□ bits per cubic centimeter, which explains the research into using synthetic DNA to archive digital data [91]. The same logic links informatics to biology. Genomics produces data volumes comparable to those of imaging, and their analysis requires the storage and management methods discussed in Chapter 3.

2.6.2 A Note on Quantum Information

Classical computers work with bits that have a value of either 0 or 1. Quantum computing uses **qubits**, which, through superposition, can be in a combination of the states 0 and 1, and which, through quantum entanglement, can be correlated with one another. A 2024 review describes applications being explored in drug discovery, molecular modeling, genomics, and diagnosis. The same review lists the current limitations: hardware, scalability, and integration into the clinical environment [92]. For the first-year student, the cautious conclusion is that, as of September 27, 2026, quantum computing remains a field of research, not a routine clinical tool.

Common Pitfalls

- Confusing data with information: a value without a unit, a date, or a measurement method cannot be interpreted correctly.
- Treating an ordinal variable coded with digits as a numeric variable (the mean "pain grade").
- Believing that a positive result always "clarifies" the diagnosis. At low prevalence, a positive test may leave the patient in greater uncertainty than before.
- Ignoring diacritics: ş with comma below (U+0219) and s with cedilla (U+015F) are different characters for a computer, and record linkage by name may fail.
- Confusing bits with bytes: 1 byte = 8 bits, and transmission speeds are usually expressed in bits per second.
- Handing over a patient without a contingency plan and without confirmation by the receiver.

Key Points

- A clinical datum has five components: patient, parameter, value, time, method.
- The amount of information of an outcome is I = -log2 p. Rare events bring more bits.
- The entropy of a binary variable is maximal (1 bit) at p = 0.5 and falls to about 0.08 bits at p = 0.01.
- The information gained from a test is the pre-test entropy minus the average post-test entropy.
- A 10 s ECG takes up about 120 kB, a CT examination about 250 MB, and the raw data of a genome about 65 GB.
- DIKW is a useful teaching model, with context-dependent boundaries.
- I-PASS has moderate-certainty evidence for reducing handover errors.
- Automated evidence synthesis works well for screening, but not yet for data extraction.

Self-Assessment Questions

1. (SA) A disease has a prevalence of 1% in the tested population. The amount of information brought by learning that a patient has the disease is approximately:

a) 0.01 bits
b) 0.08 bits
c) 1 bit
d) 6.6 bits
e) 100 bits

2. (SA) The entropy of a binary variable (ill / healthy) is maximal when the probability of disease is:

a) 0
b) 0.01
c) 0.25
d) 0.5
e) 1

3. (SA) Which of the following variables is ordinal?

a) blood group
b) NYHA functional class
c) serum creatinine
d) admitting ward
e) body weight

4. (MA) Regarding character encoding in medical data, which of the following statements are correct?

a) ASCII used 7 bits and could represent 128 characters
b) UTF-8 uses between 1 and 4 bytes per character
c) ș with comma below (U+0219) and s with cedilla (U+015F) are the same Unicode character
d) the difference between ș and s with cedilla can prevent the linkage of a patient's records by name
e) UTF-8 changes the codes of ASCII characters

5. (MA) Regarding patient handover at shift change, which of the following statements are correct?

a) in I-PASS, the last element is the synthesis performed by the receiver
b) interruptions and distractions are sources of noise described in intensive care handover
c) a 2025 systematic review found moderate certainty of evidence for I-PASS in reducing medical errors
d) SBAR has demonstrated, with high certainty, a reduction in mortality
e) the contingency plan is not part of I-PASS

CHAPTER 3

Chapter 3. Structured Data, Files, and Databases

Objectives

By the end of this chapter, the student will be able to:

- compare the CSV, JSON, and XML file formats using the same fictitious patient record
- build a simple entity-relationship diagram and identify the primary and foreign keys
- normalize a redundant clinical table up to the third normal form
- write basic SQL queries with SELECT, WHERE, JOIN, and GROUP BY
- explain the role of the OMOP common data model and of the FAIR principles in health data research
- choose an electronic data capture tool suited to a study, based on validation, audit, and compliance criteria

3.1 Files and Formats

3.1.1 File, Format, Metadata

A file is a named sequence of bytes kept on a storage medium. The format is the convention that says how these bytes must be read. The same bytes may represent text, an image, or a signal, and a program that does not know the format cannot interpret them. Metadata are "data about data": who created the file, when, with what device, in what units the values are expressed, what each column means. Without metadata, a table with the column `sbp` and the value 145 does not say whether it is systolic blood pressure in mmHg, or whether it was measured in the office or at home.

3.1.2 Text Formats: CSV, JSON, XML

Text formats can be opened with any editor and are, in general, the most portable. Portability means that the data can be read by other programs, even many years later, without the program that created them.

CSV (comma-separated values) is a table written as text: each line is a record, and the values are separated by commas (or by semicolons, in regional settings that use the comma as the decimal separator). The first line usually contains the column names. The same record of a fictitious patient looks like this:

```
patient_id,birth_date,sex,sbp,dbp,smoker
P0001,1980-03-14,F,128,82,no
```

JSON (JavaScript Object Notation) organizes data as name-value pairs, which can be nested. The same record, extended with two measurements, becomes:

```
{"patient_id": "P0001", "birth_date": "1980-03-14", "sex": "F",
 "blood_pressure": [{"date": "2026-09-20", "sbp": 128, "dbp": 82},
                    {"date": "2026-09-27", "sbp": 134, "dbp": 85}]}
```

The difference is immediately visible. CSV has a flat structure, with a single value per cell, so two blood pressure measurements require either two rows for the same patient or additional columns (sbp_1, sbp_2). JSON accepts lists and objects nested inside other objects, so it is semistructured. **XML** (eXtensible Markup Language) expresses the same hierarchy through opening and closing tags (`<sbp>128</sbp>`). It is more "verbose" but allows strict validation against a schema. The HL7 FHIR Release 5 standard allows clinical resources to be represented in XML, JSON, and Turtle [93]. Interoperability is covered in detail in Chapter 8.

Two format details frequently cause errors when CSV files are exchanged. The first is character encoding. A file saved in UTF-8 and opened by a program that assumes a different encoding displays diacritics incorrectly, and patient names can no longer be compared (Chapter 2). The second is the decimal separator. In Romanian regional settings, the comma is the decimal separator (`37,5` °C), so a file with comma-separated values may be read with shifted columns. The solution is to agree explicitly, before the exchange, on the encoding (UTF-8), the field separator, and the date format (YYYY-MM-DD, which also sorts correctly as text).

3.1.3 Binary Formats and Portability

Binary formats store values directly in the computer's internal representation. They are compact and fast to read, but unreadable without the right program. Medical examples include DICOM imaging files (Chapter 7), ECG or EEG signal files (Chapter 6), and compressed genomic formats. An open, documented binary format is much safer in the long term than a proprietary format that only the manufacturer can read. For archiving the data of a study, the practical rule is: data in an open format (CSV with UTF-8 encoding, Chapter 2), accompanied by a data dictionary (Section 3.6) and a description of the processing steps.

3.2 Data Modeling

3.2.1 The Entity-Relationship Model

Before a database is built, one describes which "things" (entities) are of interest and how they are related. In a fictitious outpatient clinic, the entities are **Patient**, **Consultation**, **Diagnosis**, **Prescription**, and **Drug**. Each entity has attributes (the patient has a date of birth and a sex, the consultation has a date). Relationships have a cardinality:

- a patient may have many consultations, but each consultation belongs to a single patient (a one-to-many relationship, 1:N)
- several diagnoses may be recorded and several prescriptions issued at one consultation (1:N)
- a drug may appear in many prescriptions (1:N from Drug to Prescription).

Figure 3.1 shows the entity-relationship diagram of this database. Clinical data are full of one-to-many relationships, which a spreadsheet handles poorly. A relational database handles them naturally [94].

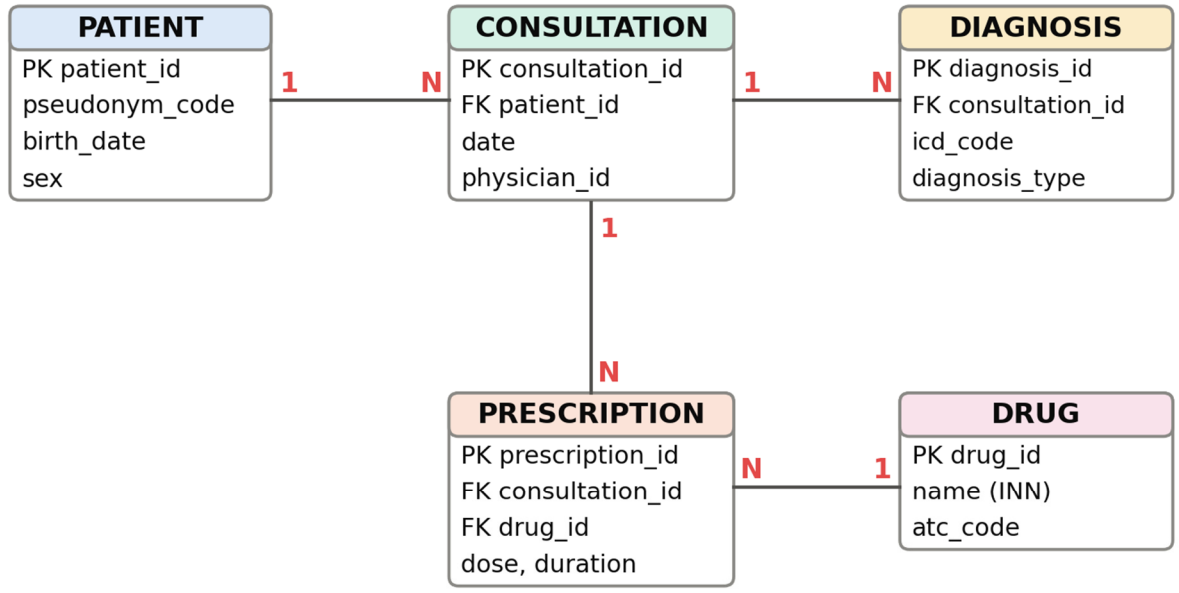


Figure 3.1. Entity-relationship diagram of the fictitious outpatient database: five entities, the primary keys (PK), the foreign keys (FK), and the 1:N cardinalities. Didactic diagram.

3.2.2 Primary Keys and Foreign Keys

Each table has a **primary key**, that is, an attribute or group of attributes that uniquely identifies each row. Its value cannot be missing, and a table has only one primary key [95]. The patient's name is not a good key, because two patients may have the same name, and a name may change or be spelled differently (see the diacritics problem in Chapter 2). An artificial identifier is therefore used, for example `patient_id`.

A **foreign key** is a column that must point to an existing row in another table [95]. The `patient_id` column in the Consultation table is a foreign key to the Patient table. The system thus refuses to record a consultation for a nonexistent patient. The fact that each record has a key also prevents the same patient from being added twice, and the fields have a declared data type, which blocks the entry of meaningless values [94].

3.2.3 Normalization

Normalization is the process by which a redundant table is broken down into smaller tables linked by keys, so that each fact is written only once. The starting point is often a spreadsheet such as the following:

Table 3.1. Unnormalized table with fictitious outpatient data. Each row repeats the patient's data, and one cell contains several diagnoses. Fictitious data.

consultation_id	patient (name, birth year)	date	diagnoses	drug prescribed
C01	Patient A, 1958	2026-09-02	I10, E11.9	metformin 850 mg
C02	Patient A, 1958	2026-09-23	I10	amlodipine 5 mg
C03	Patient B, 1990	2026-09-23	J06.9	paracetamol 500 mg

The table has three defects. Patient A's data are written twice, and a correction made in only one row produces a contradiction (an update anomaly). The "diagnoses" column contains a list, not a value. Finally, the dose and the name of the drug are mixed in the same text.

- **First normal form (1NF)**: each cell contains a single atomic value, and repeating groups are moved into separate rows. The diagnoses of consultation C01 become two rows in a Diagnosis table.
- **Second normal form (2NF)**: in addition, every attribute that is not part of the key depends on the entire primary key, not just on part of it. If a table has the composite key (consultation_id, icd_code), the name of the diagnosis depends only on icd_code and must be moved to a nomenclature table.
- **Third normal form (3NF)**: in addition, no attribute that is not part of the key depends on another attribute that is not part of the key. In the Consultation table, the year of birth depends on the patient, not on the consultation, so it is moved to the Patient table.

The result is the structure in Figure 3.1: each patient appears once, each drug once, and the consultations, diagnoses, and prescriptions refer to them through keys. To speed up analyses, data warehouses sometimes deliberately give up part of the normalization (Section 3.4), but the database into which the data are entered must be normalized.

Normalization also has a cost. To reconstruct the "record" of a consultation, the system must join four or five tables again, and queries become longer. This is why, in practice, a normalized database is kept for entering and updating data, and views or derived tables are built separately for reporting. The teaching rule remains simple: a fact is changed in only one place.

3.3 Relational Database Management Systems and SQL

3.3.1 DBMS and the SQL Language

A database management system (DBMS) is the program that stores the tables, checks the integrity rules, manages simultaneous access by multiple users, and answers queries. Examples of relational DBMSs are PostgreSQL, MySQL, Microsoft SQL Server, and Oracle. They all use SQL (Structured Query Language), which makes it possible to select precise subsets of patients from a very large number of fields [94]. On the fictitious database in Figure 3.1, a few queries look like this:

- selection with a filter: `SELECT patient_id, birth_date FROM patient WHERE sex = 'F'`
- join between tables: `SELECT p.patient_id, c.date FROM patient p JOIN consultation c ON c.patient_id = p.patient_id WHERE c.date >= '2026-09-01'`
- grouping and counting: `SELECT d.icd_code, COUNT(*) AS n FROM diagnosis d GROUP BY d.icd_code ORDER BY n DESC`

The first query returns the female patients. The second links each consultation to its patient (JOIN) and keeps only the consultations from September 2026. The third counts how many times each diagnosis was recorded, that is, it produces a frequency table. The same logic underlies pivot tables in spreadsheets (Practical Session 3), but SQL applies it to tables with millions of rows.

3.3.2 ACID Transactions and Integrity

A transaction is a group of operations that must be treated as a whole. The typical clinical example is dispensing a drug: deducting it from the pharmacy stock and recording its

administration to the patient must happen together or not at all. The ACID properties describe the guarantees of a DBMS [95]:

- **atomicity**: either all the operations of the transaction are completed as a unit or none are, and no partial results remain after a system crash
- **consistency**: the data comply with the integrity constraints at all times, and a transaction that would violate them is canceled
- **isolation**: the effects of a transaction are not visible to other concurrent transactions before it is committed
- **durability**: once committed, the change persists even after a system crash.

Integrity constraints are rules declared in the database: a mandatory field (NOT NULL), unique values, an allowed range (for example, a CHECK that requires systolic pressure to be between 50 and 300 mmHg), and foreign keys. They form the first line of defense for data quality (Section 3.6).

3.4 NoSQL Databases, Data Warehouses and Data Lakes, Big Data

3.4.1 NoSQL Models

The term NoSQL (not only SQL) covers databases that do not organize data in relational tables. There are four main models:

- **document**: each record is a JSON document with a flexible structure (for example, a patient summary with all the diagnoses nested inside it)
- **key-value**: a unique key points to an opaque value (useful for user sessions or image caches)
- **column** (column families): data are stored by column, which is efficient for very large volumes and analytical queries
- **graph**: nodes (patient, diagnosis, drug) and edges (has the diagnosis, receives the drug), efficient for multiple relationships and paths through a network.

A comparative review showed that document, column, and graph NoSQL databases efficiently handle the specific features of standardized EHR data, especially in distributed systems [96]. Graph models are increasingly used for clinical and biomedical data, for example for disease phenotyping and risk prediction, although the integration of EHR data into them has remained limited [72]. Figure 3.2 compares the three main models on the same fictitious patient.

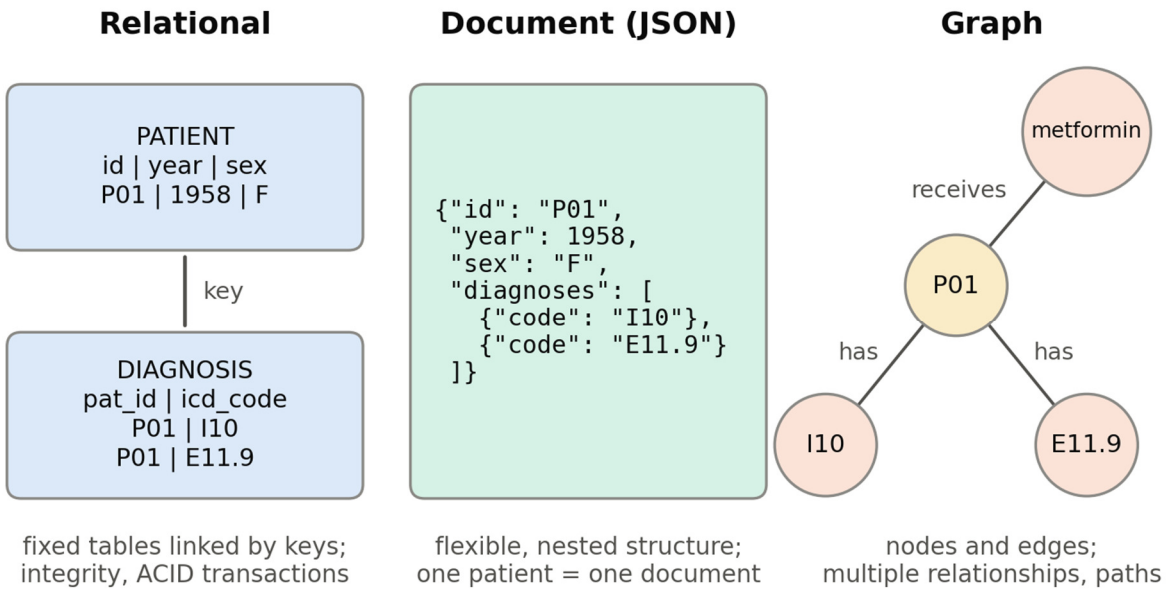


Figure 3.2. The same fictitious patient represented in three data models: relational (tables linked by keys), document (nested JSON), and graph (nodes and relationships). Didactic diagram.

The choice is not "SQL or NoSQL" but depends on the task. Clinical transactions (prescriptions, admissions, billing) need the ACID guarantees and the referential integrity of the relational model. Free-text searches and queries on very large volumes may benefit from other models. At the Bordeaux University Hospital, an i2b2 warehouse with 2.5 million patients and more than 3 billion observations was moved from the relational database to the NoSQL engine Elasticsearch. Query times decreased, with large gains for free-text searches, and less disk space was needed. On the other hand, the model had to be adapted, because Elasticsearch cannot perform joins across different indexes [97].

Table 3.2. Comparison of database models for health data. Source: synthesis based on [72,94,96,97].

Model	Unit of storage	Strength	Example of use
Relational	Table with rows and columns	Integrity, ACID transactions	Prescriptions, admissions
Document	JSON document	Flexible structure	Patient summary, FHIR messages
Key-value	Key-value pair	Direct access by key	Sessions, cache
Column	Column family	Large volumes, column-based analyses	Signal series, logs
Graph	Nodes and edges	Multiple relationships	Knowledge graphs

3.4.2 Data Warehouses and Data Lakes

Routine systems (the electronic health record, the laboratory, the pharmacy) are optimized for transactions: many small operations on a single patient. Research analyses require something else: large queries covering many patients and long periods. This is why data are periodically copied into a **data warehouse**, where they are cleaned, integrated, and organized for analysis. A clinical data warehouse provides the infrastructure for the secondary use of data, for example in research [97]. A **data lake** keeps data in their raw form (tables, text, images, signal files) and structures

them only at the time of analysis. In a modern hospital, data flow from the source systems through integration pipelines to warehouses and data lakes, which explains why a clinician sometimes cannot find a piece of information that "exists" in another system [98].

3.4.3 The Volume of Health Data

The term big data does not refer only to volume. Laboratory data, for example, have the typical characteristics of big data: volume, velocity, variety, and veracity, and their use can generate value in patient care [99]. These five "Vs" can be illustrated by a few published figures:

- the UK Biobank whole-genome sequence data set for 500,000 participants occupies 1.10 tebibytes even after conversion to an optimized, compact format [100]
- the OHDSI network includes more than 331 data sources, with 2.1 billion patient records from 34 countries [101]
- a single university hospital warehouse may contain more than 3 billion clinical observations [97].

A systematic mapping of big data architectures in health identified storage, processing, ingestion, analysis, and visualization as the most frequently addressed layers. The authors also pointed out the lack of an automated layer for checking statistical consistency [102].

3.5 Common Data Models for Research

3.5.1 The OMOP CDM and the OHDSI Network

Each hospital organizes its data in its own way: different tables, different codes, different units. A study in ten hospitals would require ten different analysis programs. The OMOP common data model (Observational Medical Outcomes Partnership Common Data Model, OMOP CDM) standardizes the structure and content of observational data so that the same analysis can run anywhere [103]. The current version, as of September 27, 2026, is 5.5 [103]. The tables have fixed names: PERSON for patients, VISIT_OCCURRENCE for visits, CONDITION_OCCURRENCE for diagnoses, DRUG_EXPOSURE for drugs, MEASUREMENT for measurements.

A common structure is not enough. The content must also use the same codes, and this is the role of the OHDSI Standardized Vocabularies: more than 10 million concepts from 136 vocabularies, mandatory for all the sources in the network [101]. Local data are converted to OMOP through an extract, transform, load (ETL) process: the data are extracted from the source systems, the local codes are mapped to standard concepts, and the data are then loaded into the OMOP tables (Figure 3.3). The conversion takes time and resources, which is why generic, metadata-driven ETL frameworks have been developed [104].

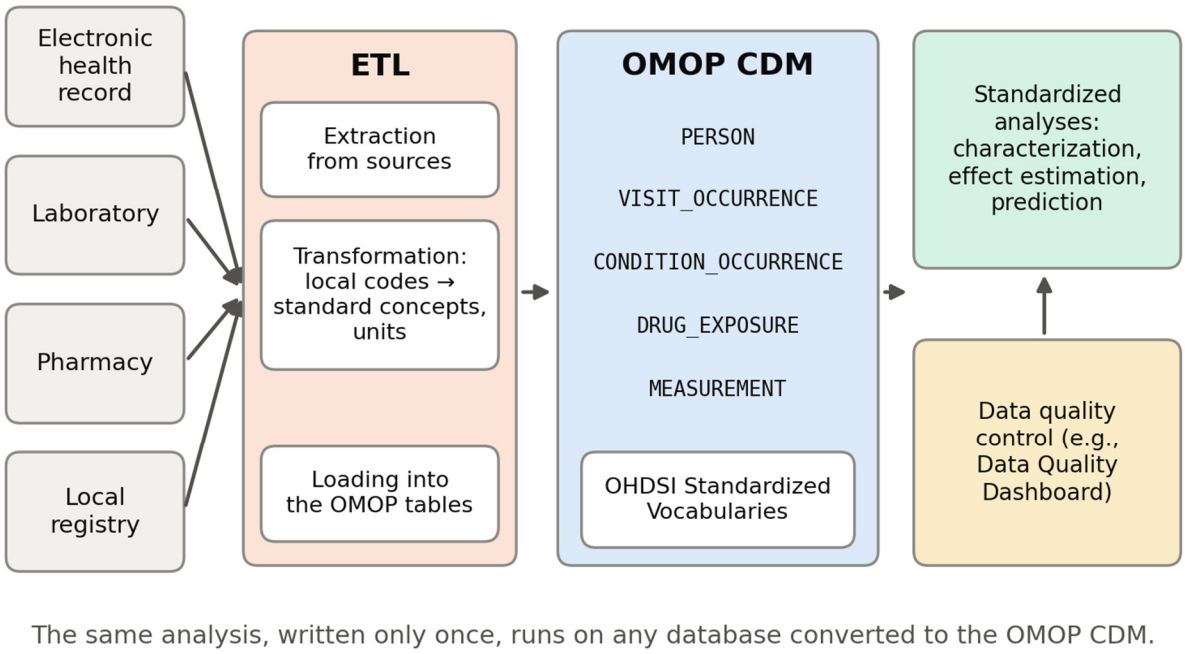


Figure 3.3. From source systems to standardized analyses: extraction, transformation (mapping to the standard vocabularies), and loading into the OMOP CDM, followed by quality control and reproducible analyses. Didactic diagram based on [101,103,105].

European experience shows what this process involves. In the EHDEN project (European Health Data & Evidence Network), 25 partners from 11 countries converted data on more than 133 million patients. The median duration of the 6 key ETL stages ranged from 4 to 115 days. The success factors were the right team, expertise in vocabulary mapping, and early preparation of data governance [105]. The OHDSI community places emphasis on reproducibility: openly published protocols, code, and results, negative and positive controls, and the analysis of multiple sources to check generalizability. The results have included studies on the treatment of hypertension and on COVID-19 [106]. The number of publications on the OMOP CDM has grown steadily since 2016 [107]. A PubMed query carried out by the author confirms the trend and shows a similar increase for the FAIR principles and for REDCap (Figure 3.4).

A simplified example shows what the transformation step does in concrete terms. A fictitious hospital records the diagnosis "essential HTN" in its own system under the internal code 4711, and the laboratory reports blood glucose in mg/dL under the local code GLU. In the ETL, code 4711 is mapped to the standard concept for essential hypertension, GLU is mapped to the standard concept for serum glucose, and the unit of measurement is kept explicitly. After loading, a query written only once ("patients with hypertension and blood glucose above a threshold") produces comparable results in all the hospitals of the network. In EHDEN, the quality of the conversion was also checked with the DataQualityDashboard tool [105].

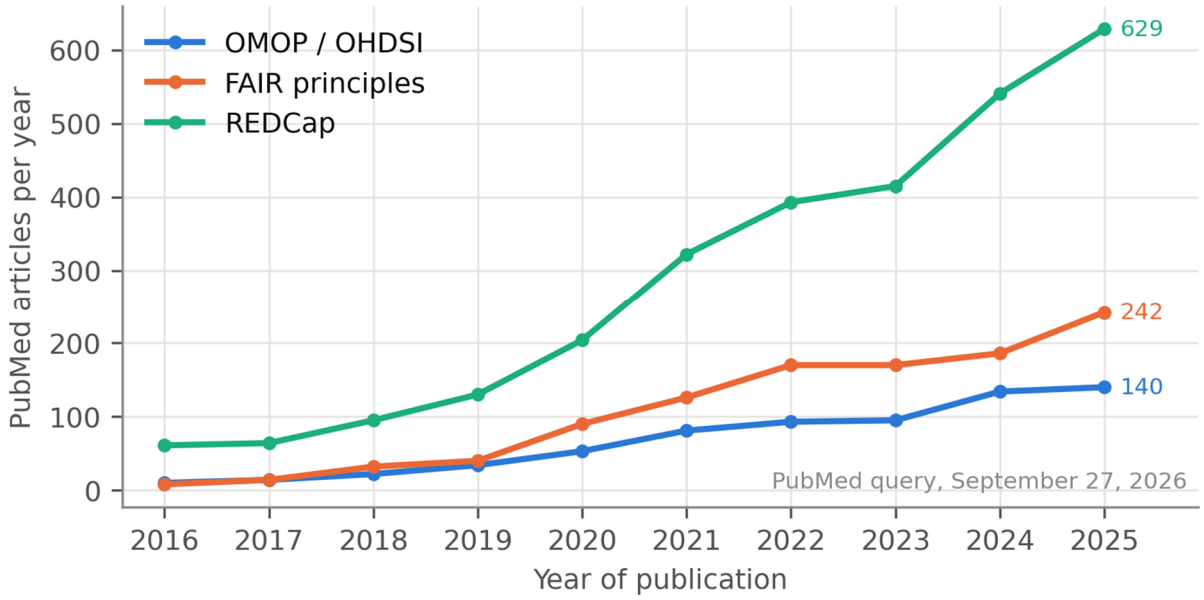


Figure 3.4. Annual number of PubMed-indexed articles that mention OMOP/OHDSI, the FAIR principles, and REDCap in the title or abstract, 2016-2025. Real data, E-utilities query of September 27, 2026 [39]; the values reflect the terms used, not the total number of studies that use these tools.

3.5.2 i2b2

i2b2 (Informatics for Integrating Biology and the Bedside) is an open-source platform for querying and exploring clinical data, designed mainly for cohort identification and feasibility studies through a graphical interface, without programming. The platform reports more than 250 user hospitals [108]. The i2b2 model is classically based on a relational database with a central table of observations [97]. The practical difference from OMOP is one of emphasis. i2b2 quickly answers the question "how many patients meet the criteria?", whereas the OMOP CDM, together with the OHDSI tools, supports complete and reproducible network analyses.

3.5.3 The FAIR Principles in Health

The FAIR principles require data to be Findable, Accessible, Interoperable, and Reusable. In concrete terms, data have unique and persistent identifiers and rich metadata, are registered in searchable resources, can be retrieved through a standard protocol, and use a formal language for knowledge representation [109]. FAIR does not mean "open to everyone". Health data may be accessible only with authentication and authorization, and the principle requires only that the conditions of access be clear.

For health data, the general "FAIRification" workflow has gaps related to curation, validation, de-identification, versioning, and indexing. The European FAIR4Health project proposed a revised workflow, with an architecture based on HL7 FHIR resources [110]. A scoping review of 34 articles found, among the approaches used, complete data dictionaries, semantic interoperability, ontologies, and data quality control. The obstacles included high initial costs, technical and administrative problems, and privacy concerns [111].

3.6 Data Quality and Medical Registries

3.6.1 Dimensions and Assessment Frameworks

Chapter 2 presented the dimensions of quality: accuracy, completeness, timeliness, consistency. Recent reviews confirm that there is still no single standard for assessing the quality

of EHR data [55] and that, among the dimensions, consistency influences all the others [56]. For observational studies, a German framework proposes 34 indicators grouped into four aspects: integrity (compliance with structural and technical requirements), completeness (presence of values), consistency (inadmissible or contradictory values), and accuracy (unexpected distributions and associations). The indicators are calculated automatically in R, based on the study data and metadata [112]. The authors recommend using them both for quality monitoring during the study and for the initial data analysis, before the statistical analyses proper [112].

In practice, the initial analysis of a study file follows a few repeatable steps. The records are counted, and duplicates are searched for by identifier. For each variable, the proportion of missing values is calculated. The minimum and maximum of the numeric variables are displayed, to catch impossible values (a height of 17 cm most likely comes from a typing error for the value 170). The consistency rules between fields are then checked. Each correction is noted in a log, so that the final file can be reconstructed from the raw one.

3.6.2 Input Validation and the Data Dictionary

The cheapest error is the one that never enters the database. Input validation uses a few types of rules:

- **type and format**: the date has the form YYYY-MM-DD, the numeric code has the correct number of digits
- **range**: systolic pressure between 50 and 300 mmHg, age between 18 and 120 years in a study of adults
- **list of values**: sex is F, M, or "unspecified", not free text
- **mandatory field**: the main study variable cannot be missing
- **cross-field consistency**: the discharge date does not precede the admission date, and the diastolic pressure is lower than the systolic pressure.

The **data dictionary** (also called a codebook) describes each variable: the short name, the full label, the type, the unit, the permitted values, and the code used for missing values. It is the part of the metadata without which a CSV file cannot be reused by someone else. This is why it appears among the approaches reported for making health data reusable [111]. A complete example of a data dictionary is given in Practical Session 3.

3.6.3 Clinical Registries

A patient registry systematically collects data on the diagnosis, demographic characteristics, treatment, and outcomes of a defined group of patients. Registries may be local, regional, or national, and general or dedicated to a single disease [113]. Frequently studied examples are the SEER cancer program and the MIMIC-III intensive care data set, which comes from a single hospital. The limitations cited include low patient involvement and few patient-reported outcomes [113]. The quality of a registry is checked systematically. The quality program of the US cardiovascular registries, NCDR, reported complete data and very good but variable accuracy, with room for improvement [114].

Hospital data used retrospectively are close in character to a registry: at a university hospital in Timișoara, the data of 867 patients operated on for colon cancer between February 2019 and February 2023 were extracted and analyzed to compare the pandemic period with the preceding one [115], and an earlier literature synthesis on the same topic brought together the results

reported by centers in several countries, with decreases in elective procedures ranging from 14% to 70% [116].

3.7 Electronic Data Capture Tools

3.7.1 From the Paper Form to EDC

The case report form (CRF) is the instrument through which the data of a study are collected. In its electronic version (electronic data capture, EDC), it applies the validation rules at the very moment of entry and eliminates transcription from paper. The ICH E6(R3) Good Clinical Practice guideline, adopted on January 6, 2025, requires changes to source data to be traceable through an audit trail, without obscuring the original entry, and requires audit trails that cannot be disabled [117]. For personal data in the European Union, the General Data Protection Regulation (GDPR) applies. Data concerning health are personal data that reveal information about a person's health status, and pseudonymization means that the data can no longer be attributed to a person without additional information kept separately [118].

3.7.2 REDCap

REDCap (Research Electronic Data Capture) was created at Vanderbilt University and launched in 2004. As of September 27, 2026, the Vanderbilt page reports about 3.8 million users from almost 8,000 institutions in 163 countries, for about 2.4 million projects [119]. The software is offered free of charge to nonprofit organizations that join the consortium, but it is not available to individuals without an institutional affiliation [120]. It provides access control, an audit trail, calculated fields, the flagging of fields that identify participants, and export to CSV or to the formats of SPSS, SAS, Stata, and R [120]. REDCap has been described as a possible solution to the challenges of collecting, managing, and storing data, in a context of increasingly complex studies and stricter legislation [121]. Its adoption requires institutional support: at the University of the Witwatersrand, the number of active accounts grew from 129 in 2013 to 3,447 in 2021, with periodic training and a local support team [122].

3.7.3 Google Forms and Epi Info

Google Forms is free with a Google account and very easy to use. It allows response validation rules for short-answer questions (a number within a range, length, regular expression) [123] and branching of the questionnaire into sections based on the answer, for multiple-choice and drop-down questions [124]. Responses can be sent automatically to a Google Sheets spreadsheet, which organizes them as a table [125]. Online platforms such as Google Forms make it possible to conduct surveys and collect responses from a large number of respondents, including in medical research [126]. They are not, however, designed as EDC systems for regulated clinical trials.

Epi Info, developed by the CDC, was a suite of public-domain programs for building forms and databases, entering data, and performing epidemiological analysis. The CDC states that product development and technical support were discontinued as of September 2025, and the desktop version remained functional on Windows 11 at that date [127]. For a new project, this situation must be weighed carefully: the program can be used for teaching, but without security updates and without support.

Case Study

- Situation A: a group of students organizes an anonymous questionnaire on sleep, with no identifying data, for a seminar project. Google Forms is sufficient, with validation rules and export to Google Sheets.
- Situation B: a clinic is preparing a multicenter study with pseudonymized clinical data, 6-month follow-up, and external monitoring. It needs an institutional EDC system with an audit trail, access roles, and controlled export, for example REDCap hosted by the university.
- Situation C: an applied epidemiology course uses a computer without a stable internet connection. Epi Info 7 can be used for the exercise, provided that the students know the product is no longer being developed.

Table 3.3. Comparison of three electronic data capture tools (status as of September 27, 2026). Source: [119,120,123,124,127,128].

Criterion	REDCap	Google Forms	Epi Info 7
Cost	Free for nonprofit institutions in the consortium	Free with a Google account	Public domain, free
Validation, branching	Required fields, calculated fields	Response rules, skipping to sections	Check Code module
Audit trail	Yes	Not found in the documentation consulted	Not mentioned on the CDC page
Declared compliance	GDPR, HIPAA, 21 CFR Part 11 (according to the authors)	Not designed for regulated studies	Development stopped in September 2025
Export	CSV, SPSS, SAS, Stata, R	Google Sheets, CSV	Analysis built into the program

The compliance claims in the table refer to the platform. The actual compliance of a study depends on the specific installation, on the institution's procedures, and on the way the team uses the tool.

Common Pitfalls

- Using the patient's name as the identification key: names repeat, change, and are spelled differently.
- Several values in a single cell ("I10, E11.9"), which prevents counting and filtering.
- CSV files without a data dictionary: a year later, no one knows what the column "sbp2" or the code 9 means.
- Confusing the data warehouse (cleaned data, structured for analysis) with the data lake (raw data).
- Believing that FAIR means public data. Health data can be FAIR even with controlled access.
- Collecting identifiers such as the name, personal numeric code (CNP), or e-mail address in a Google form for a study that does not need them.
- Choosing a tool only for its ease of use, without checking the audit trail, the export options, and long-term support.

Key Points

- CSV is flat, whereas JSON and XML are hierarchical (semistructured). FHIR R5 uses XML, JSON, and Turtle.
- The primary key uniquely identifies a row, and the foreign key links it to another table.
- Normalization (1NF-3NF) eliminates redundancy: each fact is written only once.
- SELECT, WHERE, JOIN, and GROUP BY cover most simple clinical queries.
- ACID: atomicity, consistency, isolation, durability.
- The OMOP CDM standardizes the structure, and the OHDSI vocabularies standardize the content of observational data.
- FAIR: findable, accessible, interoperable, reusable, but not necessarily open.
- An EDC tool is chosen on the basis of validation, audit trail, export, and compliance, not on cost alone.

Self-Assessment Questions

1. (SA) In a relational database, the `patient_id` column in the Consultation table, which points to a row in the Patient table, is:

a) the primary key of the Consultation table
b) a foreign key
c) a free-text index
d) file metadata
e) an encryption key

2. (SA) A table in which the "diagnoses" cell contains the value "I10, E11.9" violates:

a) the first normal form
b) the durability property
c) the FAIR principle "findable"
d) the JSON format
e) transaction isolation

3. (SA) The main role of the OHDSI Standardized Vocabularies is:

a) encrypting health data
b) standardizing the content (codes) of the data in the network
c) replacing the SQL language
d) storing DICOM images
e) collecting data through web forms

4. (MA) Regarding the ACID properties of transactions, which of the following statements are correct?

a) atomicity means that either all the operations of the transaction are completed or none are
b) durability means that a committed change survives a system crash
c) isolation means that the data are encrypted
d) consistency means compliance with the integrity constraints
e) ACID is a file format for genomic data

5. (MA) Regarding electronic data capture tools, which of the following statements are correct?

a) REDCap provides an audit trail and export to statistical software
b) Google Forms allows branching into sections based on the answer
c) ICH E6(R3) requires data changes to be traceable through an audit trail
d) the development of Epi Info was discontinued by the CDC as of September 2025
e) Google Forms is designed as an EDC system for regulated clinical trials

CHAPTER 4

Chapter 4. Medical Classifications, Nomenclatures, and Terminologies

Objectives

By the end of this chapter, the student will be able to:

- define and distinguish, using examples, a classification, a nomenclature, a clinical terminology, and an ontology
- break down an ICD-10 code into chapter, block, category, and subcategory, and explain why coding starts from the index but is verified in the tabular list
- apply the definition of the underlying cause of death to a simple death certificate
- build an ICD-11 cluster from a stem code and extension codes, using the "&" and "/" separators correctly
- compare ICD-10, ICD-11, SNOMED CT, LOINC, ATC, and MeSH by purpose, owner, and mode of use
- describe the DRG grouping workflow in Romania and recognize the risks related to coding quality and automated coding.

4.1 Why We Code

4.1.1 Four Practical Reasons

One physician writes "extensive anterior infarction, STEMI"; a colleague at another hospital writes "acute MI with ST elevation in the LAD territory". For a computer that counts cases across a country, these are two different character strings. Coding turns free wording into a standardized symbol that can be counted, compared, and transmitted.

The first reason is **mortality statistics**. Since its adoption, the International Classification of Diseases (ICD) has been the main basis for comparing statistics on causes of death and illness between countries and over time [129]. The ICD has a long history: the first edition appeared in 1893, and revisions followed roughly every ten years, first under the International Statistical Institute and later under the World Health Organization (WHO) [15].

The second reason is **morbidity statistics**: how many admissions for pneumonia a county had, or what proportion of hospitalized patients have diabetes. What matters here is the main diagnosis of the episode of care, a notion defined below (Section 4.2.3).

The third reason is **funding**. In systems based on diagnosis-related groups (DRG), the diagnosis and procedure codes determine the group into which a case falls and, through that group, the amount reimbursed (Section 4.6). In Australia, for example, ICD-10-AM underpins activity-based funding, and the authors of a literature analysis observe that this function has reoriented the objectives of clinical coding [130].

The fourth reason is **research and semantic interoperability**. Coded data are reused in epidemiological studies: a scoping review found 636 articles published between 2012 and 2022 that used Australian data coded in ICD-10 or ICD-10-AM, most of them observational studies [130]. Semantic interoperability means that the receiving system understands the data in exactly the sense in which they were sent. Messaging standards carry the structure, but the meaning comes from the codes of the terminologies.

4.1.2 Classification, Nomenclature, Terminology, Ontology

The terms *terminology*, *vocabulary*, and *nomenclature* are often used as synonyms, even by the authors of coding systems [15]. The differences between the four notions below matter when we choose a tool.

A **classification** is a set of classes that group concepts we no longer wish to distinguish in the analysis. A classification of diagnoses, such as ICD-11, consists of diagnostic classes, and a classification of procedures, such as ICHI, consists of procedure classes [131]. The classes are exhaustive (every case ends up somewhere, possibly in an "other" or "unspecified" class) and mutually exclusive, which allows counting without duplication.

A **nomenclature** provides coded terms that serve to label (index) objects in medicine: cases, patients, documents. SNOMED CT is the example given in the textbook by Winter and colleagues [131]. A nomenclature does not group. It tries to name each clinical concept as precisely as possible.

In the practice of medical informatics, a **clinical terminology** is understood as a controlled nomenclature with stable identifiers, synonyms, and relationships between concepts. Any standard terminology needs a continuous maintenance process; otherwise it quickly becomes outdated [15].

An **ontology** adds logical definitions. Concepts are linked through formal relationships ("is a", "has finding site", "is caused by"), and a program can infer relationships that were not stated explicitly. Ontologies allow structured representation and integration of medical knowledge and are used for interoperability and decision support, but they are difficult to standardize and to keep up to date [132].

A simple comparison helps. In ICD-10, unspecified bacterial pneumonia receives subcategory J15.9, a class used for counting [133]. In SNOMED CT, the concept "bacterial pneumonia" is defined through relationships: it is an infectious disease of the lung, it has the finding site "lung structure", the morphology "inflammation", and the causative agent "bacterium" [15]. The first construction answers the question "how many?"; the second answers "what exactly?".

4.2 ICD-10 and National Modifications

4.2.1 Structure

ICD-10 was published by the WHO in 1992 [15]. The current WHO version is the 2019 version, available in the official browser and accompanied by Volume 2 (the instruction manual, 6th edition) [133]. The classification has three volumes: Volume 1 contains the tabular list, Volume 2 gives the rules of use, and Volume 3 is the alphabetical index [133].

The core is a list of **three-character categories**, each of which can be divided into up to ten **four-character subcategories**. A code begins with a letter followed by digits, and the fourth character comes after a point, for example I21.0 [133]. The three-character level is mandatory for reporting to the WHO mortality database and for international comparisons [133]. Categories are

grouped into blocks (ranges, such as I20-I25), and blocks into chapters numbered with Roman numerals. The 2019 version has 22 chapters, the last being Chapter XXII, Codes for special purposes (U00-U85), which also contains the codes for COVID-19 [133]. Figure 4.1 follows a code from the chapter down to the subcategory.

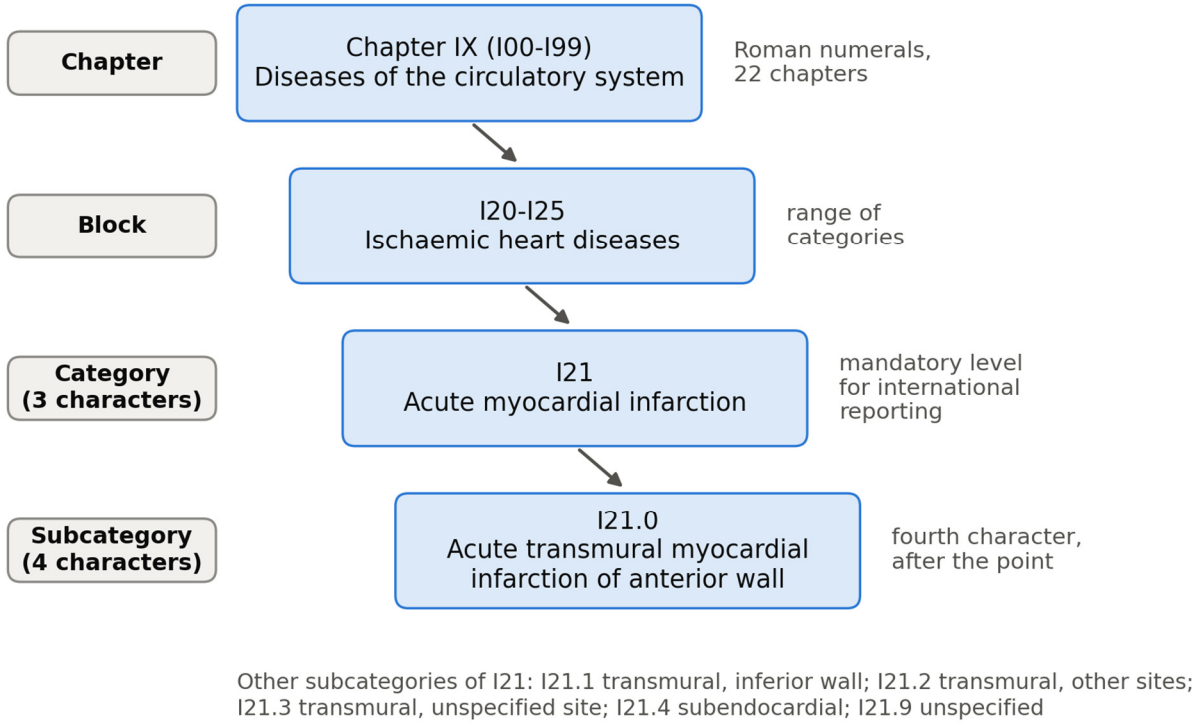


Figure 4.1. Hierarchical structure of ICD-10, illustrated by acute transmural myocardial infarction of anterior wall (I21.0). Codes and titles verified in the WHO ICD-10 browser, 2019 version [133].

Two conventions appear often in exercises. The ending .8 usually means "other", and .9 "unspecified" [15,133]. The "dagger and asterisk" system assigns two codes to a diagnostic statement that contains both a generalized underlying disease and its manifestation in a particular organ: the code for the underlying disease is marked with a dagger (†), and the optional code for the manifestation with an asterisk (*) [133]. *Diabetic polyneuropathy in a patient with type 2 diabetes is thus written E11.4† G63.2*.

4.2.2 Basic Rules: Index First, Then Tabular List

The alphabetical index contains many terms that do not appear in Volume 1. Coding therefore requires consulting both parts before a code is assigned [133]. In brief, the steps the WHO recommends for the occasional user are [133]:

1. identify the type of statement (disease or injury, or external cause) and the corresponding section of the index
2. locate the lead term, usually the noun that names the pathological condition
3. read the notes, the modifiers in parentheses (which do not change the code), and the indented terms (which may change the code)
4. follow the cross-references "see" and "see also"
5. verify the code found in the tabular list, including the inclusion and exclusion notes
6. assign the code.

A code in the index followed by a dash (for example "I21.-") warns that there is a fourth character, which must be looked up in the tabular list [133].

4.2.3 Main Diagnosis and Underlying Cause of Death

For single-condition morbidity analysis, the WHO defines the **main condition** as the condition diagnosed at the end of the episode of healthcare that is primarily responsible for the patient's need for treatment or investigation. If there is more than one such condition, the one that used the most resources is selected. If no diagnosis was made, the main symptom is selected [133].

For mortality, the **underlying cause of death** is tabulated, defined as "(a) the disease or injury which initiated the train of morbid events leading directly to death, or (b) the circumstances of the accident or violence which produced the fatal injury" [133]. The definition has a public health rationale. The most effective prevention breaks the chain at its beginning [133].

An example from the WHO guide shows the difference between the terminal cause and the underlying cause. The certificate reads: (a) pneumonia, due to (b) hip fracture, due to (c) tripping on a carpet. The starting point of the sequence is the tripping, that is, the external event [134]. Statistics will not count this death under "pneumonia". Selection is done in two steps: the starting point is identified, and then one checks whether the special rules modify it [133].

4.2.4 National Modifications and ICD-10-AM

The WHO stopped maintaining ICD-10 in 2018, but countries may use it for as long as they need [135]. Many countries have built their own clinical modifications: the USA has ICD-10-CM, Canada ICD-10-CA, Korea the KCD, and Australia ICD-10-AM [130,136-138]. For ICD-10-CM, the difference in granularity is large. Fung and colleagues counted 10,607 leaf codes in ICD-10 (WHO) versus 71,932 in ICD-10-CM [136]. Australia uses ICD-10 for mortality and ICD-10-AM for morbidity [130].

Romania in 2026

- Inpatient diagnoses are coded in ICD-10-AM, 3rd edition (Australia, July 2002), and procedures in the RO.vi.DRG Classification of Procedures, 1st edition (Clasificarea RO.vi.DRG a procedurilor, ed. I), derived from the Australian ACHI list [139].
- The DRG classification in force is RO DRG v.1.1 (March 2024), with 669 groups: 651 in the 23 major diagnostic categories, 12 pre-MDC groups, and 6 error groups [139].
- The classification is maintained by the National Institute for Health Services Management (Institutul Național de Management al Serviciilor de Sănătate, INMSS). Its history goes back to 2001, and it has been used in DRG funding since 2003 [139].
- As of September 27, 2026, DRG grouping in Romania is still based on ICD-10-AM, 3rd edition. We have not identified any legal act or official timetable for the adoption of ICD-11.

4.3 ICD-11

4.3.1 Adoption and Status in 2026

ICD-11 was adopted by the World Health Assembly in May 2019 [129] and came into effect on January 1, 2022 [140]. The current version in the official browser is ICD-11 MMS release 2026-01 [141]. The WHO fact sheet of May 2026 reports 132 countries at various stages of

implementation and more than 17,000 diagnostic categories [142]. A breakdown from May 2024, still displayed by the WHO, showed that only 14 of these countries and territories were already collecting or reporting data in ICD-11 [140].

4.3.2 Architecture: The Foundation and the Linearizations

The major change compared with ICD-10 lies not in the number of codes but in the construction. ICD-11 has a semantic knowledge base, the **Foundation** (Foundation Component), a biomedical ontology linked to it, and classifications derived from the Foundation. The main derived classification, the successor to ICD-10, is **ICD-11 for Mortality and Morbidity Statistics** (MMS) [129]. The process by which Foundation entities are extracted into a classification for a given purpose is called **linearization**. The content and meaning of an entity do not change through extraction [134]. Specialty linearizations and national modifications can be generated from the same Foundation, the latter only by agreement with the WHO [134]. Figure 4.2 summarizes the architecture.

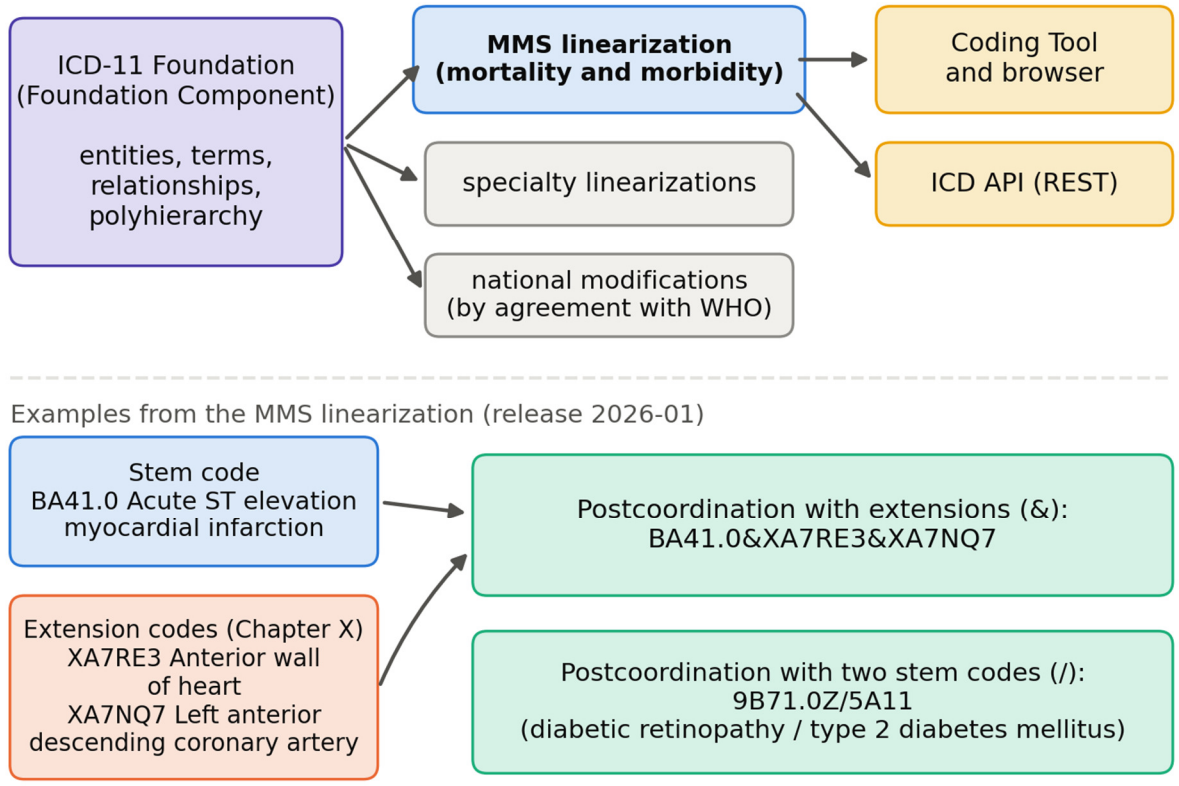


Figure 4.2. ICD-11 architecture: the Foundation, the linearizations, and the WHO tools, with examples of stem codes, extension codes, and clusters from the MMS linearization, release 2026-01. Synthesis based on [129,134]. The cluster examples are taken from the WHO reference guide [134].

4.3.3 The ICD-11 Code: Structure, Stem Codes, and Extension Codes

MMS codes are alphanumeric and cover the range 1A00.00-ZZ9Z.ZZ. The first character always indicates the chapter, the second is always a letter (so that the code cannot be confused with an ICD-10 code), and the letters O and I are omitted so that they are not confused with the digits 0 and 1 [134]. The terminal letter Y is reserved for the residual category "other specified", and the terminal letter Z for the category "unspecified" [134]. Thus, BA41.Z means acute myocardial infarction, unspecified.

Codes that can be used on their own are called **stem codes** [134]. **Extension codes** begin with the letter X, form Chapter X and add detail: anatomy, infectious agent, histopathology, severity, medicament [134]. They are not used alone in statistics and can never occupy the first position of a cluster [134].

4.3.4 Postcoordination

Postcoordination links, when needed, two stem codes, or a stem code with one or more extension codes. The resulting group is called a **cluster** and describes a single clinical concept [134]. The syntax is fixed [134]:

- "&" links a stem code to its extension codes: stem code&extension1&extension2
- "/" links two stem codes: code1/code2
- when a cluster contains two stem codes, each with its own extensions, the extensions are written immediately after the stem code to which they belong.

The opposite is **precoordination**: a single stem code already contains all the elements. Squamous cell carcinoma of bronchus or lung, 2C25.2, combines site and histology in a single code [134].

Example

Diagnosis (fictitious case): acute ST elevation myocardial infarction, anterior wall, lesion of the left anterior descending artery.

- stem code: BA41.0 Acute ST elevation myocardial infarction
- anatomy extension: XA7RE3 Anterior wall of heart
- anatomy extension: XA7NQ7 Left anterior descending coronary artery
- cluster: **BA41.0&XA7RE3&XA7NQ7**

Diagnosis (fictitious case): closed subcapital fracture of the left femoral neck, after tripping and falling on the same level.

- cluster: **NC72.20&XK8G&XJ44E/PA60** (the fracture, left laterality, closed fracture, then the external cause, "Unintentional fall on the same level or from less than 1 metre")

The codes were verified in the ICD-11 MMS browser, release 2026-01 [141]. The first example reproduces an example from the WHO reference guide [134]. In ICD-10, the first case receives I21.0, without the information on the affected artery, and the second S72.0, with W01 as the external cause code [133].

The order of the stem codes depends on the purpose. In morbidity, the first code is the main condition. In mortality, the first code is the underlying cause. The WHO guide gives the example of type 2 diabetes complicated by diabetic coma: on admission, the cluster is 5A23/5A11 (the coma is the reason for admission), whereas on the death certificate it is 5A11/5A23 (the diabetes initiated the chain) [134].

4.3.5 WHO Tools: Browser, Coding Tool, API

In ICD-11, the printed index is replaced by the online Coding Tool, and the application programming interface (API) provides remote access to ICD-11 content and services [129]. The Coding Tool searches by terms and can directly return an assembled cluster, not just an isolated code [134]. Through the API, a hospital can integrate ICD-11 search directly into the discharge form of the electronic health record.

4.3.6 New Chapters

ICD-11 has five new chapters. Diseases of the blood (Chapter 03) and of the immune system (Chapter 04) are separated. Sleep-wake disorders have their own chapter (07). Conditions related to sexual health are grouped in Chapter 17. Traditional medicine has the supplementary Chapter 26 [134]. There is also a supplementary section for functioning assessment (V), aligned with the ICF, and Chapter X of extension codes [134].

Table 4.1. ICD-10 and ICD-11 compared. Source: [129,133-136].

Feature	ICD-10	ICD-11 (MMS)
Approval / entry into force	published 1992; final WHO version 2019	adopted 2019; in effect since Jan. 1, 2022
Code format	letter + digits, 3-4 characters (A00.0-Z99.9)	4-6 characters, 1A00.00-ZZ9Z.ZZ; second position always a letter
Leaf codes (without extensions)	10,607	14,622 (excluding Chapters 26, V, and X)
Chapters	22 (2019 version)	26 + Section V + Chapter X
Additional detail	dagger/asterisk, multiple codes	postcoordination with "&" and "/"
Search	alphabetical index (Volume 3)	Coding Tool, API
Maintenance	stopped in 2018	annual updates

4.3.7 The Transition: What the Evaluations Show

The field test in Calgary (Canada) recoded, in the 2018 beta version of ICD-11, inpatient charts previously coded in ICD-10-CA. The main condition code matched in 86.3% of cases, and the mean coding time fell from 23.6 to 9.9 minutes per chart as the coders gained experience [137]. In Korea, 27 health information management specialists coding isolated diagnostic terms achieved a mean accuracy of 71.6% in ICD-11, compared with 80.2% in the national classification derived from ICD-10. Clusters, changed terms, and deleted codes contributed to the lower agreement [138].

ICD-11 content covers the older classifications well if postcoordination is used. For 1,725 ICD-10-CM codes, full coverage was 35.2% with stem codes alone, 46.5% after adding Foundation entities, and 89.4% with postcoordination (Figure 4.3) [143].

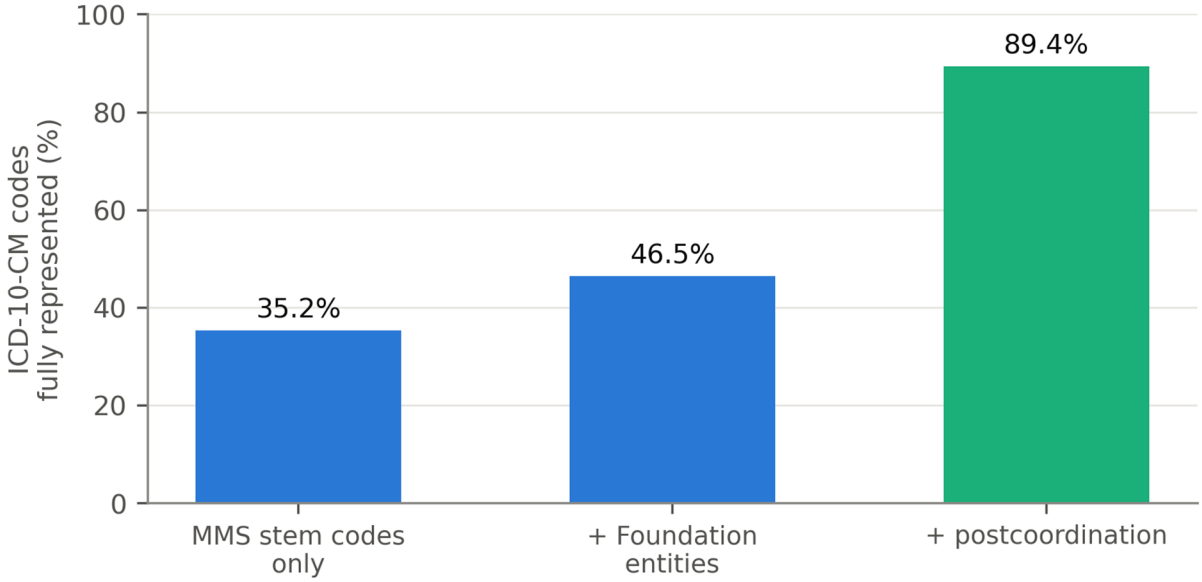


Figure 4.3. Proportion of ICD-10-CM codes (n = 1,725) fully represented in ICD-11, cumulatively, by stage: MMS stem codes, Foundation entities, postcoordination. Real data from [143].

The ICD-10-ICD-11 mapping tables can be downloaded from the "Info" tab of the browser. The WHO warns that they serve to compare data across revisions, not to convert data automatically [135]. Applied to 24,645 deaths in Alberta, the tables successfully mapped 96.6% of the ICD-10 codes for causes of death in adults [144]. By contrast, in the Canadian map from ICD-11 to ICD-10-CA, 65% of the target ICD-10-CA codes were broader than the source ICD-11 code, and only 26% were equivalent [145]. In other words, detail is lost when moving back to ICD-10.

4.4 Other WHO Classifications

The **International Classification of Functioning, Disability and Health** (ICF) was endorsed by the World Health Assembly in 2001, through resolution WHA54.21. It describes body functions and structures, activities and participation, as well as environmental factors, within a biopsychosocial model [146]. The ICD states which disease the patient has; the ICF states what the patient can do with that disease.

The **International Classification of Health Interventions** (ICHI) describes interventions along three axes: the target (the entity on which the action is carried out), the action, and the means. The Beta-3 version was published in October 2020, and the component for public health interventions was finalized in 2023 [147]. ICD, ICF, and ICHI use the same set of extension codes [146].

4.5 Clinical Terminologies

4.5.1 SNOMED CT

SNOMED International presents SNOMED CT as the most comprehensive multilingual clinical terminology in the world. The International Edition exceeds 360,000 concepts, is released monthly, and is used in more than 80 countries [148]. It has three components: **concepts** (each with a unique numeric identifier), **descriptions** (the human-readable terms, including synonyms), and the **relationships** between concepts [148]. The "is a" relationship builds the hierarchies,

starting from a root concept, while attributes such as "finding site" and "associated morphology" define the concept logically [149]. A concept may have several parents. "Bacterial pneumonia", for example, is both an infectious disease of the lung and an inflammation of the lower respiratory tract [15].

SNOMED CT also allows postcoordination, through **expressions** written in a compositional grammar. The expression "80146002 |appendectomy| : 260870009 |priority| = 25876001 |emergency|" describes an emergency appendectomy [149]. Freedom has a price. The same "acute appendicitis" can be coded as a single concept, as a modifier plus a disease, or as a modifier plus a morphology plus a topography, and two coders may arrive at different expressions for the same meaning [15].

A review of 651 articles published between 2020 and 2025 shows that SNOMED CT has moved from the theoretical phase to implementation. Implementation and evaluation accounted for 26.4% and 32.0% of the studies, respectively, and the most frequent use was the extraction of patient data from registries [150]. The same authors note that coding uniformity remains a problem and that evidence of improved patient outcomes is lacking [150]. SNOMED International has 54 national members and has issued affiliate licenses to more than 50,000 individuals and organizations. Romania is not among the members as of September 27, 2026 [151].

4.5.2 LOINC

LOINC (Logical Observation Identifiers Names and Codes) names and codes laboratory tests and other clinical observations. The fully specified name of a LOINC term has five or six parts, separated by ":" [152]:

1. the **component** (the analyte: glucose, creatinine)
2. the **property** (mass concentration, substance concentration, number)
3. the **time aspect** (point in time or interval, for example 24-hour urine)
4. the **system** (the specimen type: serum, plasma, urine)
5. the **scale** (quantitative, ordinal, nominal)
6. the **method**, only when it matters for interpretation.

Glucose in serum or plasma has the code 2345-7, with the fully specified name Glucose:MCnc:Pt:Ser/Plas:Qn [153]. The numeric code has no internal structure, except for the last digit, which is a check digit [152].

4.5.3 ATC/DDD for Medicines

The Anatomical Therapeutic Chemical (ATC) classification is maintained by the WHO Collaborating Centre for Drug Statistics Methodology, hosted by the Norwegian Institute of Public Health. It has five levels and 14 main anatomical groups [154]. Metformin, A10BA02, reads as follows: A alimentary tract and metabolism, A10 drugs used in diabetes, A10B blood glucose lowering drugs, excluding insulins, A10BA biguanides, A10BA02 metformin [154]. The defined daily dose (DDD) is the assumed average maintenance dose per day for a drug used for its main indication in adults. It is a unit of measurement for drug utilization studies and does not necessarily reflect the recommended or prescribed dose [155].

4.5.4 MeSH and UMLS

MeSH (Medical Subject Headings) is the controlled, hierarchical vocabulary of the US National Library of Medicine (NLM), used for indexing, cataloging, and searching biomedical

information and updated annually [156]. A term may appear in several places in the hierarchy [15]. MeSH is not used to code patients. It is the tool for searching the literature in PubMed (Chapter 1).

The UMLS (Unified Medical Language System) brings together vocabularies and standards, including ICD-10-CM, LOINC, MeSH, RxNorm, and SNOMED CT. It has three knowledge sources: the Metathesaurus, the Semantic Network, and the SPECIALIST Lexicon [157].

Table 4.2. The main classification systems and terminologies. Source: [140,146-148,152,154,156,157].

System	Type	Owner	Domain	Main use
ICD-10 / ICD-11	classification	WHO	diseases, causes of death	statistics, funding
ICF	classification	WHO	functioning, disability	rehabilitation, social policy
ICHI	classification	WHO	interventions	reporting of interventions
SNOMED CT	clinical terminology	SNOMED International	clinical concepts	electronic health record
LOINC	terminology	Regenstrief Institute	tests, observations	laboratory results
ATC/DDD	classification	WHO Collaborating Centre (Oslo)	active substances	drug utilization statistics
MeSH	thesaurus	NLM	biomedical subjects	literature searching
UMLS	metathesaurus	NLM	multiple vocabularies	mappings between vocabularies

4.6 Diagnosis-Related Groups (DRG)

4.6.1 The Principle

The DRG is a patient classification scheme that relates the number and types of patients treated in a hospital to the resources they consume [139]. Each group has a **relative weight**, which expresses the expected resource use for that type of hospital stay relative to the average stay [158]. The mean of the weights of the cases discharged by a hospital is the **case-mix index** (CMI) [159]. A higher CMI describes a hospital whose cases consume, on average, more resources [158,159]. Nevertheless, in an American academic medical center, DRG weights explained only about 37% of the variability in the direct costs of hospitalization [158].

In a meta-analysis of 29 studies, DRG-based payment was associated with a shorter length of stay, with no statistically significant effect on 30-day readmissions or on mortality [160].

4.6.2 Grouping Variables and Steps in RO DRG v.1.1

The RO DRG v.1.1 manual lists the variables used for grouping: diagnoses, procedures, sex, age, discharge mode, length of stay, admission weight (for newborns), and day admission [139]. The grouper, a dedicated program, goes through the steps in order: demographic and clinical edits, assignment of the major diagnostic category (MDC), pre-MDC processing, partitioning into surgical, "other", and medical partitions, assignment of the DRG group (GDRG), calculation of the complication and comorbidity levels and, finally, DRG assignment [139] (Figure 4.4).

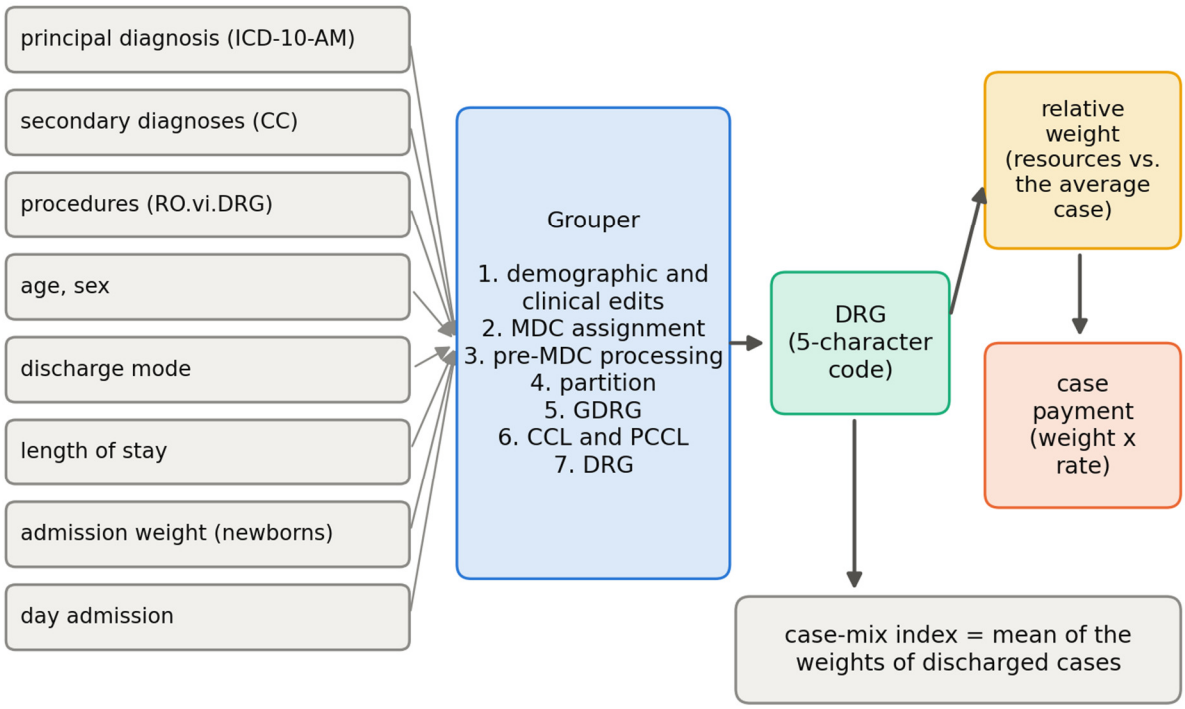


Figure 4.4. The DRG grouping workflow: from discharge data to group, relative weight, and payment. Simplified diagram based on the RO DRG v.1.1 manual [139] and [158].

Among other things, the clinical edits check that the principal diagnosis is not a manifestation code or an external cause code [139]. Pre-MDC processing removes very costly cases from the usual pathway, for example transplantation, ECMO, tracheostomy, or mechanical ventilation for more than 95 hours [139]. Secondary diagnoses receive a complication or comorbidity level from 0 to 4 (from 0 to 3 for medical episodes), and these values are used to calculate the patient's clinical complexity level [139]. The DRG code has five characters: the major category, the type (1 surgical, 2 other, 3 medical), the sequence number, and the complexity grade [139]. A case with an invalid age or with a principal diagnosis incompatible with the patient's sex ends up in one of the error groups, for example 93010 "ungroupable" [139].

The consequence for the physician is direct. An omitted secondary diagnosis can move the case down into a group with a lower weight, and a wrongly chosen principal diagnosis can move the case into another major category.

4.6.3 Risks: Undercoding, Upcoding, Data Quality

Coded administrative data do not mirror clinical reality exactly. A systematic review of 280 validation studies found wide interquartile ranges for the accuracy of ICD codes: 44.0-91.4% for sensitivity and 59.4-91.1% for positive predictive value [161]. In an American university hospital, algorithms applied to the structured data in the electronic health record found the disease without the corresponding code in 32.9% of admissions, and 5.8% of admissions could have received a higher DRG [162].

Conversely, **upcoding** means coding a case at a higher severity than the real one in order to obtain a higher payment. It may be fraud, but it may also reflect more complete documentation of real severity [163]. In five US states, discharges in the MS-DRG with the highest coding intensity increased by 41% between 2011 and 2019. After adjusting for patient characteristics, the authors estimate that, without the change in coding behavior, the increase would have been 13% [163].

4.7 Computer-Assisted Coding and Artificial Intelligence

Manual coding is slow and expensive. A systematic review of 54 studies published between 2019 and 2024 describes the shift from classical machine learning to deep learning, with convolutional, recurrent, transformer, and hybrid networks. The models perform better on frequent codes than on the full code set, and rare codes remain unresolved [164].

Large language models (LLMs) seem well suited at first sight, since they "read" clinical text. The evaluations show otherwise. In a test on more than 27,000 unique diagnosis and procedure codes extracted from the electronic health records of the Mount Sinai Health System, the models were given the official description of each code and had to generate the code. The best model, GPT-4, reproduced exactly 45.9% of the ICD-9-CM codes, 33.9% of the ICD-10-CM codes, and 49.8% of the CPT codes [165]. On 300 discharge summaries, even reasoning models predicted the ICD-9 code correctly in at most 45.3% of cases, and coding was the weakest of the tasks tested [166]. In a test on synthetic clinical vignettes, the SNOMED CT codes proposed by LLMs were systematically invalid, even though the text labels looked plausible [167]. This illustrates a typical risk: the model generates a code that looks correct but does not exist.

Useful help comes from systems that propose codes while a human decides. In a randomized crossover study from Sweden and Norway, the Easy-ICD tool reduced the median coding time for long clinical notes by 46%. Accuracy increased (from 62% to 67% for long notes), but the difference was not statistically significant [168]. The study analyzed only 15 participants in an online setting, and the authors call for evaluations in the real hospital workflow [168].

Common Pitfalls

- Coding directly from memory or from a list of "usual codes", without verification in the tabular list. Exclusion notes often change the code.
- Confusing the terminal cause with the underlying cause of death. The pneumonia on the first line of the certificate is not automatically the cause that is counted.
- Using an ICD-11 extension code alone or in the first position of a cluster. Extensions are always attached to a stem code.
- Believing that the ICD-10-ICD-11 mapping tables convert data automatically. The WHO recommends them for comparison, not for conversion.
- Trusting a code proposed by an LLM without verification in the official browser. The code may be plausible and yet nonexistent.
- Omitting secondary diagnoses at discharge, with effects on DRG assignment and on hospital statistics.

Key Points

- A classification groups cases into exhaustive, mutually exclusive classes. A nomenclature names concepts as precisely as possible.
- ICD-10: chapter, block, 3-character category, 4-character subcategory. Look up in the index, verify in the tabular list.
- The underlying cause of death is the disease that initiated the chain of events, or the circumstances of the accident or violence.
- ICD-11 has been in effect since January 1, 2022. The MMS is the statistical linearization extracted from the Foundation.
- In ICD-11, "&" adds extension codes, and "/" links two stem codes.
- Romania codes inpatient diagnoses in ICD-10-AM, 3rd edition, and groups cases in RO DRG v.1.1 (669 groups).
- SNOMED CT, LOINC, ATC, and MeSH serve different purposes: clinical concepts, laboratory tests, medicines, literature.
- General-purpose language models frequently get codes wrong. Every automatically proposed code must be verified.

Self-Assessment Questions

1. (SA) In ICD-10, the code I21.0 represents:

a) a chapter
b) a block of categories
c) a three-character category
d) a four-character subcategory
e) an extension code

2. (SA) A death certificate reads: (a) pneumonia, due to (b) hip fracture, due to (c) tripping on a carpet. According to the WHO guide, the starting point of the sequence is:

a) the pneumonia
b) the hip fracture
c) the tripping on a carpet (the external event)
d) the patient's age
e) terminal respiratory failure

3. (SA) In the ICD-11 cluster BA41.0&XA7RE3&XA7NQ7, the codes XA7RE3 and XA7NQ7 are:

a) stem codes for other associated diseases
b) extension codes that add anatomical detail
c) ICD-10 codes kept for compatibility
d) ICHI procedure codes
e) SNOMED CT codes

4. (MA) Regarding the DRG system in Romania as of September 27, 2026, which of the following statements are correct?

a) diagnoses are coded in ICD-10-AM, 3rd edition
b) the classification in force is RO DRG v.1.1, with 669 groups
c) groups are determined solely by the principal diagnosis, without other variables
d) age, sex, and discharge mode are grouping variables
e) ICD-11 officially replaced ICD-10-AM in hospitals in 2022

5. (MA) Regarding medical terminologies, which of the following statements are correct?

a) the fully specified name of a LOINC term has five or six parts, including the component, the system, and the scale
b) in the ATC classification, metformin has the code A10BA02, at the fifth level
c) MeSH is used mainly for hospital billing
d) SNOMED CT allows postcoordinated expressions
e) the UMLS is a statistical classification of causes of death

CHAPTER 5

Chapter 5. Ethics, Data Protection, and Information Security in Health Care

Objectives

By the end of this chapter, the student will be able to:

- apply the four principles of biomedical ethics to concrete situations involving the use of health data
- recognize data concerning health within the meaning of the GDPR and identify the legal basis for processing
- distinguish anonymization from pseudonymization and estimate the re-identification risk of a data set
- describe the confidentiality-integrity-availability triad and the basic measures that protect it in a hospital
- explain how algorithmic bias arises and how it can affect the equity of care
- cite correctly, avoid plagiarism, and disclose the use of generative artificial intelligence (AI) tools

Today, a patient's data circulate between the hospital, the laboratory, the family physician, the county health insurance house, and, increasingly, research databases. What does the physician owe the patient once these data become electronic? The answer has an ethical layer, a legal layer built around the General Data Protection Regulation (GDPR), and a technical layer: without information security, ethical and legal obligations remain declarations of intent.

5.1 Ethical Principles Applied to Health Data

5.1.1 The Four Principles

The values commonly accepted in medical ethics are summarized by the four principles of Beauchamp and Childress [169]. They carry over directly into the world of data.

Autonomy means that the person decides what happens to information about him or her: the person knows who sees the data and for what purpose, and consent, refusal, or withdrawal is respected. **Beneficence** requires that data be used in the interest of the patient and of the community: a complete electronic health record can reduce errors, and aggregated data enable studies that improve treatments. **Non-maleficence** (*primum non nocere*) requires the prevention of harm caused by disclosure, discrimination, or decisions based on incorrect data. **Justice** concerns the fair distribution of benefits and risks, for example ensuring that an algorithm trained on a single population does not disadvantage other groups (Section 5.5).

The principles often conflict. A recent example is that of synthetic medical data, generated by AI as a substitute for real data. Friedman and Smajdor show that when privacy is treated as the

supreme value, it can come into conflict with autonomy and beneficence, so the ethical advantage of synthetic data cannot be assumed without careful analysis [169].

5.1.2 Confidentiality and Professional Secrecy

The terms "privacy" and "confidentiality" are not synonymous. Privacy refers to the person, confidentiality to information: someone who follows a patient into a clinic violates the patient's privacy, whereas someone who opens the patient's record without authorization violates confidentiality [170]. The same chapter of the Shortliffe and Cimino textbook states the basic idea for the medical student: protecting confidentiality is not a favor done to the patient but a duty, regardless of the disease and of the medium on which the information is stored [170]. The reason is also practical: people who fear disclosure seek medical help less often, which increases the risk that infections will spread and diseases will remain untreated [170].

In Romania, Law No. 46/2003 on patients' rights (Legea nr. 46/2003 a drepturilor pacientului) provides that all information on the patient's condition, the results of investigations, the diagnosis, the prognosis, and the treatment is confidential even after death (Art. 21). This information may be disclosed only with the patient's explicit consent or where the law expressly requires it (Art. 22), but consent is no longer required when the information is needed by other accredited providers involved in the treatment (Art. 23) [171]. The GDPR explicitly links the processing of medical data to professional secrecy: the data may be processed for the purposes of diagnosis and care by or under the responsibility of a professional subject to the obligation of professional secrecy (Art. 9(3)) [118]. Confidentiality is not absolute. The care team has legitimate access to the record, and certain communicable diseases are reported to the public health authorities [170].

5.1.3 Informed Consent, from Form to Platform

Informed consent requires adequate information, understanding, and a free decision. For medical interventions, Law 46/2003 devotes a whole chapter to it (Chapter III) [171]. In research, the discussion has moved toward the electronic form (eConsent). A systematic review of 35 studies with 13,281 participants found no study in which the paper form gave better results than the electronic version for comprehension, acceptability, or usability [172].

The reuse of clinical data for research has produced two alternative models. **Broad consent** covers a category of future research rather than a single study. In Germany, the Medical Informatics Initiative developed a GDPR-compliant national broad consent with mandatory safeguards: independent review of every project, protection against unintended disclosure, and public transparency [173]. **Dynamic consent** uses a digital platform through which participants update their preferences over time and can withdraw consent. A scoping review identified multiple options and continuous communication as facilitators, and consent fatigue, the digital divide, technical complexity, and security risks as barriers [174]. A 2026 legal analysis shows that the European Health Data Space shifts the governance of secondary use toward institutional authorization, through data access bodies, which reduces the role of individual consent for many large-scale uses [175].

5.2 The GDPR for Physicians

5.2.1 What Data Concerning Health Are

Regulation (EU) 2016/679 (GDPR) defines **personal data** as any information relating to an identified or identifiable natural person, who can be identified, directly or indirectly, for example

by reference to a name, an identification number, location data, or factors specific to the physical, physiological, or genetic identity of that person (Art. 4(1)). **Data concerning health** are personal data related to the physical or mental health of a natural person, including the provision of healthcare services, which reveal information about his or her health status (Art. 4(15)) [118]. The definition is broad: an oncology appointment or a prescription reveals health status.

Data concerning health belong to the special categories of personal data, whose processing is prohibited in principle (Art. 9(1)). The prohibition is lifted only in the situations listed in Art. 9(2). For physicians, the most relevant are point (h) (preventive or occupational medicine, medical diagnosis, the provision of healthcare, and the management of healthcare services), point (i) (reasons of public interest in the area of public health), point (j) (scientific research, with the safeguards of Art. 89), point (c) (the vital interests of a person who is incapable of giving consent), and point (a) (explicit consent) [118]. In addition, every processing operation needs a legal basis under Art. 6(1), for example a legal obligation of the controller or a task carried out in the public interest. Figure 5.1 summarizes the reasoning.

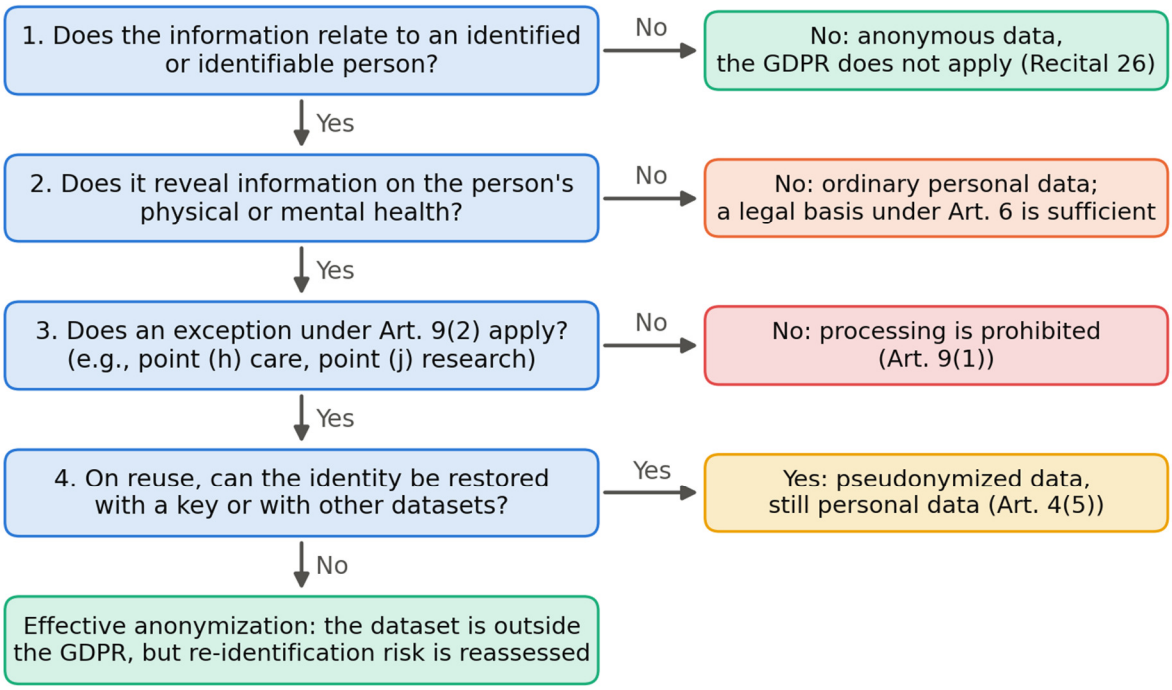


Figure 5.1. Decision tree for classifying a piece of information: personal data, data concerning health, legal basis, pseudonymized or anonymous data. Based on Regulation (EU) 2016/679, Art. 4, Art. 9, and Recital 26 [118].

A frequently overlooked detail: consent is not the usual legal basis for routine care. The hospital processes the data of an inpatient because they are necessary for diagnosis and treatment (Art. 9(2)(h)), not because the patient ticked a box. Consent to a medical intervention, as required by Law 46/2003, is a separate legal institution.

5.2.2 Principles and Rights

Article 5 lists the principles of processing: lawfulness, fairness and transparency, purpose limitation, data minimization, accuracy, storage limitation, integrity and confidentiality, to which is added accountability, the controller's responsibility for demonstrating compliance [118]. **Purpose limitation** means that data collected for care are not further processed in an incompatible

manner, research with appropriate safeguards being considered compatible. **Data minimization** requires that an export for a study contain only the necessary variables. The rights of the patient as a data subject are presented in Table 5.1.

Table 5.1. Rights of the data subject under the GDPR, with examples from medical practice. Source: [118,176].

Right (GDPR article)	What the patient may request	Note for the physician
Information (Art. 13-14)	who processes the data, for what purpose, for how long	information notice at admission
Access (Art. 15)	a copy of the data in the record	also Law 46/2003, Art. 24
Rectification (Art. 16)	correction of inaccurate data	clinical errors are corrected with an audit trail
Erasure (Art. 17)	erasure of the data	limited by the legal obligation to archive
Restriction (Art. 18)	temporary blocking of processing	e.g., while a dispute is pending
Portability (Art. 20)	the data in a structured, machine-readable format	applies to processing based on consent or contract
Objection (Art. 21)	cessation of a processing operation	derogations for research (Law 190/2018, Art. 8)
Automated decisions (Art. 22)	not to be subject to a decision based solely on automated processing	Law 190/2018, Art. 3(1)

5.2.3 The DPO, Impact Assessment, and Personal Data Breaches

A public hospital must designate a **data protection officer** (DPO), both because it is a public body and because it processes special categories of data on a large scale (Art. 37(1)(a) and (c)) [118]. The DPO informs and advises the controller and the employees, monitors compliance, and acts as the contact point for the supervisory authority, which in Romania is the National Supervisory Authority for Personal Data Processing (Autoritatea Națională de Supraveghere a Prelucrării Datelor cu Caracter Personal, ANSPDCP).

Before processing that is likely to result in a high risk, the controller carries out a **data protection impact assessment** (DPIA). The Regulation explicitly requires it for processing on a large scale of special categories of data (Art. 35(3)(b)) [118], for example when a regional disease registry is set up.

When a **personal data breach** occurs (unauthorized access, loss, encryption by ransomware), the controller notifies the supervisory authority without undue delay and, where feasible, not later than 72 hours after having become aware of it. Notification may be omitted only if the breach is unlikely to result in a risk to individuals (Art. 33) [118]. If the risk is high, the affected patients are also informed (Art. 34). For the physician, the practical rule is simple: a lost laptop, an e-mail with results sent to the wrong recipient, or a suspicious access is reported to the DPO immediately, because the deadline runs from the moment the organization becomes aware of it.

5.2.4 Law 190/2018

Law No. 190/2018 lays down the national measures implementing the GDPR. Art. 3(1) allows genetic, biometric, or health data to be processed for automated decision-making or profiling only with the explicit consent of the data subject or on the basis of express legal provisions, with protective measures. Art. 3(2) prohibits data processed for the purpose of ensuring public health from being subsequently used by third parties for other purposes. Art. 4 sets conditions for the processing of the personal numeric code (CNP). Art. 8 allows, for scientific research, derogations

from the rights of access, rectification, restriction, and objection, only if exercising these rights would make the research impossible or seriously impair it, and only with the safeguards of Art. 89 of the GDPR. For public bodies, and therefore for public hospitals, Art. 13-14 provide first for a warning with a remediation plan and then, if the plan is not carried out, for fines of 10,000 to 100,000 lei, or of 10,000 to 200,000 lei for infringements of the basic principles of processing and of the rights of data subjects [176].

5.2.5 Anonymization and Pseudonymization

Pseudonymization is the processing of personal data in such a manner that they can no longer be attributed to a specific person without the use of additional information, provided that such additional information is kept separately and is subject to technical and organizational measures (Art. 4(5)). Pseudonymized data remain personal data. By contrast, **anonymous** data no longer relate to an identifiable person and fall outside the GDPR (Recital 26), but the test is demanding: account must be taken of all the means reasonably likely to be used by the controller or by another person, including the cost, the time required, and the available technology [118]. The differences are summarized in Table 5.2.

Table 5.2. Anonymization and pseudonymization: a practical comparison. Source: [118,177].

Criterion	Pseudonymization	Anonymization
Definition	identity replaced by a code, key kept separately	link to the person irreversibly removed
GDPR status	personal data (Art. 4(5))	outside the GDPR (Recital 26)
Reversibility	yes, by the key holder	no (if successful)
Usefulness for research	high; allows linkage and follow-up over time	reduced by generalization and suppression
Example	registry with a study code instead of the CNP	published table with age groups, without rare data

Pseudonymization is recommended by the GDPR as a security measure (Art. 32) and as a safeguard for research (Art. 89) [118]. In practice, the linkage key is kept separately by a trusted person or service, and the researcher receives only the code.

Genuine anonymization is harder to achieve than it seems. A rare diagnosis combined with age, sex, and postal code can act as a fingerprint of the patient even after the name has been removed [170]. A review of anonymization techniques showed, through a reconstruction attack, that anonymization is necessary but not sufficient to protect patients [177]. Medical images raise risks of their own. On the public ChestX-ray14 data set (112,120 radiographs from 30,805 patients), a neural network recognized whether two chest radiographs belonged to the same person with an area under the ROC curve (AUC) of 0.994, even when the examinations were more than ten years apart [178].

Nor should the risk be overestimated. On data from more than one million patients, a model that takes into account the real resources of the attacker showed that, compared with the most pessimistic scenario (median re-identification risk 0.987), the median risk reduction was 90.1% when the attacker used voter lists and 100% when the attacker used social media posts [179].

5.3 The European Health Data Space

Regulation (EU) 2025/327 on the European Health Data Space (EHDS) entered into force on March 26, 2025, and applies, in general, from March 26, 2027 [23]. It complements the GDPR and distinguishes two uses of the same data.

Primary use means the use of data for care, including in another Member State. Patients have the right to access the priority categories of data (Art. 3) and may restrict health professionals' access to certain information (Art. 8), and Member States may introduce a right to opt out of primary use (Art. 10). The priority categories are patient summaries, electronic prescriptions and electronic dispensations, medical imaging studies with their reports, test results, and discharge reports (Art. 14). The obligations apply from March 26, 2029, for patient summaries, electronic prescriptions and electronic dispensations, and from March 26, 2031, for the other categories [23]. Electronic health record systems must include a European interoperability software component and a European logging software component (Art. 25), which is precisely the audit log discussed in Section 5.4 [23]. The MyHealth@EU infrastructure and cross-border exchange are described in Chapter 8.

Secondary use means the reuse of data for research, public health, statistics, policy-making, or education (Art. 53). Chapter IV of the Regulation applies from March 26, 2029, and Art. 51 lists 17 minimum categories of data, from data in electronic health records to registries and biobanks [23]. Each Member State designates one or more **health data access bodies**, which issue data permits (Art. 55, Art. 68), and a national contact point connected to the European HealthData@EU platform (Art. 75). Data are made available only to the extent necessary for the purpose, in anonymized form when the purpose can be achieved in that way (Art. 66), and only within a secure processing environment that restricts access to authorized persons and minimizes the risk of unauthorized copying or extraction (Art. 73) [23].

Two safeguards are important for patients. Art. 54 prohibits the use of data for decisions detrimental to the person, for exclusion from insurance or credit, for advertising, or for developing harmful products, such as tobacco products. Art. 71 gives every person the right to opt out, at any time and without stating reasons, of the secondary use of his or her personal electronic health data, and the decision is reversible [23].

5.4 Information Security

5.4.1 The Confidentiality-Integrity-Availability Triad

Information security is described by three properties. **Confidentiality** means that data are accessible only to authorized persons. **Integrity** means that data are accurate and complete and have not been altered without authorization: a dose changed in a prescription is an integrity problem. **Availability** means that data and systems can be used when needed, and in healthcare "when" may mean three o'clock in the morning in the on-call room. The GDPR adopts the same triad and requires the controller to take measures ensuring the ongoing confidentiality, integrity, availability, and resilience of processing systems (Art. 32) [118]. There are trade-offs among the three properties. Every access restriction has a cost, sometimes explicit (security equipment and software) and sometimes implicit, through the time demanded of users [170].

5.4.2 Threats

The first report of the European Union Agency for Cybersecurity (ENISA) devoted to the health sector analyzed 215 publicly reported incidents between January 2021 and March 2023. Health care providers were targeted in 53% of the incidents, and hospitals in 42%. Ransomware, which encrypts data and demands a ransom, accounted for 54% of threats, and patient data, including electronic health records, were the most targeted asset (30%) (Figure 5.2) [180]. The same report notes that only 27% of the health organizations surveyed had a dedicated ransomware defense program and that 40% had no awareness program for non-IT staff [180].

The entry point is often a message. In the hospitals analyzed by Jalali and colleagues (397 employees whose questionnaire responses were compared with their actual clicks in test campaigns), the stated intention to comply with the rules did not predict behavior, and the only factor associated with clicking on a phishing link was workload [181]. The insider threat is just as real. Systems that make access easy for clinicians also make it easy for the curious employee who has no need for the information [170]. In a survey of 41 professionals from Spanish medical centers, 97.6% authenticated with a password only, and more than 40% sent sensitive data by personal e-mail or messaging applications [182].

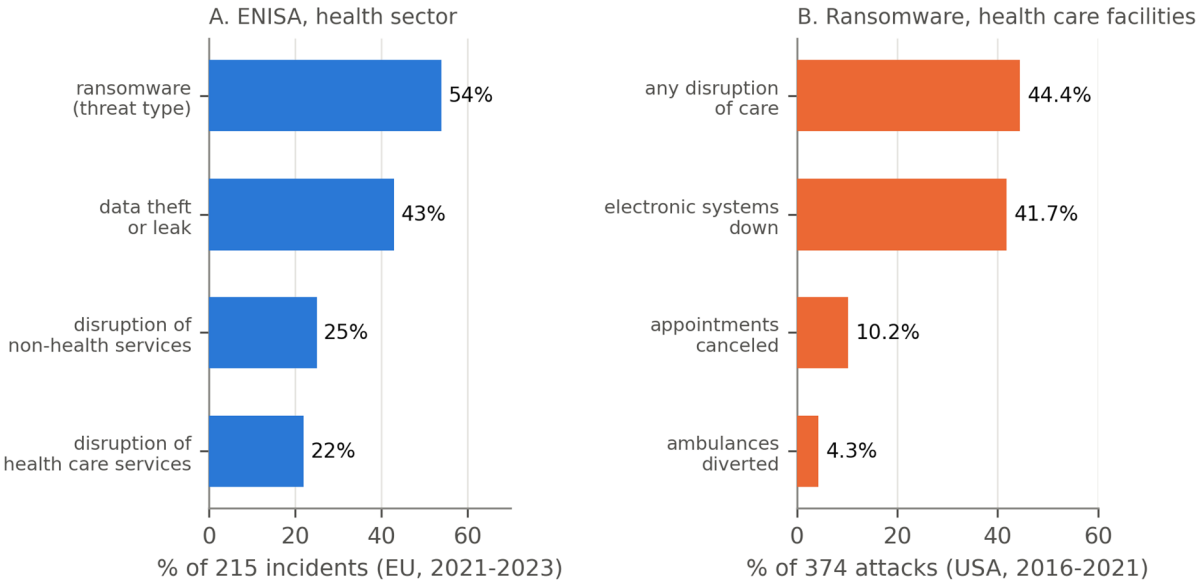


Figure 5.2. Security incidents in healthcare, real data. A: share of ransomware and of the main consequences in 215 incidents in the EU and neighboring countries, January 2021-March 2023 (the percentages do not add up, as one incident may fall into several categories). B: consequences for care in 374 ransomware attacks on US healthcare delivery organizations, 2016-2021. Data from [180,183].

5.4.3 When an Attack Becomes a Clinical Event

Between 2016 and 2021, 374 ransomware attacks on US healthcare delivery organizations exposed the data of nearly 42 million patients, and the annual number doubled, from 43 to 91. Almost half (44.4%) disrupted care, most often through electronic system downtime (41.7%), canceled appointments (10.2%), or ambulance diversion (4.3%). Over time, attacks were less and less often resolved by restoring from backups [183].

The effects also spread to neighboring hospitals. When a hospital network in the San Diego area was attacked in May 2021, the emergency departments of a neighboring academic center, not itself affected by the attack, received on average 251.4 patients per day, compared with 218.4 before, and stroke code activations rose from 59 to 102 [184]. In the two hospitals of the same

center, cardiac arrests rose from 21 to 38 during the attack period, and survival with a favorable neurological outcome after out-of-hospital cardiac arrest fell from 40.0% to 4.5% [185]. These are observational studies, with small numbers of events, which show associations, not causality. They nevertheless share a common conclusion: a cyberattack should be treated as a regional disaster.

Romania in 2026

- In February 2024, an attack with the Backmydata ransomware, from the Phobos family, encrypted the servers running the HIPOCRATE hospital information system in several hospitals. On February 12, the National Cyber Security Directorate (Directoratul Național de Securitate Cibernetică, DNSC) confirmed 21 affected hospitals, including the Timişoara Institute of Cardiovascular Diseases, and another 79 facilities had been disconnected from the internet for investigation [186].
- The attackers demanded 3.5 BTC (about 157,000 euros). The DNSC recommended not paying the ransom and not contacting the attackers. On February 15, the DNSC alert referred to 26 hospitals [186].
- Most of the affected hospitals had recent backups, 1-3 days old, with a single exception of 12 days, which, according to the DNSC, could facilitate the restoration of data and services [186].
- The preamble of Government Emergency Ordinance (OUG) No. 155/2024, which transposes NIS2, invokes this incident, which "affected 26 hospitals nationwide through a managed service provider" (our translation) [187].

5.4.4 NIS2 and Its Transposition in Romania

Directive (EU) 2022/2555 (NIS2) has applied since October 18, 2024, and includes health among the sectors of high criticality: healthcare providers, reference laboratories, and manufacturers of medicines and of critical medical devices [188]. Large entities in these sectors are "essential" entities [188]. They must implement cybersecurity risk-management measures (Art. 21), and their management bodies approve these measures and can be held liable for infringements (Art. 20) [180,188]. A significant incident is reported in three stages: an early warning within 24 hours, an incident notification within 72 hours, and a final report within one month at the latest (Art. 23) [188]. In Romania, the Directive was transposed by OUG No. 155/2024, published on December 31, 2024. The DNSC is the competent authority and the national computer security incident response team (CSIRT), and the deadlines of 24 hours, 72 hours, and one month are taken over in Art. 15(7) [187]. The same breach may therefore trigger two parallel notifications: to the DNSC under NIS2, and to the ANSPDCP under the GDPR, if personal data are affected.

5.4.5 Basic Measures

The ISO 27799 standard, whose third edition was published in December 2025, adapts the controls of ISO/IEC 27002:2022 to healthcare [189]. The recommendations below, summarized in Figure 5.3, are taken from the draft submitted to ballot in 2025 [190].

- **Multifactor authentication.** For systems that process personal health data, at least two factors are recommended, for example a password plus a code on a phone or a card [190].
- **Least privilege.** Access is granted by role and limited to the needs of the role, and users access a record only if they belong to the care team or have another legitimate task. Temporary access for students, residents, and locum physicians is closed when they leave [190].

- **Audit logs.** Every consultation and modification is recorded. The logs later show who accessed the record and deter curiosity, especially when the sanctions for violations are known and enforced [170].
- **Backups.** Health data are backed up in encrypted form, and offline or immutable copies are recommended against ransomware [190]. A backup that has never been test-restored remains a hypothesis.
- **Continuity plan.** Paper forms and procedures for working without the system are prepared in advance, and the data recorded during the outage are then re-entered with additional checks [190].

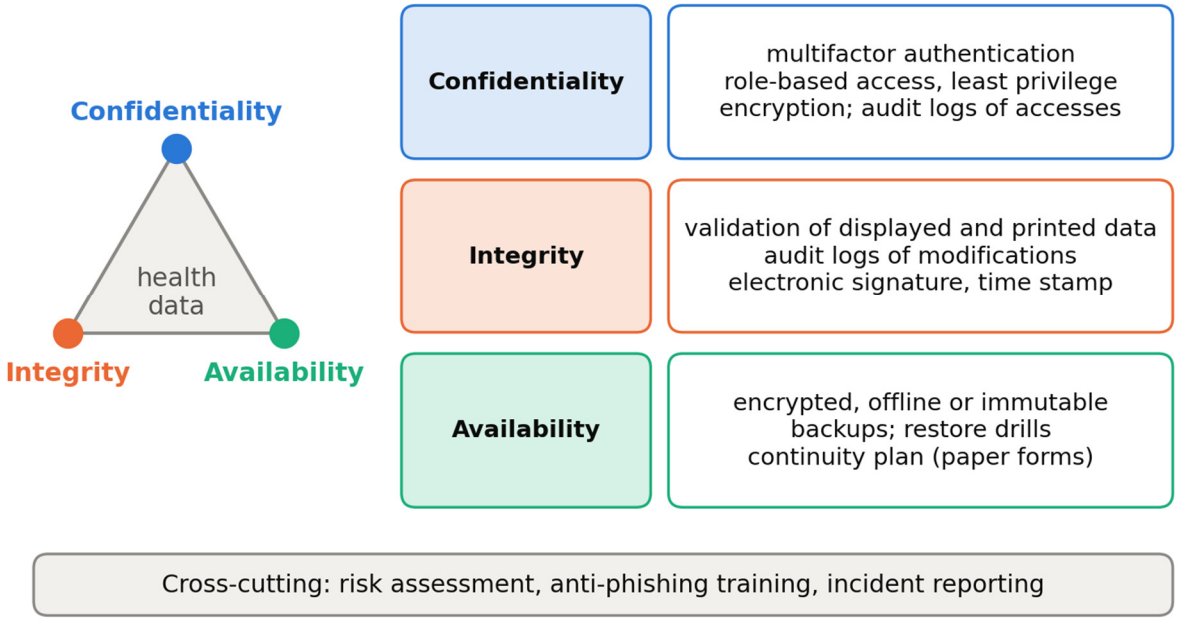


Figure 5.3. The confidentiality-integrity-availability triad and the controls that support it in a healthcare facility. Based on [118,190].

5.5 Equity and Algorithmic Bias

An algorithm can be accurate on average and still be inequitable. **Algorithmic bias** means a systematic error whose consequences fall disproportionately on a particular group of patients. The risk arises at every stage, from the choice of the problem and of the training data to use after deployment, and models can amplify existing inequities [191].

Pulse oximetry is the most frequently cited example, even though it does not involve AI. In 10,789 paired measurements from the University of Michigan, an arterial oxygen saturation below 88% that went undetected by a pulse oximeter reading 92-96% (occult hypoxemia) occurred in 11.7% of the measurements in Black patients, compared with 3.6% in White patients. In the multicenter cohort, the proportions were 17.0% and 6.2% [192]. A 2024 meta-analysis of 15 studies with 207,464 hospitalized patients confirmed a higher prevalence of occult hypoxemia in Black patients (prevalence ratio 1.67, 95% CI 1.47-1.90) and recommended that clinicians exercise increased caution in patients with darker skin pigmentation [193]. The lesson for informatics is direct: a sensor calibrated on a homogeneous population produces systematically erroneous data for other groups, and any model built on those data inherits the error.

The second example comes from imaging. State-of-the-art classifiers for chest radiographs, trained on three large data sets, selectively underdiagnosed underserved groups (women, Black patients, patients of low socioeconomic status), with higher rates for intersectional subgroups, for example Hispanic women [194]. Classic formulas can also be corrected. The 2021 equations for estimated glomerular filtration rate no longer use race, and those based on creatinine and cystatin C were more accurate and reduced the differences between groups compared with the single-marker equations [195].

Bias mitigation starts with the data. The STANDING Together consensus recommendations, developed with input from more than 350 representatives from 58 countries, call for transparent documentation of the composition and limitations of data sets and for evaluation of performance across population subgroups [196]. Detecting bias sometimes requires precisely sensitive data, such as ethnic origin. Regulation (EU) 2026/1744, which amends the AI Act, introduced an Art. 4a that allows, strictly to the extent necessary and with safeguards, the processing of special categories of personal data for the detection and correction of bias [197]. The other requirements for high-risk AI systems are discussed in Chapter 13 [22].

5.6 Academic Integrity

5.6.1 Plagiarism and Self-Plagiarism

The Law on Higher Education No. 199/2023 (Legea învățământului superior nr. 199/2023) defines **plagiarism** as the presentation, as one's own creation, of texts, ideas, demonstrations, data, theories, results, or methods taken from the works of other authors, without acknowledging this and without reference to the original sources. **Self-plagiarism** is the republication of substantial parts of one's own previous publications, including translations, without proper citation of the original (Art. 169(d) and (e)) [198]. The definition also covers ideas, not only copied sentences.

The phenomenon is not rare even in the training of health professionals. A meta-analysis of 31 studies with 9,175 nursing students found plagiarism to be the most frequent form of academic dishonesty (55.3%). The most common form was paraphrasing without reference to the source (39.5%), and plagiarism was a predictor of clinical misconduct [199]. Anti-plagiarism software calculates a similarity score. An identical passage may be a correct quotation, with quotation marks and a source. Nor is the score an adequate index of plagiarism: there is no threshold above which similarity means plagiarism, the software also flags common phrases, and the matches must be examined carefully by an evaluator [200].

5.6.2 Correct Citation

Citation has two functions: it acknowledges the contribution of others, and it allows the reader to verify the claim. Medical journals use the Vancouver style, based on the recommendations of the International Committee of Medical Journal Editors (ICMJE) and on the formats of the US National Library of Medicine [201]. Sources are numbered in the order of their first appearance in the text, and the formats are practiced in Practical Session 5.

5.6.3 Generative AI in Student Work

Generative AI tools can correct language or summarize an article, but they can also invent references. The ICMJE Recommendations, updated in January 2026, require authors to disclose, in the cover letter and in the appropriate section of the work, how they used AI-assisted

technologies. Chatbots cannot be authors, because they cannot be responsible for the accuracy, integrity, and originality of the work. Authors must be able to assert that text and images produced with AI contain no plagiarism, and AI-generated material cannot be cited as a primary source [202]. The Committee on Publication Ethics (COPE) adopted the same position in February 2023 and requires a description of the tool and of how it was used in the methods section [203].

Publishers' policies are not uniform. In 2023, 24% of the 100 largest publishers and 87% of the 100 highest-ranked journals had guidance on generative AI. Almost all prohibited AI authorship, but the form of disclosure varied widely [204]. Detectors of generated text do not offer a reliable solution either. They frequently classify texts written by non-native speakers of English as AI-generated [205]. For the student, the working rule is the same as for the author of an article: AI may assist, but it cannot sign. Its use is disclosed, every claim is verified against the original source, and every reference is actually looked up in PubMed or on the issuer's website.

Common Pitfalls

- "I deleted the name and the CNP, so the data are anonymous." Age, postal code, and a rare diagnosis can identify the patient, and the data remain personal data.
- Confusing informed consent to an intervention with the GDPR legal basis for care. Routine care is based on Art. 9(2)(h), not on a consent form.
- Sending patients' results or images through personal messaging applications. This is a breach of confidentiality and often of hospital policy.
- Postponing the reporting of an incident until "things become clear". The 72-hour deadline runs from the moment the organization became aware of it.
- Leaving the session open on the ward computer or using the same password everywhere.
- Believing that a low similarity score proves originality or that text produced by a chatbot need not be disclosed.

Key Points

- Data concerning health are a special category: their processing is prohibited, except in the cases listed in Art. 9(2) of the GDPR.
- Pseudonymized data remain personal data. Genuine anonymization is hard to achieve and must be demonstrated.
- A personal data breach is notified to the authority within 72 hours at most, and patients are informed when the risk is high.
- Confidentiality, integrity, and availability are protected together: multifactor authentication, least privilege, audit logs, offline backups.
- A cyberattack on a hospital is a clinical event, not merely a technical one.
- Algorithmic bias arises from data and from design decisions. Performance is checked across subgroups.
- AI tools cannot be authors. Their use is disclosed, and sources are verified manually.

Self-Assessment Questions

1. (SA) According to Art. 4(5) of the GDPR, pseudonymized data:

a) are no longer personal data
b) remain personal data, because they can be attributed to the person with the help of additional information
c) are identical to anonymous data
d) may be published without any restriction
e) cannot be used in research

2. (SA) On a Friday evening, a hospital discovers that a server holding laboratory results has been encrypted by ransomware. The deadline by which, where feasible, it must notify the supervisory authority is:

a) 24 hours from the time of the attack
b) 72 hours from the moment it became aware of the breach
c) 30 days from the end of the investigation
d) only if patients file a complaint
e) 72 hours, counting working days only

3. (SA) The typical GDPR legal basis for processing the data of an inpatient for the purposes of diagnosis and treatment is:

a) consent to marketing
b) the legitimate interest of the hospital
c) Art. 9(2)(h), processing necessary for healthcare, subject to the obligation of professional secrecy
d) data made public by the patient
e) no legal basis is required

4. (MA) Which of the following measures primarily protect the confidentiality of the data in the electronic health record?

a) multifactor authentication
b) role-based access, following the principle of least privilege
c) an uninterruptible power supply
d) audit logs of accesses, reviewed periodically
e) server redundancy

5. (MA) Regarding algorithmic bias and academic integrity, which of the following statements are correct?

a) pulse oximeters missed hypoxemia (occult hypoxemia) more often in Black patients, according to studies from 2020-2024
b) a generative AI tool may be listed as a coauthor if it wrote part of the text
c) a low similarity score proves that a paper contains no plagiarism
d) self-plagiarism includes republishing substantial parts of one's own work without citing the original
e) detectors of AI-generated text may misclassify texts written by non-native speakers of English

CHAPTER 6

Chapter 6. Biomedical Signals: Acquisition, Processing, and Assisted Interpretation

Objectives

By the end of this chapter, the student will be able to:

1. define the biomedical signal and classify signals by origin and by their mathematical properties.
2. describe the acquisition chain and calculate the minimum sampling frequency for a given signal.
3. explain aliasing, quantization, and the signal-to-noise ratio using simple numerical examples.
4. choose the appropriate filter type for baseline wander, 50 Hz power-line interference, and high-frequency noise.
5. calculate heart rate, SDNN, and RMSSD from a series of RR intervals.
6. critically appraise artificial intelligence algorithms for ECG and wearable devices (false positives, validation population, randomized trials).

6.1 The Biomedical Signal

6.1.1 Definition and Origin

A biomedical signal is a time-varying physical quantity, produced by the body or obtained through the interaction of a device with it, that carries information about a physiological function. The electrocardiogram (ECG) tracks the electrical activity of the heart. The electroencephalogram (EEG) tracks the electrical activity of the cortex, and the electromyogram (EMG) that of the skeletal muscles.

Classification by origin helps in understanding which sensor is needed:

- **bioelectric**: potentials generated by excitable membranes (ECG, EEG, EMG), picked up with electrodes.
- **mechanical**: pressures, displacements, forces (invasive arterial pressure, respiratory movements).
- **acoustic**: sounds produced by the heart, lungs, or vessels, picked up with a digital stethoscope or a microphone.
- **optical**: variations in the light absorbed or reflected by tissue, as in photoplethysmography (PPG) and pulse oximetry.
- **chemical**: concentrations of substances measured continuously (for example, interstitial glucose) with electrochemical sensors.

Most signals reach the computer as voltages, because a transducer converts the physical quantity (pressure, light) into a voltage. The pulse oximeter, for example, sends red and infrared

light through the tissue, and a photodiode records the variation in intensity with each arterial pulse. From these variations, after calibration, the device estimates oxygen saturation (SpO2) [206].

6.1.2 Useful Classifications

An **analog** signal can take any value within a continuous interval, whereas a **digital** signal takes only discrete values from a specified range [62]. Moreover, the digital signal is known only at the sampling instants. The voltage on the skin of the chest is analog. The file saved by the electrocardiograph is digital.

The second classification concerns predictability. A **deterministic** signal can be described completely by a formula (a 50 Hz sine wave produced by the power grid). A **random** (stochastic) signal can be described only statistically, by its mean, variance, and spectrum. Real physiological signals are mixtures: the ECG of an adult at rest looks almost **periodic**, with a waveform that repeats with every beat. However, the interval between beats varies continuously, and this variation itself has clinical value (Section 6.3.6). Biological signals are therefore, at most, quasi-periodic.

6.1.3 Examples and Typical Values

Table 6.1 brings together the signals most often encountered in clinical practice. The values are indicative.

Table 6.1. Characteristics of some common biomedical signals. Source: [206-211].

Signal	Origin	Amplitude or unit	Usual band or settings
ECG	bioelectric (myocardium)	approximately 1 mV on the skin	high-pass 0.05 Hz for the standard ECG, 0.5 Hz or more on monitors
Scalp EEG	bioelectric (cortex)	microvolts (beta rhythm: 10-20 µV)	display 0.5-70 Hz, minimum sampling 256 Hz
Surface EMG	bioelectric (motor units)	varies with the muscle and the effort	band-pass 20-500 Hz, dominant energy 20-150 Hz
PPG	optical (pulsatile blood volume)	relative units	main component 0.5-5 Hz
Invasive arterial pressure	mechanical (catheter and transducer)	mmHg	pulsatile components up to 20 Hz
SpO2	optical (red and infrared)	percent saturation	value derived from the pulsatile component

The ECG signal picked up on the skin is only about 1 mV, which is why it is sensitive to artifacts caused by movement of the patient or the electrodes and to power-line interference. Its quality has improved through analog and digital processing techniques [206]. For the EEG, with amplitudes on the order of microvolts, the problem is even greater. In EEG, the clinician works with frequency bands designated by Greek letters (Table 6.2). The posterior alpha rhythm, at 8-12 Hz, defines the normal background of the EEG of an awake adult. It is most evident with the eyes closed and attenuates when the eyes open [209].

Table 6.2. EEG frequency bands and their typical physiological context. Source: [208,209].

Band	Frequency	Typical physiological context
Delta (δ)	0.5-4 Hz	deep sleep
Theta (θ)	4-7 Hz	drowsiness, stages N1 and N2
Alpha (α)	8-12 Hz	relaxed wakefulness, eyes closed, occipital region
Sigma (σ)	12-16 Hz	sleep spindles in stage N2
Beta (β)	13-30 Hz	wakefulness, frontal and central regions
Gamma (γ)	30-80 Hz	sensory processing and cognition

6.2 The Acquisition Chain

6.2.1 From the Electrode to the Converter

Figure 6.1 shows the path of a signal from the biological source to its interpretation. The first element is the **sensor** (the electrode, the pressure transducer, the photodiode). It is followed by a **differential amplifier**, which amplifies the voltage difference between two electrodes and rejects the voltages common to both, such as the one induced by the power grid. The rejection capacity depends on the balance of the electrode impedances: the IFCN-ILAE minimum standard for EEG recommends impedances below 5 kΩ and considers values below 10 kΩ acceptable [208].

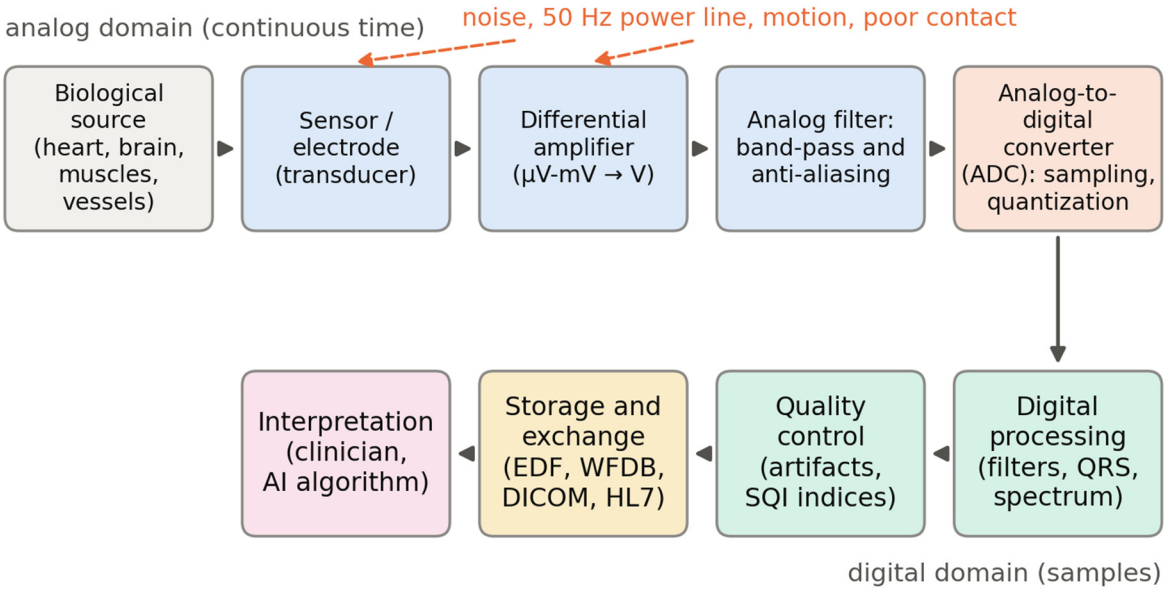


Figure 6.1. The acquisition and processing chain of a biomedical signal. The top row works with voltages that are continuous in time, the bottom row with numerical samples. Original diagram.

The **analog filter** placed before the converter has two roles. It removes components with no diagnostic value (slow drift, very high-frequency noise), and it limits the signal bandwidth to below half the sampling frequency, a role called anti-aliasing. The last analog block is the **analog-to-digital converter** (ADC), which transforms the sampled values of a continuous signal into a discrete numerical representation [62].

For invasive arterial pressure, the chain includes the catheter, the saline-filled tubing, and the transducer. Two sources of error depend on the staff: zeroing at heart level and the dynamic response, checked with the fast flush test, because small air bubbles distort the waveform [206].

6.2.2 Sampling and the Nyquist-Shannon Theorem

Sampling means measuring the value of the analog signal at equal time intervals. The number of measurements per second is the sampling frequency, fs, expressed in hertz. The sampling theorem (Nyquist-Shannon) states that a signal containing no components above the frequency fmax can be reconstructed without loss if fs > 2·fmax. The textbook by Shortliffe and Cimino calls this minimum threshold the Nyquist frequency [62]. Stated the other way around, the highest frequency that is correctly reproduced is fs/2 (at 256 Hz, this is 128 Hz) [208]. Some texts give the name Nyquist frequency to fs/2 itself, so the meaning of the term should be checked in each source.

If the rule is violated, **aliasing** occurs: a high-frequency component "disguises" itself as a false, lower frequency. Figure 6.2 shows the phenomenon for a 5 Hz sine wave. At 50 samples per second, its shape is preserved. At 6 samples per second, the points describe a 1 Hz wave that does not exist in the real signal.

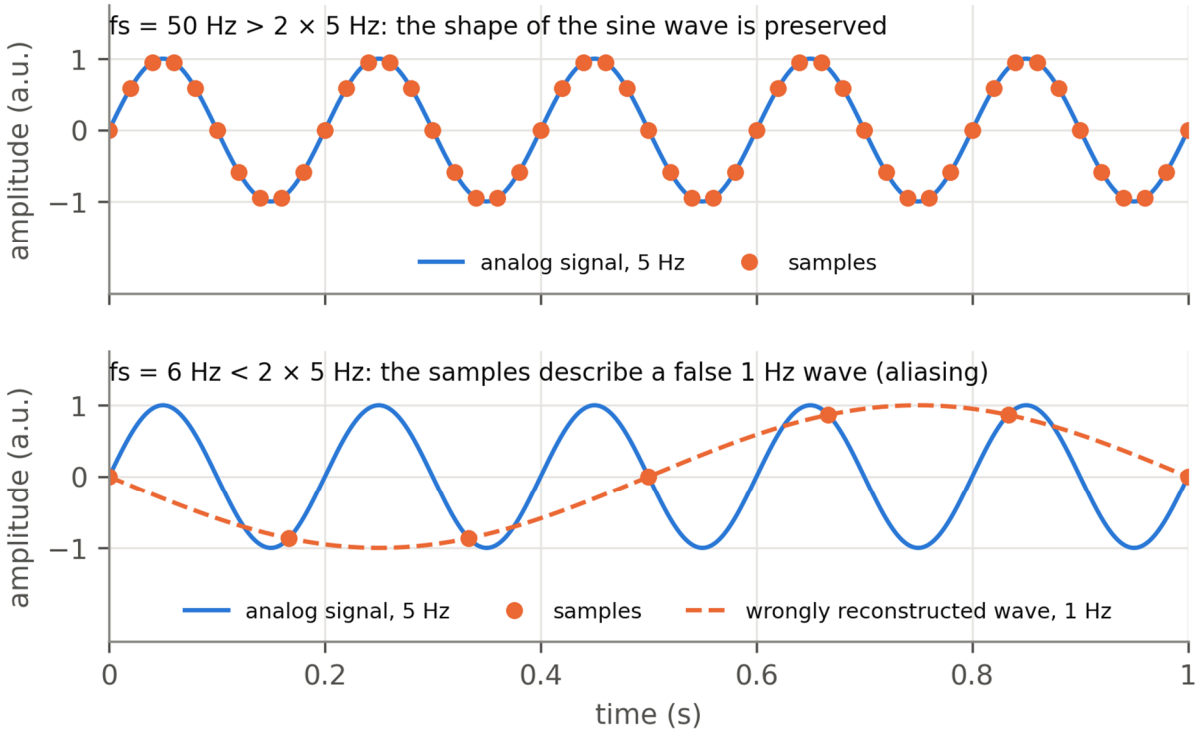


Figure 6.2. Correct sampling (top) and aliasing (bottom) for a 5 Hz sine wave. The apparent frequency with insufficient sampling is |5 - 6| = 1 Hz. Computed graph; synthetic, illustrative data.

Example

- **Minimum frequency for the EEG.** The clinical components of the routine EEG are displayed up to 70 Hz. The theorem requires fs > 140 Hz. Because of phase alignment, the IFCN-ILAE standard recommends sampling at no less than 2.5 times the maximum frequency and proposes a minimum of 256 Hz [208].
- **Aliasing of the power line.** 50 Hz interference sampled at 60 Hz without an anti-aliasing filter appears at |50 - 60| = 10 Hz, that is, exactly in the alpha band of the EEG. A technical artifact could be mistaken for a brain rhythm.
- **Timing precision.** At 250 Hz, the R wave is localized in time with an error of about ±4 ms. At 1000 Hz, the error falls to ±1 ms, which is why 1000 Hz is the recommended frequency, especially when the R wave must be linked to other synchronized events [212].

6.2.3 Quantization and Resolution

After sampling, each value must be written with a finite number of bits. This process is called quantization. An N-bit ADC divides the input range into 2^N levels. The quantization step (the resolution) is the range divided by the number of levels.

The classic example comes from the MIT-BIH Arrhythmia Database, available on PhysioNet: the signals were digitized at 360 samples per second per channel, with 11-bit resolution over a 10 mV range [213]. This gives 2^11 = 2048 levels and a step of 10 mV / 2048 ≈ 4.9 µV. A 1 mV QRS complex thus spans approximately 200 levels. With 16 bits, the same range would have 65,536 levels, with a step of approximately 0.15 µV.

Quantization adds a small error, of at most half a step. For a signal that occupies the full range, the theoretical ratio between the signal and the quantization noise is approximately 6.02·N + 1.76 dB. At 11 bits, this gives about 68 dB, well above biological and electrode noise.

The volume of data is the product of the number of channels, the sampling frequency, and the duration. A 24-hour Holter recording, with 2 channels at 360 Hz, has 2 × 360 × 86,400 ≈ 62 million samples, about 124 MB uncompressed at 2 bytes per sample.

6.2.4 The Signal-to-Noise Ratio

The signal-to-noise ratio (SNR) compares the power of the useful signal with the power of the noise: SNR = 10·log10(Ps/Pn), in decibels. When amplitudes are compared, the formula becomes 20·log10(As/An). An ECG with an amplitude of 1 mV over a noise of 0.1 mV has an SNR of 20·log10(10) = 20 dB. With a noise of 0.5 mV, the SNR drops to approximately 6 dB. The SNR estimated over short signal windows is one of the statistical indices used to decide automatically whether an ECG segment is usable [214].

6.3 Digital Processing

6.3.1 Filters

A digital filter modifies the spectrum of the signal: it lets some frequencies pass and attenuates others. Four types are used in clinical practice:

- **low-pass**: keeps the low frequencies and attenuates high-frequency noise (for example, muscle activity superimposed on the ECG).

- **high-pass**: removes very slow components, such as the baseline wander caused by breathing, sweating, or electrode movement.
- **band-pass**: the combination of the two, used in surface EMG in the typical range of 20-500 Hz [210].
- **notch (band-stop)**: removes a narrow band around a fixed frequency, usually the power-line frequency, 50 Hz in Europe and 60 Hz in North America [210].

Filtering comes at a price. An aggressive high-pass filter can distort the ST segment of the ECG. The standard 12-lead ECG uses a 0.05 Hz high-pass filter, whereas bedside monitors often use 0.5 Hz or more to reduce baseline wander. In a 2026 analysis of 150 patients, the 1.0 Hz filter increased the odds of a false ST-segment elevation of at least 1 mm in V1 almost tenfold (OR 9.14) [207]. The practical rule: a suspicious ST segment on the monitor is checked on a standard 12-lead ECG. Even medical informatics textbooks note that the filtering required for continuous monitoring "slightly distorts" the ECG [206].

Figure 6.3 shows the effect of a simple chain of filters on a synthetic ECG. After filtering, the R peaks can be detected with a simple threshold.

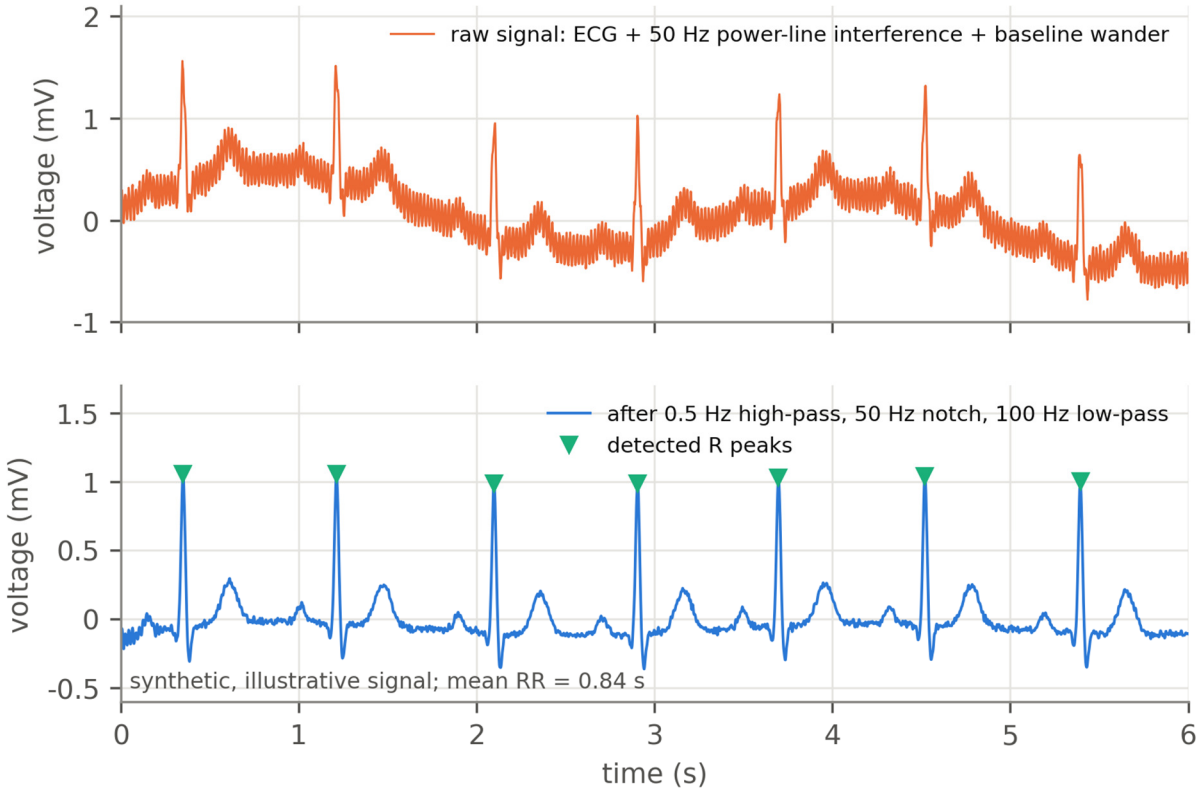


Figure 6.3. Synthetic ECG with 50 Hz interference and baseline wander (top), then after 0.5 Hz high-pass, 50 Hz notch, and 100 Hz low-pass filters (bottom), with automatically detected R peaks. Synthetic, illustrative signal generated from Gaussian waves.

6.3.2 Averaging

Some signals are much smaller than the background noise but repeat identically after a stimulus: visual or auditory evoked potentials are the typical example. If N responses are aligned to the moment of the stimulus and averaged, the synchronized component is preserved, while random noise decreases by a factor of √N. With 100 repetitions, the noise decreases tenfold, that is, the SNR increases by 20 dB. The method assumes an identical response at each repetition.

6.3.3 The Fourier Transform and the Spectrum

Any signal can be decomposed into a sum of sine waves of different frequencies. The discrete Fourier transform, computed efficiently with the fast Fourier transform (FFT) algorithm, takes the signal from the time domain to the frequency domain. The result, the power spectrum, shows how much energy the signal has at each frequency. The IFCN recommendations for resting-state EEG include band filtering and the calculation of the amplitude or power spectrum as basic steps in feature extraction [215].

On the spectrum, power-line interference appears as a narrow peak at 50 Hz, and baseline wander as energy concentrated below 0.5 Hz (Figure 6.4). The same figure shows that most of the energy of the synthetic ECG lies below 40 Hz.

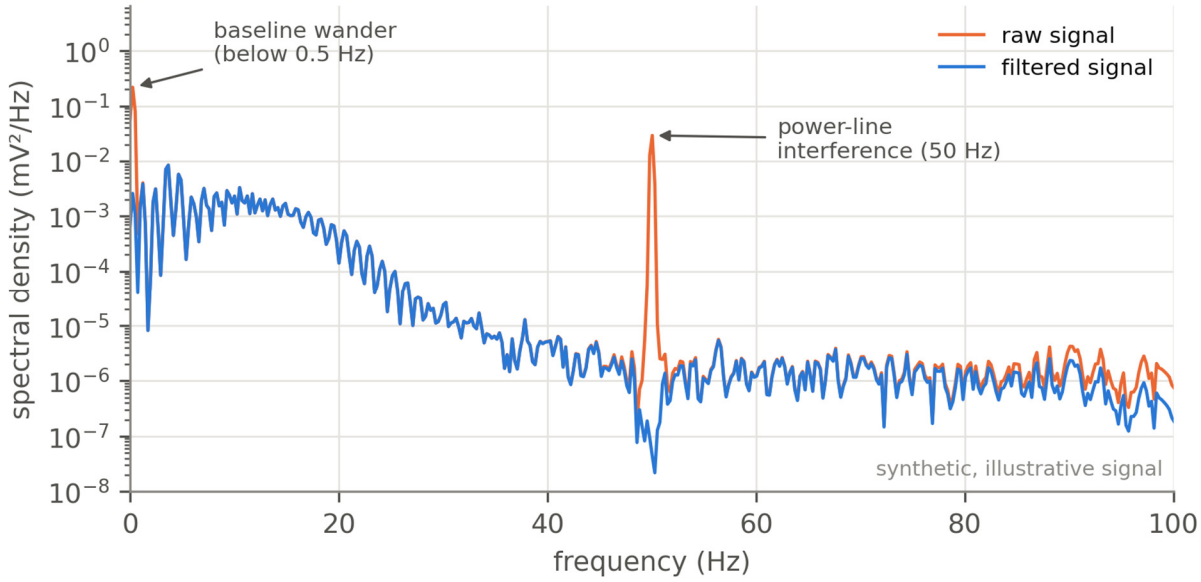


Figure 6.4. Power spectral density (Welch method) of the synthetic ECG in Figure 6.3, before and after filtering. The 50 Hz peak and the energy below 0.5 Hz disappear after filtering. Synthetic, illustrative signal.

6.3.4 Time-Frequency Analysis in Brief

The classic spectrum assumes a signal with constant properties, but an epileptic seizure starts abruptly and a sleep spindle lasts one second. Time-frequency analysis shows "which frequencies are present and when." The spectrogram applies the Fourier transform to short windows that slide in time, whereas the wavelet transform uses short windows for high frequencies and long ones for low frequencies. Time resolution and frequency resolution cannot both be arbitrarily fine at the same time.

6.3.5 QRS Complex Detection

The QRS detector is the first step of almost any automated rhythm analysis. A classic detector of the Pan-Tompkins type works in a few steps: it band-pass filters the signal to keep the energy of the QRS complex, differentiates the signal to emphasize steep slopes, squares the values, integrates them over a moving window on the order of 150 ms, and compares the result with an adaptive threshold. Free viewing programs such as EDFbrowser include such a detector and allow the RR intervals to be exported [216].

Detectors are not interchangeable. A 2026 study compared eight algorithms, Pan-Tompkins among them, on recordings from 25 volunteers at rest, under cognitive load, and during physical activity. No algorithm was the best under all conditions, and the agreement of heart rate

variability indices with the reference values ranged from almost perfect to none, depending on the detector and the acquisition mode [217].

6.3.6 Heart Rate Variability

The RR interval is the time between two successive R peaks. The instantaneous heart rate is HR = 60/RR (RR in seconds): an RR of 0.8 s corresponds to an HR of 75 beats per minute. Heart rate variability (HRV) describes the beat-to-beat fluctuations of these intervals. It is calculated only on NN intervals, that is, between normal sinus beats. Ectopic beats and artifacts must be identified and handled before the calculation, preferably with visual inspection after automated detection, because values flagged automatically as outliers are not always artifacts, and an ectopic beat cannot simply be "corrected" [212].

Table 6.3. Basic heart rate variability indices. Source: [212].

Index	Definition	Remark
SDNN (ms)	standard deviation of the NN intervals	overall variability, dependent on recording length
RMSSD (ms)	square root of the mean of the squared successive differences	short-term variability, abolished by vagal blockade
pNN50 (%)	percentage of successive differences greater than 50 ms	related to RMSSD
HF	spectral power in the respiratory band, typically 0.12-0.40 Hz	reflects respiratory sinus arrhythmia
LF	spectral power in the 0.05-0.15 Hz band (sometimes from 0.04 Hz)	controversial physiological interpretation

Example

- Five consecutive NN intervals: 800, 820, 790, 810, and 830 ms. The mean is 810 ms, so the mean HR = 60,000/810 ≈ 74 bpm.
- **SDNN**: the deviations from the mean are -10, 10, -20, 0, and 20 ms. The sum of squares is 1000; divided by n - 1 = 4, it gives 250, and the square root gives SDNN ≈ 15.8 ms.
- **RMSSD**: the successive differences are 20, -30, 20, and 20 ms. Their squares (400, 900, 400, 400) have a mean of 525, and the square root gives RMSSD ≈ 22.9 ms.
- Five intervals are too few for a clinical conclusion. The example shows only the mechanics of the calculation, which is taken up again in Practical Session 6.

For spectral analysis, the RR series, which is unevenly spaced in time, is resampled at equal intervals, preferably at 4 Hz [212]. Interpretation calls for caution: LF power is not a simple index of sympathetic activity, and the standard HF band may miss the very slow breathing of an adult or the faster breathing of children [212].

6.4 Artifacts and Signal Quality

6.4.1 Types of Artifacts

An artifact is any component of the signal that does not originate from the physiological phenomenon being monitored. For the ECG, the sources described in the literature include detached electrodes, movement, broken wires, external or implanted devices, and even other

physiological signals [218]. ECG artifacts can mimic ventricular tachycardia or acute myocardial infarction, which is why recognizing them is part of the physician's clinical competence [218].

The main categories are:

- **movement**: displacement of the electrode relative to the skin changes the contact potential and produces slow or abrupt oscillations. In EMG, movement artifacts typically lie between 0 and 20 Hz [210].
- **electrode contact**: dried gel, high or unbalanced impedance, an electrode that "pops" (electrode pop). High-impedance electrodes are more affected by sweating and movement [208].
- **the power line**: a narrow sinusoidal component at 50 or 60 Hz, induced by cables and equipment powered from the electrical grid [210].
- **muscle activity (EMG)** superimposed on the ECG or EEG: high-frequency noise, common in agitated patients and in those with shivering or tremor.
- **physiological artifacts in the EEG**: eye movements, the pulse, the ECG, and sweating. Some of them, such as EMG and eye movements, even help with interpretation [219].

In PPG, the motion artifact (approximately 0.01-10 Hz) overlaps the main band of the signal, 0.5-5 Hz, and can distort it, which is why it must be removed or corrected before analysis [211]. The pulse oximeter has similar limitations: movement, anemia, low peripheral perfusion, and cold skin. It is among the intensive care devices with the most false alarms [206].

6.4.2 Quality Control

Quality control begins at the bedside, with skin preparation and the securing of electrodes and cables. In EEG, the impedance of each electrode is checked at the start of the recording [208]. For invasive pressure, zeroing and the flush test depend on the staff, and the catheter-transducer system requires continuous supervision: drawing blood through the catheter, for example, temporarily makes the pulsatile waveform disappear [206].

Quality control continues in software, through signal quality indices (SQI): the relative power of the QRS complex, skewness, kurtosis, and the SNR. A 2022 analysis showed that these statistical indices, used with fixed thresholds, behave differently from one database to another, which is why the authors recommend adaptive thresholds or machine learning methods [214]. A systematic review of automated ECG quality assessment techniques published in 2012-2022 reported accuracies between 85.75% and 97.15%, but also almost no real-time testing in humans [220].

The clinical stake is alarm fatigue. Data cited in the textbook by Shortliffe and Cimino show between 1.6 and 14.6 alarms per patient per hour in intensive care, of which up to 90% are false. A few seconds of artifact on the ECG can trigger, for example, a false ventricular tachycardia alarm [206]. The proposed solutions include short alarm delays and the combination of several signals that estimate the same heart rate [206].

6.5 Formats and Open Databases

6.5.1 File Formats

Without metadata (sampling frequency, unit, gain, lead), the samples are almost unusable. Standard formats keep them together.

- **EDF and EDF+**. The European Data Format was published in 1992, and EDF+ in 2003. EDF+ can contain discontinuous recordings, annotations, stimuli, and events. The format is used mainly in clinical neurophysiology, sleep medicine, and cardiology, and it has become the de facto standard for EEG and polysomnography [221]. The IFCN-ILAE minimum standard recommends that EEG data be exportable to CSV, EDF, or DICOM [208].
- **WFDB** (Waveform Database) is a set of open standards and software for storing and analyzing physiological signals and annotations, developed at MIT starting in 1989. A WFDB record has a header file (.hea), one or more signal files (.dat), and annotation files, and the format underlies many of the databases on PhysioNet [222].
- **DICOM waveform**. The DICOM standard, known from imaging, also defines waveform objects: the 12-lead ECG, general and ambulatory ECG, hemodynamic waveforms, arterial pulse, respiration, routine scalp EEG, sleep EEG, EMG, and electrooculogram [223].
- **HL7 aECG** (annotated ECG) is an XML standard created in response to the 2001 FDA initiative on the digital ECG. Its initial purpose was the systematic assessment of waveforms and measurement points (for example, the QT interval) in clinical trials of noncardiac drugs [224].
- **ISO/IEEE 11073** is the family of standards for point-of-care medical device communication. The historical bedside variant (the Medical Information Bus) had poor commercial uptake [206]. Part 10103 (2025 edition) defines the terms for the interrogation summary of implantable cardiac devices (pacemakers, defibrillators, resynchronization devices, implantable monitors), transferred from the manufacturers' systems to the electronic health record [225].

6.5.2 PhysioNet, an Open Resource

PhysioNet, maintained by the MIT Laboratory for Computational Physiology, hosts databases of physiological signals, open-source software, and scientific competitions (PhysioNet Challenges) [226]. In 2026, the resource's team described it as a global platform for biomedical research, and this paper is the standard citation requested of users [227].

The MIT-BIH Arrhythmia Database contains 48 half-hour excerpts from 47 subjects, with approximately 110,000 beat annotations independently verified by at least two cardiologists. Access is open, under the Open Data Commons Attribution License [213]. Practical Session 6 uses this database and the LightWAVE viewer, which runs directly in the browser [228].

6.6 Interpretation Assisted by Artificial Intelligence

6.6.1 AI-Enabled ECG

Automated interpretation of the 12-lead ECG has been in use since the 1980s. The algorithms are fairly accurate, especially in healthy individuals, but disagreements with cardiologists can be clinically important, which is why a cardiologist usually confirms the result [206]. Artificial intelligence (AI) models based on deep learning have broadened the questions put to the ECG: left ventricular systolic dysfunction, hypertrophic cardiomyopathy, cardiac amyloidosis, aortic stenosis, or the risk of future atrial fibrillation. Before widespread adoption, a 2024 review calls for evidence of patient benefit and for implementations that do not overwhelm the system with false-positive results [229].

The reference example is EAGLE, a pragmatic cluster-randomized trial in which 120 primary care teams from 45 clinics or hospitals either received or did not receive the result of an ECG

algorithm that flagged the probability of a low ejection fraction (≤ 50%). A total of 22,641 adults without known heart failure were included. In the AI arm, new diagnoses of low ejection fraction within 90 days rose from 1.6% to 2.1% (OR 1.32). Among patients with a positive AI result (6% of the cohort), they rose from 14.5% to 19.5%. The total number of echocardiograms remained similar (19.2% vs. 18.2%) [230].

The idea has also been tested in other health systems. A pragmatic randomized trial of 13,631 inpatients in Taiwan, cared for by noncardiologist physicians, found more new diagnoses of low ejection fraction within 30 days (1.5% vs. 1.1%, HR 1.50). The use of echocardiography did not increase, but the echocardiograms ordered confirmed the disease more often (34.2% vs. 20.2%) [231]. In Nigeria, a randomized trial of 1,232 pregnant and postpartum women compared AI-guided screening with usual care. The AI-enabled digital stethoscope doubled the detection of systolic dysfunction (4.1% vs. 2.0%, OR 2.12). The model applied to the 12-lead ECG, however, had an effect in the same direction that did not reach statistical significance (3.4% vs. 2.0%, p = 0.125) [232]. The difference suggests that a model that performs well in one population does not automatically bring the same benefit in another.

AI can also extract information from the ECG that is invisible to the naked eye. In 1,936 older adults without known atrial fibrillation, the output of a model applied to an ECG in sinus rhythm predicted future atrial fibrillation with a C statistic of 0.69, equal to that of the CHARGE-AF clinical score. The combination of the two reached 0.72 [233]. A C statistic of 0.69 is modest: the model ranks risk; it does not make a diagnosis.

6.6.2 Atrial Fibrillation Detected by Smartwatches and Fitness Bands

Smartwatches use two technologies. Optical PPG measures the pulse continuously and flags an irregular rhythm. The single-lead ECG is recorded on demand, with a finger on the watch case. PPG sensors already track heart rate and rhythm, sleep, and exercise, and the field's 2023 roadmap describes the research needed for their clinical use [234].

A large study, the Fitbit Heart Study, enrolled 455,699 participants in the United States. Irregular rhythm notifications occurred in 1% of them and in 4% of those aged 65 years or older. In 1,057 of them, an ECG patch worn for one week confirmed atrial fibrillation in 32.2% of cases. When the notification recurred during ECG monitoring, the positive predictive value was 98.2% [235]. A subsequent analysis of the same cohort showed that the probability of confirmation depends on the duration of monitoring: 7.6% for a single 10-second ECG, 32.2% for a one-week monitor, and 60.8% for a four-week monitor [236].

A randomized trial with the Apple Watch (EQUAL), published in 2026, included 437 patients aged 65 years or older at high risk of stroke. After six months of monitoring with the watch, newly diagnosed atrial fibrillation was more frequent than in the standard care group (9.6% vs. 2.3%, HR 4.40) [237]. Accuracy meta-analyses give high figures. One from 2025, with 26 studies and 17,349 patients, found a sensitivity of 95% and a specificity of 97%, comparable between PPG and ECG models, but with very high heterogeneity between studies (I^2 of 95% and 99%) [238].

6.6.3 Limitations

High sensitivity and specificity do not automatically mean clinical benefit. The AF-SCREEN Collaboration notes that randomized trials of screening led by health professionals have not demonstrated a reduction in stroke. A suspicion raised by PPG is confirmed with an ECG, and an

ECG recorded by a consumer is read by a professional. Among the young early adopters of the technology, screening is unlikely to be cost-effective [239]. A 2026 analysis, explicitly hypothesis-generating, synthesized the randomized trials of wearable-type monitoring in asymptomatic people at high risk. Continuous monitoring increased the odds of an atrial fibrillation diagnosis (OR 2.84) and doubled the use of anticoagulants, without a clear reduction in adverse events, which suggests overdiagnosis [240].

The second limitation is the population. A model trained on a homogeneous population may perform differently in another. For the ECG model that detects low ejection fraction, derived from a cohort that was 96.2% non-Hispanic white, performance remained consistent across racial and ethnic groups (AUC 0.93-0.96). The authors nevertheless recommend reporting performance by age, sex, race, and ethnicity subgroups for any new tool [241]. For optical sensors, the problem is a physical one: melanin absorbs light. A 2024 meta-analysis found that pulse oximeters overestimate SpO2 in both lighter and darker skin, although this overestimation may not be clinically relevant. Pulse rate measured by wearable devices showed no systematic error related to pigmentation [242].

The third limitation is external validation: the results from the center that developed the model must be reproduced on other devices, in other hospitals, and in other populations. The Nigerian trial applied models developed at the Mayo Clinic and trained on data independent of those of the trial [232]. The Taiwanese trial, by contrast, tested a model previously developed by the same team, so it shows benefit in a different clinical setting, not external validation in the strict sense [231].

Continuous monitoring is not limited to cardiology. A systematic review by the author included 9 studies (778 patients) with wearable sensors after colorectal surgery and found that the step count in the first postoperative days was associated with length of stay, complications, and readmission, while calling for standardized definitions and multicenter validation [243].

6.7 Romania in 2026

As of September 27, 2026, the open databases on PhysioNet, the EDF and WFDB formats, and the free software described in this chapter can also be used in Romanian universities, under open licenses that require citation of the source. On the systems side, the terms of reference (caiet de sarcini) published by the Ministry of Health for the integrated information system of the anesthesia and intensive care departments (SIEC-ATI) cover 18 hospitals, 649 intensive care bed positions, and 275 anesthesia positions. The document requires the automatic recording, at intervals of no more than 90 seconds, of the parameters measured by the devices connected to the patient (vital signs monitor, ventilator, pumps), automatic alarm generation, and HL7/FHIR interfaces [244]. The requirement concerns periodic numerical values (trends), not waveforms sampled at hundreds of hertz, a difference that the student can explain after studying Section 6.2. This chapter does not assess the implementation status of the system.

Common Pitfalls

- Confusing the signal frequency with the sampling frequency: a signal with components up to 70 Hz requires more than 140 Hz.
- Believing that aliasing can be corrected after digitization. The false frequency can no longer be separated, which is why the anti-aliasing filter is placed before the converter.
- Interpreting an ST-segment elevation on the monitor without checking the high-pass filter setting.
- Calculating HRV on all RR intervals, including those around extrasystoles or artifacts.
- Treating a smartwatch notification as a diagnosis of atrial fibrillation. A suspicion raised by PPG is confirmed by ECG.
- Reading sensitivity and specificity as evidence of clinical benefit. Benefit is demonstrated by patient outcomes in randomized trials.

Key Points

- The sampling frequency must exceed twice the highest useful frequency. In practice, a larger margin is chosen.
- The number of bits of the converter sets the smallest voltage difference that can be distinguished.
- Filters clean the signal, but they can also distort it. Their settings are part of the interpretation.
- HRV is calculated on clean NN intervals. SDNN and RMSSD are the starting point.
- Open formats (EDF, WFDB, DICOM, HL7 aECG) keep together the samples and the metadata needed for interpretation.
- AI applied to the ECG has increased the number of diagnoses in randomized trials, but the prognostic benefit and the behavior in other populations must be demonstrated separately.

Self-Assessment Questions

1. (SA) The IFCN-ILAE minimum standard (2023) for routine EEG recommends a minimum sampling frequency of:

a) 70 Hz
b) 100 Hz
c) 140 Hz
d) 256 Hz
e) 1000 Hz

2. (SA) Sinusoidal 50 Hz interference is sampled at 60 Hz without an anti-aliasing filter. In the digital signal, it will appear as a component at:

a) 50 Hz
b) 110 Hz
c) 10 Hz
d) 60 Hz
e) 0 Hz

3. (SA) In the randomized EAGLE trial, giving family physicians access to the result of an AI algorithm applied to the ECG:

a) reduced one-year mortality
b) increased the proportion of new diagnoses of low ejection fraction (2.1% vs. 1.6%)
c) halved the total number of echocardiograms
d) replaced echocardiography as a diagnostic method
e) did not change the number of new diagnoses

4. (MA) Regarding heart rate variability, the correct statements are:

a) RMSSD is calculated from the differences between successive RR intervals
b) SDNN is the standard deviation of the intervals between normal sinus beats
c) ectopic beats are kept unchanged in the series used for the calculation
d) the HF band is typically defined between 0.12 and 0.40 Hz and reflects the influence of breathing
e) HRV can be calculated from the SpO2 values displayed by the pulse oximeter

5. (MA) Regarding the detection of atrial fibrillation with smartwatches, the correct statements are:

a) a suspicion based on PPG must be confirmed with an ECG
b) in the Fitbit Heart Study, the positive predictive value of a repeated notification during ECG monitoring was 98.2%
c) continuous screening with wearable devices clearly reduced adverse events in asymptomatic people
d) a 2025 meta-analysis found very high heterogeneity between accuracy studies
e) the probability of confirming atrial fibrillation does not depend on the duration of ECG monitoring

CHAPTER 7

Chapter 7. Digital Medical Imaging: DICOM, PACS, and Computer-Assisted Analysis

Objectives

By the end of this chapter, the student will be able to:

- calculate the uncompressed size of a CT series from the matrix, the bit depth, and the number of images;
- convert a value in Hounsfield units into a displayed gray level for a window defined by its level and width;
- describe the DICOM information model (patient, study, series, instance) and recognize the main attributes in a file header;
- explain the path of an examination through the RIS, the PACS, and the electronic health record, including the role of the IHE Scheduled Workflow profile;
- compare lossless and lossy compression from the standpoint of diagnostic safety;
- critically interpret a published result on artificial intelligence in imaging, taking into account external validation and automation bias.

7.1 The Digital Image

7.1.1 Pixel, Voxel, Matrix

To a computer, a medical image is an array of numbers. Each element of the two-dimensional array is called a pixel (*picture element*) and represents the intensity of a small square area of the image. The pixel corresponds to a volume element of the examined body, called a voxel (*volume element*), and a CT or MRI volume is represented as a three-dimensional array, that is, a stack of slices [24]. Voxel dimensions matter. If the slice thickness differs from the in-plane pixel size, the voxels are anisotropic, and coronal or sagittal reconstructions obtained from axial images lose sharpness. Most modern CT and MRI scanners, however, produce isotropic voxels [24].

The number of rows and columns forms the image matrix. It varies widely between modalities: from 64 x 64 pixels for some nuclear medicine images to approximately 5000 x 4000 pixels for mammography [68].

7.1.2 Bit Depth and Resolution

Each pixel stores a number using a given number of bits. With 8 bits, 2^8 = 256 values can be represented; with 12 bits, 4096; and with 16 bits, 65,536. For CT images, the DICOM standard requires 16 bits allocated per pixel, of which 12 to 16 bits actually hold the value [245]. Color images (scanned pathology slides, dermatology photographs) usually have three channels (red, green, and blue), each with its own depth.

The quality of a digital image is described by three quantities. Spatial resolution shows how well two nearby points can be separated and depends on the number of pixels per unit area. Contrast resolution shows how small an intensity difference can be and still be distinguished, and it is limited by the number of bits per pixel. Temporal resolution, finally, is the time needed to form an image: from about 30 images per second upward, the beating heart appears without blurring [24]. Microcalcifications on a mammogram demand high spatial resolution. An isodense liver lesion demands contrast resolution.

7.1.3 How Much Space a CT Series Takes: A Worked Example

The calculation takes three steps and is worth practicing, because it explains why hospitals need dedicated archives.

1. An axial CT image has a matrix of 512 x 512 = 262,144 pixels.
2. At 16 bits (2 bytes) per pixel, one image takes 262,144 x 2 = 524,288 bytes, or about 0.52 MB, without a header and without compression.
3. An axial chest-abdomen-pelvis series reconstructed at 1 mm over a length of 60 cm has roughly 600 images: 600 x 0.52 MB ≈ 315 MB.

The same examination is often reconstructed several times (soft-tissue kernel, lung kernel, coronal planes), and the volume per patient multiplies. Published reference values are on the order of 500 images and 250 MB for a CT examination, 63 MB for MRI, and 153 MB for a mammogram, with the caveat that the number of images and the resolution increase from year to year [68].

7.1.4 The Hounsfield Scale and the Display Window

The numerical values of CT pixels are expressed in Hounsfield units (HU). The scale is defined so that water has 0 HU and air approximately -1000 HU. In practice, values range from -1000 to approximately +3000 [24]. In the DICOM file, the stored value (SV) is converted into HU with the relation $HU = m \cdot SV + b$, where m is the Rescale Slope (0028,1053) and b is the Rescale Intercept (0028,1052) [245]. For example, with m = 1 and b = -1024, a stored value of 1064 corresponds to 40 HU.

The human eye cannot tell apart more than about 100 shades of gray. A monitor usually displays 256. If all of the approximately 4000 HU values were compressed linearly into 256 levels, each gray level would cover about 16 HU, and the small density differences between two soft tissues would become hard to see. This is why windowing is used: the physician chooses the level L (*window level* or *center*, the middle of the interval) and the width W (*window width*), and only the values in the interval [L - W/2; L + W/2] are stretched over the entire gray scale [24]. Values below the interval appear black, and those above it appear white. The DICOM standard stores the recommended window in the attributes Window Center (0028,1050) and Window Width (0028,1051), with a function that is linear by default [245].

The simplified form of the transformation is: $G = 255 \cdot (HU - (L - W/2)) / W$, clipped to the interval 0-255.

Numerical example. The mediastinal window, with L = 40 and W = 400, covers the interval from -160 to +240 HU. A lymph node of 20 HU receives $255 \cdot (20 + 160) / 400 \approx 115$, a medium gray. The lung parenchyma, at about -850 HU, falls below -160 and appears black. The lung window, with L = -600 and W = 1500, covers -1350 to +150 HU. The same parenchyma at -850 HU now receives $255 \cdot 500 / 1500 = 85$, a dark gray on which the vessels and a possible nodule

are visible. Soft tissues at 40 HU become almost white (237) and can no longer be compared with one another. Figure 7.1 shows the two transfer functions and their effect on a synthetic phantom.

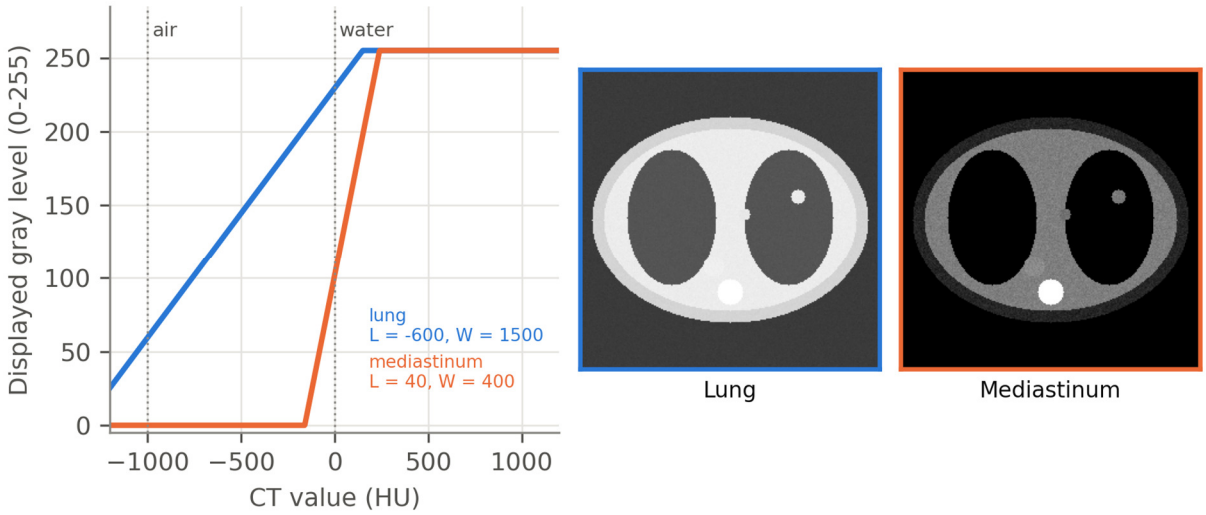


Figure 7.1. Conversion of CT values (HU) into gray levels for the lung window (L = -600, W = 1500) and the mediastinal window (L = 40, W = 400). On the right, the same synthetic slice displayed in the two windows. Synthetic, illustrative data; the HU values of the phantom were chosen for teaching purposes.

The window does not modify the data. It changes only the display, and a lesion that is "invisible" in one window may be obvious in another.

7.2 Imaging Modalities from an Informatics Perspective

From the informatician's point of view, modalities differ in what exactly they measure, in the size and number of their images, and in the type of associated data. Radiography projects all structures traversed by the beam onto a single plane; it has very high spatial resolution and poorer contrast resolution than CT or MRI. Almost all departments have moved from film to digital radiography, either with phosphor plates read afterward (CR) or with detectors that produce the image directly (DR) [24]. Computed tomography (CT) calculates the density of each voxel through reconstruction algorithms. The CT image is therefore computed, not photographed. Ultrasound produces digital images from the outset, often as dynamic sequences (*cine loops*), to which Doppler flow information is added. Magnetic resonance imaging (MRI) generates several sequences (T1, T2, and others) for the same region, while nuclear medicine and PET show the distribution of a radiotracer, with functional information and low spatial resolution, which is why they are frequently fused with CT [24].

Digital breast tomosynthesis (DBT) acquires several low-dose projections of the compressed breast, from which slices are reconstructed that can be scrolled through layer by layer at the workstation [246]. Digital pathology scans glass slides in their entirety (*whole slide imaging*, WSI). Such an image is in color, organized as a pyramid of "tiles" at several magnifications, and takes between 500 MB and 5 GB per slide, compared with 5-200 MB for a typical radiology study [247]. Before diagnostic use, each laboratory validates the WSI system on at least 60 cases, comparing the on-screen diagnosis with the glass-slide diagnosis [248].

Table 7.1. Characteristics of the data produced by the main modalities. Approximate values. Source: [24,68,246,247].

Modality	What it produces	Images / examination	Size / examination
Digital radiography (CR/DR)	2D projection, high spatial resolution	about 3	about 29 MB
CT	volume of slices in HU	about 500	about 250 MB
MRI	several sequences (T1, T2, etc.)	about 500	about 63 MB
Ultrasound	images and cine loops, Doppler	about 50	about 25 MB
2D mammography	4 views, very large matrix	4	about 153 MB
Tomosynthesis (DBT)	slices reconstructed from multiple projections	multiple slices	variable
Nuclear medicine, PET	radiotracer distribution, often fused with CT	about 25 (scintigraphy)	about 1 MB (scintigraphy)
Digital pathology (WSI)	gigapixel color image, pyramid of tiles	1 per slide	500 MB - 5 GB per slide

7.3 Image Compression

The volumes in Table 7.1 explain the interest in compression. There are two families of methods. Lossless compression encodes the redundancy of the image so that the original can be reconstructed exactly, bit for bit. Lossy (irreversible) compression produces an image that is only visually similar to the original, with a degree of fidelity chosen by the user. Lossless methods reach compression ratios of about 2.5:1, while lossy methods can reach, for certain types of images, 40:1 without perceptible or diagnostic loss [68]. The 315 MB CT series from the example above would take about 126 MB with lossless compression.

The problem with lossy compression is that no universal setting guarantees both a good ratio and the absence of visible artifacts. The methods are tuned to the diagnostic task at hand, and compression is more widely accepted in cardiology and pathology than in radiology. As the price of storage falls faster than the number of examinations grows, the pressure toward irreversible compression decreases [68]. The DICOM standard allows both variants through transfer syntaxes (JPEG, JPEG-LS, JPEG 2000, and others), and the file keeps the information on the type of compression applied [245].

Recent evidence is mixed. In a study of 9111 mammograms, deep learning models maintained their performance under moderate compression, but the area under the ROC curve fell from 0.87 to 0.79 at extreme ratios of 10,000-11,000:1 [249]. Twelve radiologists and technologists rated minimally compressed JPEG mammograms as equivalent to the DICOM original, but the conclusion concerns use in training, not in diagnosis [250]. In digital pathology, lossy JPEG compression is already routine and can introduce differences in color and texture, and neural network-based compressors do not transfer well between different compression schemes [251]. The practical rule for the student: the diagnostic archive keeps the original image or a losslessly compressed one, and irreversible compression is used only where a validated protocol allows it.

7.4 The DICOM Standard

7.4.1 Origin and Status

The first CT and MRI scanners saved data on magnetic tapes, in formats proprietary to each manufacturer, which researchers had to decode on their own [68]. The American College of Radiology (ACR) and the National Electrical Manufacturers Association (NEMA) published the ACR-NEMA standard in 1985, and in 1993 it was replaced by DICOM (*Digital Imaging and Communications in Medicine*), which introduced communication over TCP/IP networks, exchange on physical media (CD, DVD, USB), and a formal description of services [252]. The parts of the standard are numbered from PS3.1 (introduction) to PS3.22 (real-time communication), two of them (PS3.9 and PS3.13) having been retired. Among the others, PS3.3 describes the information objects, PS3.4 the services, PS3.6 the data dictionary, and PS3.18 the web services [252]. Approved supplements are consolidated into new editions several times a year. On September 27, 2026, the current edition published on the official website was 2026d [245].

DICOM has also been adopted as an international ISO standard. The draft ISO/DIS 12052:2025, put to the vote in March-June 2025, announced the third edition, which replaces ISO 12052:2017 [252]. The published edition is ISO 12052:2026 [253]. Any product that claims DICOM conformance must have a conformance statement that specifies exactly which services and objects it supports [252]. When purchasing equipment, the hospital compares these statements.

7.4.2 The Information Model

Most DICOM objects follow a common model with the entities patient, study, series, equipment, frame of reference, and the instance itself (Figure 7.2). A study corresponds approximately to a requested procedure, and a series to a performed acquisition step [252]. Each instance, for example an image, receives a globally unique identifier (UID) that accompanies it in all exchanges. If the meaning of the content changes, a new instance with a new UID is created [252].

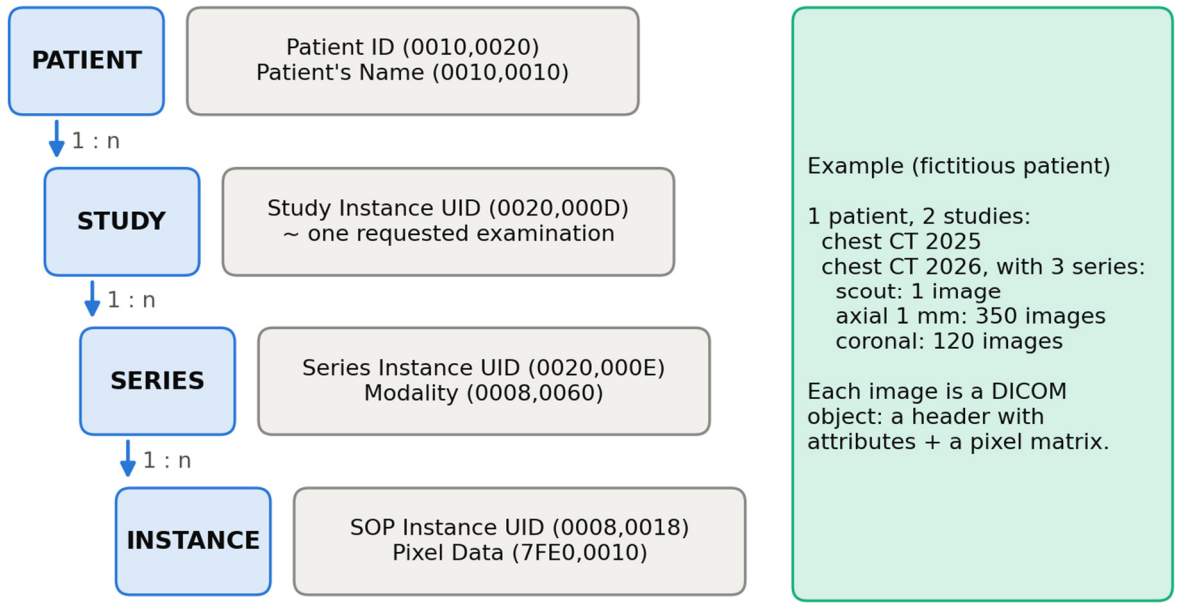


Figure 7.2. Hierarchy of the DICOM information model: a patient has one or more studies, a study has one or more series, and a series contains instances (images). On the right, a fictitious example. The attributes and tags are taken from PS3.6 [245].

7.4.3 Attributes, Tags, IODs, and SOP Classes

A DICOM file has a header made up of attributes and, usually, a pixel matrix. Each attribute has a tag consisting of two hexadecimal numbers (the group and the element), a name, and a value representation (VR). A few examples from the data dictionary: (0010,0010) Patient's Name, of type PN (person name), (0010,0020) Patient ID, (0008,0060) Modality, (0020,000D) Study Instance UID, (0028,0030) Pixel Spacing, which gives the pixel size in millimeters, and (7FE0,0010) Pixel Data, which contains the image itself [245]. Pixel Spacing makes measurements in millimeters possible, something you will verify in Practical Session 7.

An Information Object Definition (IOD) is an abstract description of a type of real-world object, for example "CT image" or "structured report," together with its attributes. The combination of an IOD and a service forms a SOP class (*Service-Object Pair*), for example "storage of a CT image" [245,252]. The SOP class of an instance is recorded in the SOP Class UID attribute (0008,0016) [245].

7.4.4 Network Services and DICOMweb

Classic DICOM services run over TCP/IP between two applications that establish an association. Table 7.2 summarizes the services a student encounters most often.

Table 7.2. Common DICOM services and their DICOMweb equivalents. Source: [245,254,255].

Service	What it does	Example of use	Web equivalent
C-ECHO	checks the connection between two nodes	test when a scanner is installed	-
C-STORE	sends and stores an instance	CT sends the series to the PACS	STOW-RS
C-FIND	searches by criteria	all studies of patient X	QIDO-RS
C-MOVE / C-GET	retrieves the instances found	workstation fetches the prior study	WADO-RS
Modality Worklist	worklist for the scanner	patient data taken from the RIS	UPS-RS
MPPS	reports procedure status	"examination completed" sent to the RIS	UPS-RS

DICOMweb brings together the RESTful services defined in PS3.18: QIDO-RS for searching, WADO-RS for retrieval, STOW-RS for storage, and UPS-RS for worklists, with attributes represented in JSON or XML [254]. The advantage is that a web viewer or a mobile application can access the images with the usual tools of web developers. A workflow built solely on FHIR and DICOMweb interfaces has been described for ordering, scheduling, acquisition, viewing, and reporting [256]. In digital pathology, WADO-RS allows individual tiles of a multi-gigabyte image to be requested without downloading the whole file [247]. At the 2025 DICOM WSI Connectathon, 32 implementations (scanners, archives, viewers, annotation sources) exchanged images and annotations through the standard DICOM and DICOMweb mechanisms [257].

7.4.5 De-Identification

The DICOM header contains personal data, and sometimes the pixels contain it as well: the patient's name "burned" into the image on an ultrasound scan or a screenshot, or facial features that can be reconstructed from a head CT. Part PS3.15 defines the Basic Application Level Confidentiality Profile, which removes or replaces identifying attributes, with additional options

such as cleaning the pixel data (*Clean Pixel Data*), cleaning recognizable visual features, and retaining longitudinal temporal relationships [245]. The last option matters in research, where the interval between two examinations must be preserved even if the actual dates are not.

A workshop organized in 2023 by the US National Cancer Institute brought together the recommendations of the MIDI working group (*Medical Image De-Identification Initiative*), the available tools, and the international approaches to de-identifying images intended for research and sharing, including for AI applications [258]. De-identification is part of a broader set of security measures for PACS: encrypted transport, selective encryption of the header, digital signatures, and avoiding exposure of PACS servers directly to the internet [259]. The legal aspects (GDPR, pseudonymization) are covered in the chapter on data protection.

7.5 PACS, RIS, VNA, and the Radiology Workflow

7.5.1 The Components

The Picture Archiving and Communication System (PACS) receives, stores, distributes, and displays images. The Radiology Information System (RIS) manages the textual side: scheduling, the status of each examination, and the report. Where a hospital information system or an electronic health record (EHR) exists, ordering and billing usually go through it, and the RIS takes over these functions mainly in outpatient imaging centers [68]. The report is dictated, and speech recognition has shortened the time to the final report from a week, which was common before, to less than an hour [68].

The first PACS tightly coupled the archive to the display workstations of the same manufacturer. When the system was replaced, data migration required the cooperation of the old vendor. The vendor neutral archive (VNA) separates storage from viewing, uses DICOM for access, and can also host images from outside radiology. In large hospitals, the VNA has become the rule [68]. Open-source PACS also exist. A comparative evaluation identified Orthanc, DCM4CHE, DCMTK, and Dicoogle as the most mature [260].

7.5.2 The Workflow and the IHE Scheduled Workflow Profile

Figure 7.3 follows a CT examination from order to result. The referring physician places the order in the EHR. The RIS schedules the examination and publishes it in the scanner's worklist (*Modality Worklist*), so the technologist no longer types the patient's name and the demographic data reach the DICOM header correctly. The scanner reports the start and completion of the procedure (*Modality Performed Procedure Step*, MPPS), sends the images to the PACS, and requests confirmation of their receipt (*storage commitment*). The radiologist opens the study at the reporting workstation, compares it with prior examinations, and signs the report, which reaches the RIS and the EHR, usually with a link to the images.

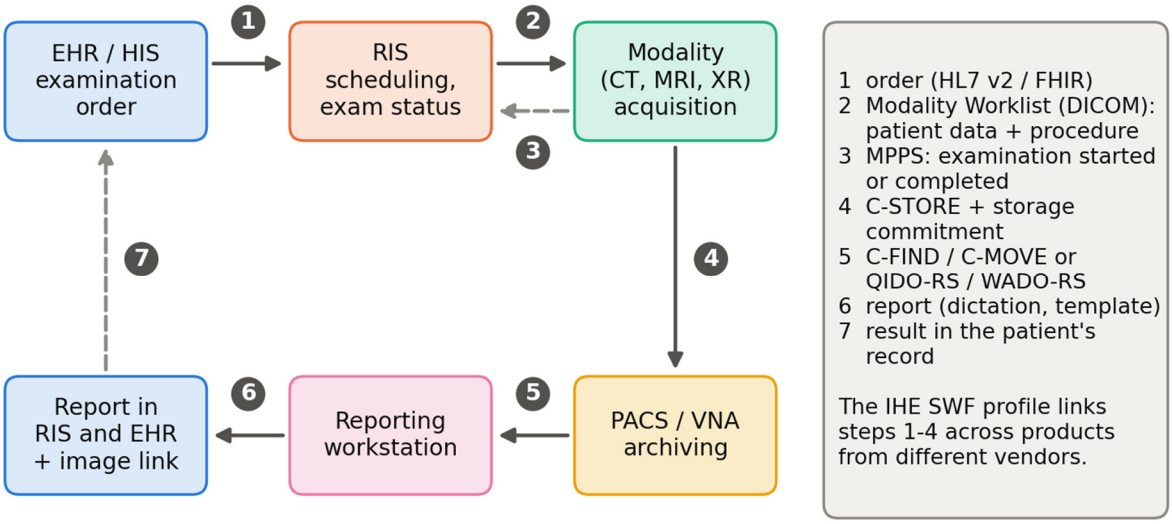


Figure 7.3. The path of an examination through the EHR-RIS-modality-PACS network, with the HL7/FHIR messages and DICOM services for each step. Adapted from [68,255].

Standards alone do not guarantee that different products work together. The IHE initiative (*Integrating the Healthcare Enterprise*) defines integration profiles that specify which standards and which options are used for a given scenario. The Scheduled Workflow (SWF) profile integrates the ordering, scheduling, acquisition, storage, and viewing of radiology examinations. It uses the worklist, MPPS, and storage commitment, and the last of these transfers responsibility for the images from the scanner to the PACS and prevents their loss [255]. Profiles may not contradict the DICOM standard; they may only add requirements [252]. Revision 23.0 of the IHE Radiology Technical Framework, published on August 8, 2025, includes the SWF.b variant and profiles under trial implementation for the results of AI algorithms [261].

7.5.3 Teleradiology and Cloud PACS

Teleradiology means interpreting images at a site other than the one where they were acquired. It began with the overnight transmission of examinations to the on-call radiologist at home and has become an organized activity, with teams reporting for several hospitals. The main obstacles are no longer technical; they concern the licensing and credentialing of the physician in the patient's jurisdiction [68]. The organizational and legal aspects of telemedicine are discussed in Chapter 9.

Economies of scale push many small imaging providers toward cloud storage, and analysis algorithms with high computing requirements have pushed in the same direction [68]. A cloud-native, vendor-neutral PACS architecture, with an on-premises appliance that keeps the classic DICOM interface, stored more than 450 TB of images in production and served more than 4000 DICOM nodes [262]. Moving to the cloud does not relieve the hospital of its responsibility for confidentiality, contracts, and continuity of access.

7.6 Image Processing and Analysis

7.6.1 Global Operations: Histogram, Contrast, Filters

A typical processing pipeline includes global operations, image enhancement, visualization, quantification, segmentation, registration (the alignment of two images), and, finally, classification [24]. The histogram counts how many pixels have each intensity value and is the first thing to look at in a new image. Histogram equalization redistributes the most frequent values to increase global contrast. Convolution filtering applies a small matrix of coefficients (a kernel) to each pixel and its neighbors. A Gaussian filter smooths the image and reduces noise, while a median filter removes isolated points and preserves edges better. Unsharp masking increases local contrast and enhances fine detail [24]. The CT window from Section 7.1 is also a global operation.

7.6.2 Segmentation and Measurements

Segmentation delineates a structure: an organ, a tumor, a nodule. The simplest method is thresholding, by which pixels above a value become "object" and the others "background." Next come region-based, edge-based, and clustering methods and, in recent years, deep neural networks [24]. nnU-Net is a method that configures its own preprocessing, architecture, and training for each new data set, and it outperformed most specialized solutions on 23 public data sets from international competitions [263]. Built on nnU-Net, TotalSegmentator segments 104 anatomical structures on CT, with a Dice coefficient of 0.943 on the test set [264]. The Dice coefficient measures the overlap between the automatic segmentation and the reference segmentation, from 0 (no overlap) to 1 (identity).

Segmentation is followed by measurements: lengths, areas, volumes, mean density in HU. All of them depend on the spatial calibration in the header (Pixel Spacing, slice thickness). A file exported as JPEG loses this information and, with it, the possibility of measuring in millimeters.

7.6.3 Radiomics in Brief

Radiomics extracts a large number of quantitative features (shape, intensity statistics, texture) from images and analyzes them with statistical or machine learning methods, for example for disease classification or risk stratification [265]. The typical workflow includes segmentation of the region of interest, image processing (harmonizing pixel size and gray levels), feature extraction and selection, after which the model is built [265]. The weak point is reproducibility: features vary with the scanner, the acquisition protocol, the reconstruction, and the software used [265]. The Image Biomarker Standardisation Initiative (IBSI), with 25 teams, standardized 169 of 174 radiomic features, which allows different software packages to be calibrated against one another [266].

7.7 Artificial Intelligence in Imaging, 2026

7.7.1 How Many Devices, in Which Specialties

The US Food and Drug Administration (FDA) publishes a list of medical devices that use AI and have received marketing authorization in the United States. The list is not exhaustive: devices are identified mainly by AI-related terms in the authorization documents, and the list is updated periodically [267]. The page consulted on September 27, 2026, had been updated on September 4, 2026. The downloadable file contained 1614 devices, with final decisions between September

1995 and June 29, 2026 [267]. Radiology was the lead panel for 1230 of them (76.2%), followed by the cardiovascular field, with 154 (9.5%), and by neurology, with 73 (4.5%). In 2025 alone, 335 authorizations were recorded (Figure 7.4).

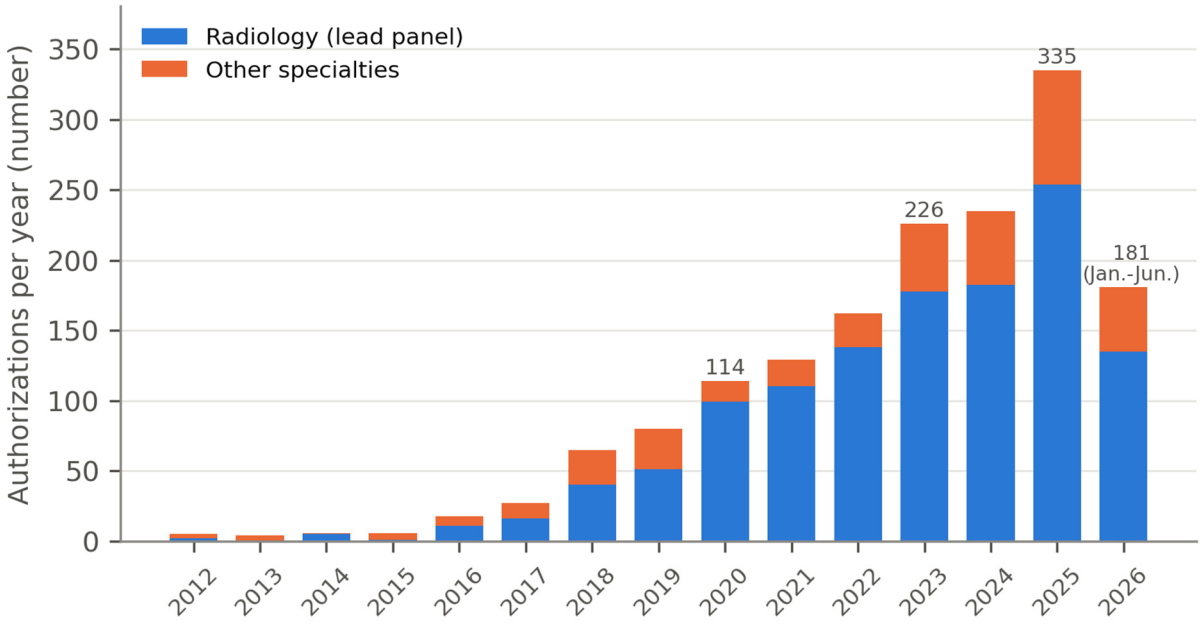


Figure 7.4. AI-enabled medical devices authorized by the FDA, by year of final decision (2012-2026), radiology versus other specialties (lead review panel). For 2026, decisions from January-June are included. Our own calculation based on the official list downloaded on September 27, 2026 (page updated on September 4, 2026) [267].

A longitudinal analysis of the 1430 devices authorized through the end of 2025 showed that the share of radiology peaked at 85.5% in 2020-2022 and fell to 77.5% in 2023-2025, the first statistically significant decline in the registry [268]. The current list yields 1433 devices for the same period, the small difference coming from subsequent updates. A large number of authorizations, however, does not mean solid clinical evidence. In the same analysis, 96.2% of the devices had been authorized through the 510(k) pathway, which relies on equivalence with a device already on the market [268].

7.7.2 Mammography Screening: Randomized and Real-World Evidence

For screening mammography, evidence from randomized trials already exists. The Swedish MASAI trial randomized women from the national program to AI-supported reading or standard double reading. The AI system assigned each examination a risk score from 1 to 10. Examinations with a score of 1-9 went to a single radiologist, and those with a score of 10 to two radiologists. The readers saw the score and, for scores of 8-10, also the algorithm's marks on the image. The safety analysis of 80,033 women found a detection rate of 6.1 per 1000 with AI versus 5.1 per 1000 without AI, with a 44.3% reduction in the reading workload [269]. In the full cohort of 105,934 women, detection was 6.4 versus 5.0 per 1000 (ratio 1.29), with no significant increase in recalls or false-positive results, and the additional cancers were mostly small invasive cancers without lymph node involvement [270]. The primary outcome, the rate of interval cancers (those arising between two screening rounds), was published later: 1.55 versus 1.76 per 1000, that is, noninferiority, with a sensitivity of 80.5% versus 73.8% and the same specificity, 98.5%, in both arms [271].

The German PRAIM study analyzed real-world implementation at 12 centers, with 463,094 women and 119 radiologists who chose for themselves whether to use the AI system. Detection

was 6.7 per 1000 with AI versus 5.7 per 1000 without, and the recall rate did not increase (37.4 versus 38.3 per 1000) [272]. PRAIM is an observational study, so the groups are not as strictly comparable as in an RCT.

7.7.3 Pulmonary Nodules on CT

For CT lung screening programs, a systematic review of CE-marked AI software included 11 studies with 19,770 patients, all at high risk of bias. Assisted reading took less time and generally increased sensitivity (by 5-20 percentage points for nodules requiring management), at the cost of a 3-7 point decrease in specificity. At a cancer prevalence of 0.5%, the authors estimated 150-750 additional cancers detected per million people screened, but also 59,700-79,600 people without cancer sent unnecessarily to CT surveillance [273]. The extended technology assessment, with 27 studies, confirmed the high risk of bias in all included studies [274]. The example shows why "higher sensitivity" does not automatically mean a benefit for the population.

7.7.4 Limitations: Generalization, Spectrum Bias, Overreliance

An algorithm that performs well on data from the hospital that developed it may be mediocre in another hospital. A systematic review of 86 externally validated deep learning algorithms found a drop in performance for 81% of them, substantial (at least 0.10 on the 0-1 scale) for 24% [275]. Another review, of 535 radiology studies, found that 98% were retrospective, and in the 77 studies that allowed the comparison, performance dropped on average by 6% at external validation [276]. Digital pathology shows the same trend. A meta-analysis of 48 studies on scanned slides reported a mean sensitivity of 96.3% and a specificity of 93.3%, but 99% of the studies included in the review had at least one domain at high or unclear risk of bias or with applicability concerns [277], and differences between institutions (staining, scanners) remain one of the main problems in the field [278].

Spectrum bias occurs when the patients in the test set do not resemble those seen in practice, for example when the positive cases are large, obvious tumors and the negative ones are healthy volunteers. The reported performance is then overestimated. Biases can enter at every stage, from data collection to the way metrics are chosen and results are published [279]. In the statistical sense, bias means a systematic error: the model's results do not represent real-world conditions [280].

The third risk concerns the human reader. Automation bias is the tendency to accept the suggestion of an automated system even when it contradicts other evidence [280]. In an experiment with 27 radiologists who read mammograms with a purported AI system, the proportion of correct BI-RADS assessments fell from about 80% when the suggestion was correct to 20-46% when it was wrong, and inexperienced readers were the most vulnerable [281]. In a study on chest radiographs, incorrect AI results increased the radiologists' false-negative rate from 2.7% to 33.0%, an effect that was attenuated when the algorithm framed the suspicious area [282]. A large study, with 140 radiologists and 15 tasks on chest radiographs, showed that experience and familiarity with AI do not predict who benefits from assistance, whereas algorithm errors impair readers' performance [283].

7.7.5 The Regulatory Framework in the European Union

In the European Union, software intended for a medical purpose, including an image detection algorithm, is a medical device and falls under Regulation (EU) 2017/745 on medical devices (MDR) [284]. Regulation (EU) 2024/1689 on artificial intelligence (AI Act) classifies as high-risk

systems, among others, AI systems that are medical devices subject to conformity assessment by a notified body [22], and the timetable for applying the obligations for these systems was amended by Regulation (EU) 2026/1744 of July 8, 2026 (*Digital Omnibus on AI*) [197].

Common Pitfalls

- Believing that the window modifies the image. It changes only the display, and the HU values in the file remain the same.
- Exporting a DICOM image as JPEG in order to measure it. The calibration (Pixel Spacing), the HU values, and, in most cases, some of the quality are lost.
- Confusing the study with the series. A CT study may have several series (scout, axial, coronal), and a series has hundreds of images.
- Considering a file de-identified just because the name has been deleted. The ID, the date of birth, text "burned" into the pixels, and the face reconstructible from CT all remain.
- Reading "AI increases sensitivity" without asking what happens to specificity and to the number of false-positive results.
- Taking the number of authorized devices as proof of clinical benefit.

Key Points

- A 512 x 512 pixel CT image at 16 bits takes about 0.5 MB, and a series of 600 slices about 315 MB uncompressed.
- The window (L, W) displays the interval [L - W/2; L + W/2] over the entire gray scale.
- DICOM organizes data on four levels: patient, study, series, instance, each with its own UID.
- C-STORE stores, C-FIND searches, C-MOVE retrieves; in DICOMweb, STOW-RS, QIDO-RS, and WADO-RS do the same over HTTP.
- The RIS manages scheduling and the report, the PACS manages the images, and the IHE SWF profile links them through the worklist, MPPS, and storage commitment.
- In the MASAI RCT, AI-supported reading detected more cancers, with 44% less reading workload and a noninferior interval cancer rate.
- External validation and automation bias are the first two questions to ask of any algorithm.

Self-Assessment Questions

1. (SA) A CT series has 400 images with a 512 x 512 matrix and 16 bits per pixel. The approximate uncompressed size (without headers) is:

a) 21 MB
b) 52 MB
c) 105 MB
d) 210 MB
e) 1.6 GB

2. (SA) For a window with level L = 40 HU and width W = 400 HU, a voxel with a value of -300 HU will be displayed as:

a) white
b) light gray
c) medium gray
d) dark gray, but distinct from black
e) black

3. (SA) In the DICOM information model, the entity that corresponds approximately to a requested imaging procedure is:

a) the patient
b) the study
c) the series
d) the instance
e) the equipment

4. (MA) Which of the following statements about DICOM and DICOMweb services are correct?

a) C-STORE sends an instance to another node, for example from the CT scanner to the PACS
b) QIDO-RS is the web service for searching DICOM objects
c) Modality Worklist transmits images from the PACS to the reporting workstation
d) WADO-RS allows DICOM objects to be retrieved over HTTP
e) DICOMweb can represent attributes in JSON

5. (MA) In the MASAI randomized trial of AI-supported mammography screening:

a) the cancer detection rate was higher in the AI arm
b) the radiologists' reading workload decreased by approximately 44%
c) the interval cancer rate was noninferior to that of standard double reading
d) specificity was much lower in the AI arm
e) the radiologists never saw the algorithm's marks

CHAPTER 8

Chapter 8. The Electronic Health Record and Interoperability

Objectives

By the end of this chapter, the student will be able to:

- distinguish the electronic medical record, the electronic health record, and the personal health record by who manages them and what data they contain;
- list the components of an electronic record and argue, using an example, the advantages and limitations of structured documentation compared with free text;
- assign a data exchange problem to one of the four levels of interoperability and indicate the appropriate standard for each level;
- read a PID segment from an HL7 v2 message and a FHIR Patient resource in JSON format, identifying the patient's name, date of birth, and identifier;
- describe the European Health Data Space timeline for the patient summary and e-prescription;
- recognize patterns of unsafe use of the electronic record (copy-paste, alert overload) and propose a preventive measure.

8.1 The Electronic Record: Concepts and Effects

8.1.1 EMR, EHR, PHR

The patient record serves several purposes: to preserve observations, to inform colleagues, to train students, to generate new knowledge, to allow activity to be monitored, and to justify interventions [285]. The electronic health record (EHR) is an electronic repository of information about a person's health status and care, equipped with functions that support clinical care and the other legitimate uses of the data [285]. The difference from the paper record lies not only in the medium. An EHR can check data at the very moment they are entered (for example, a dose range), can display trends, and can flag dangerous values.

Three related terms circulate in the literature. The electronic medical record (EMR) usually designates the documentation of a single institution, such as a practice or a hospital. The EHR gathers data from several care settings over a person's lifetime. The authors of the field's reference textbook note, however, that the terms are often used interchangeably, and many have moved from "medical" to "health" because the systems are used throughout the continuum of care by numerous categories of professionals [285]. The personal health record (PHR) is controlled by the patient: patients see their results, receive messages from the care team, and can add data measured at home, such as blood pressure or blood glucose [285]. Patient portals, linked to an institution's EHR, are related to the PHR and are often analyzed together with it as patient-centered digital records [286] (Table 8.1).

Table 8.1. The electronic medical record, the electronic health record, and the personal health record. Source: [285,286].

Characteristic	EMR (electronic medical record)	EHR (electronic health record)	PHR (personal health record)
Who manages it	an institution (practice, hospital)	the provider network, the region, or the state	the patient
Data source	the institution's clinicians	several institutions, through data exchange	EHR, patient, home devices
Main purpose	ongoing care and billing	continuity of care, secondary uses	informing and engaging the patient
Example	the record in the family practice software	the regional or national record	patient portal, mobile app

8.1.2 Components and Documentation

A mature EHR is not a mere archive. The Shortliffe textbook describes six functional components: capture, aggregation and review of patient data; computerized provider order entry (CPOE); clinical decision support; access to knowledge resources; communication within the team and with the patient; and coding and billing [285]. The data themselves fall into easily recognizable blocks:

- demographic and administrative data (identifier, name, date of birth, sex, address, insurance);
- the problem list (active and resolved diagnoses, preferably coded);
- current medication and prescription history;
- allergies and intolerances, with the type of reaction;
- laboratory results and vital signs;
- clinical notes (history, examination, progress, discharge summary);
- images and imaging reports, linked to the PACS described in Chapter 7.

The problem list is often incomplete. A randomized trial conducted in 4 health systems with 3 different EHRs used an alert that suggested adding missing problems. The acceptance rate was 22.1%, and 4.6 times more problems were added in the intervention arm than in the control group, with no change in quality-of-care indicators [287].

Documentation is done in two ways, with opposite advantages. Free (narrative) text is easy to write and preserves clinical reasoning and nuance, but it is hard to search, and abbreviations and variant phrasings make it difficult to process automatically [285]. Structured documentation uses forms with fields and controlled vocabularies. The computer can check the consistency of the data immediately and can reuse them for alerts, statistics, or research. The price is time: in a study cited by the same textbook, entering a diagnosis coded in SNOMED CT took up to 20 seconds [285]. A common practical solution is a mixed (semi-structured) approach, with structured fields for the data that will be reused (allergies, medication, vital signs) and free text for reasoning [285].

8.1.3 What Effects Have Been Documented

The benefits are not automatic, but they exist. A 2026 meta-analysis of 74 randomized trials in primary care found that digital safety interventions, from decision support integrated into the EHR to audit and feedback, improved process indicators for medication safety (OR 1.46; 95% CI 1.30-1.64) and modestly reduced adverse events (OR 1.17; 95% CI 1.06-1.30, estimated from only 12 studies). Multicomponent interventions worked better than single-component ones [288]. In patients with chronic diseases, favorable effects of portals and PHRs were reported mainly for the

use of recommended services, disease knowledge, patient engagement, and treatment adherence, but only 16 of 81 studies were of high methodological quality [286].

The most visible cost is time spent in front of the screen. A systematic review of 18 studies (half of them observational time-motion studies, almost all the others based on system logs) estimated that physicians spend approximately 37% of their working day in the EHR, and hospital nurses approximately 22% [289]. Audit logs record every user action and make such measurements possible on a large scale.

The figures obtained from logs converge. In data from approximately 155,000 US physicians using the same commercial product, the mean time in the EHR was 16 minutes and 14 seconds per visit, of which 33% was spent on chart review, 24% on documentation, and 17% on orders [290]. A comparison across 371 health systems using the same vendor showed that US clinicians (physicians and other professionals with their own schedules) worked actively in the EHR for 90.2 minutes per day, compared with 59.1 minutes outside the United States, and received 33.8 messages per day, compared with 12.8 [291]. The product was the same, so the difference also reflects the context of use. The trend is upward: among 141 primary care physicians at an academic center, EHR time per 8 hours of scheduled patient time increased by 28.4 minutes (7.8%) between 2019 and 2023, and messages in which patients requested medical advice increased by 55.5% [292]. Time-motion studies add a dimension that logs do not capture. Among 47 directly observed clinicians, task switching occurred on average 1.4 times per minute, and almost half of the task switches involved viewing or entering data [293].

The link with burnout is documented, although almost exclusively through cross-sectional studies. A meta-analysis of 37 studies with 66,556 participants estimated a pooled burnout prevalence of 40.4% among professionals who use EHRs. For those who spent more time in the EHR after hours, the odds ratio for burnout was 2.43 [294]. In a survey of 2067 family physicians, good system usability (data entry, finding information, useful alerts) was associated with satisfaction with the EHR, and satisfaction was associated with a lower frequency of burnout [295]. One solution recently evaluated on a large scale is the artificial intelligence scribe, which transcribes and summarizes the visit. In a cohort study of 8581 clinicians at 5 US academic centers, adopting such scribes was associated with 16 fewer minutes of documentation per 8 hours of scheduled patient time, but after-hours EHR time did not change significantly [296].

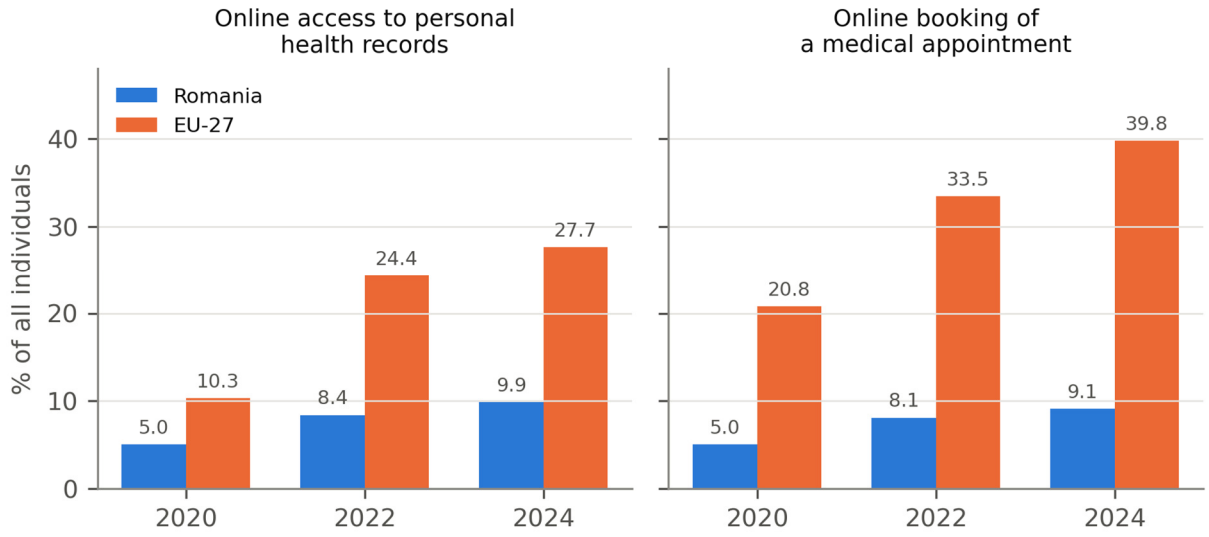


Figure 8.1. Online access to personal health records and online booking of medical appointments, Romania versus the EU-27, 2020-2024 (% of all individuals, with the 2020 EU-27 value estimated by Eurostat). Eurostat data isoc_ci_ac_i, extracted on September 27, 2026 [297].

The patient's side of the equation remains small in Romania. According to Eurostat, in 2024 only 9.9% of individuals in Romania had accessed their health records online, compared with 27.7% on average in the EU-27, and online booking of medical appointments was used by 9.1%, compared with 39.8% in the EU (Figure 8.1) [297].

8.2 Interoperability

8.2.1 The Four Levels

Interoperability is the ability of two systems to exchange data and to use them without manual intervention. A scoping review of 24 papers on EHRs found four levels (technical or foundational, syntactic or structural, semantic, and organizational or process) and identified the semantic level, that is, preserving the meaning of the data, as the main problem to be solved [298]. Figure 8.2 attaches a simple question to each level.

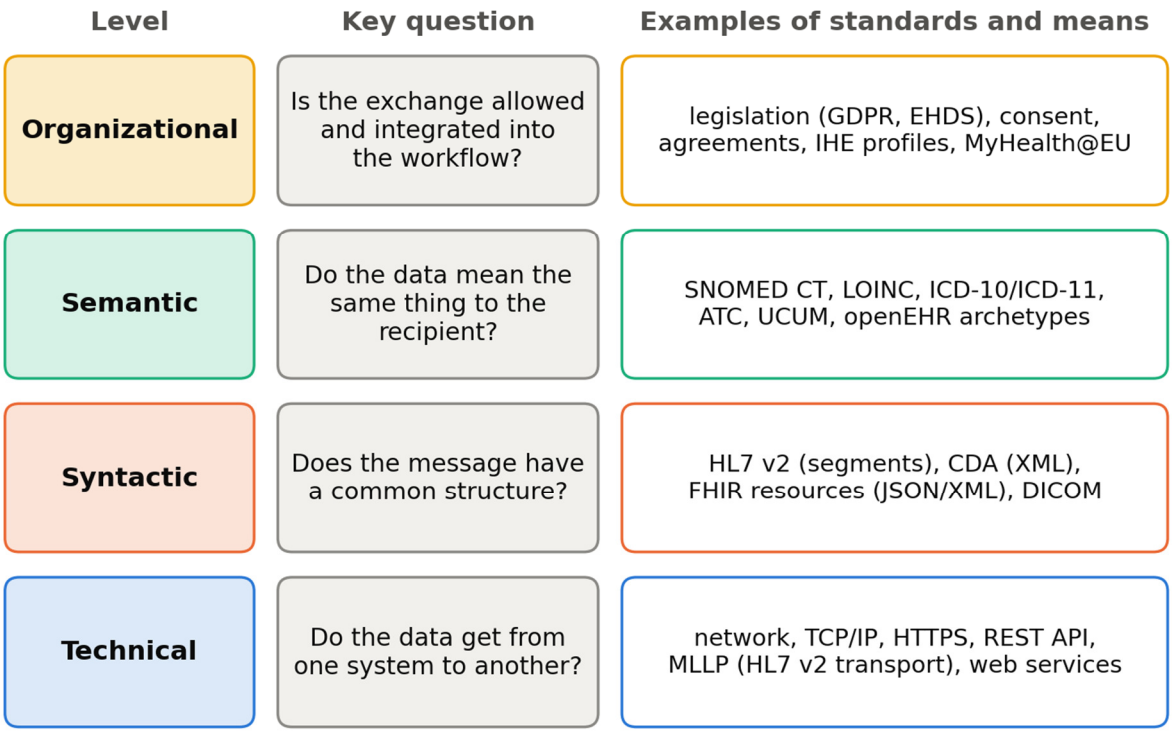


Figure 8.2. The levels of interoperability, the question each one answers, and examples of standards. The higher levels presuppose that the lower ones have been resolved. Adapted from [298].

The technical level ensures transport: networks, protocols, application programming interfaces. The syntactic level fixes the format, that is, the order of the fields, the separators, and the document structure. The semantic level requires the sender and the recipient to understand the same thing by the same code, which presupposes shared terminologies such as SNOMED CT for clinical concepts, LOINC for laboratory tests, and ICD for diagnoses. The organizational level covers legislation, consent, agreements between institutions, and the integration of exchange into the workflow.

An example shows why the syntactic level is not enough. One study followed 9 laboratory tests transmitted between two institutions via HL7 v2 messages. Although the exchange used standard messaging, safety-relevant information was lost: the specimen source was systematically missing, units and reference ranges were inconsistent, and LOINC codes were used differently for the same tests [299].

8.2.2 HL7 Version 2 and CDA

HL7 v2 is the most widely used messaging standard in hospitals. HL7 International describes it as probably the most widely implemented standard in healthcare, used by 95% of US healthcare organizations and implemented in more than 35 countries [300]. A v2 message starts from a trigger event (for example, the admission of a patient) and consists of segments. Each segment contains fields, and the fields contain components of a given data type. The message identifies the sender, the recipient, the message type, and the patient [15]. Here is a fragment of an admission message with fictitious data:

```
MSH|^~\&|HIS|SPITAL_DEMO|LIS|LAB_DEMO|20260927103000||ADT^A01|MSG00001|P|2.5.1
PID|1||123456^^^SPITAL_DEMO^MR||POPESCU^ION^^^^^L||19800115|M|||STR. EXEMPLU
1^^TIMISOARA^^300000^ROU
```

The vertical bar separates the fields, and the caret (^) separates the components. The MSH segment (the header) states that this is an ADT^A01 (admission) message in version 2.5.1. In the PID segment, field 3 contains the identifier 123456 assigned by SPITAL_DEMO, with type MR (medical record number), field 5 the name (family name, given name, type L for legal name), field 7 the date of birth in YYYYMMDD format, field 8 the sex, and field 11 the address. The format is compact, but many fields are optional, and local variants mean that every v2 interface requires extensive testing.

The Clinical Document Architecture (CDA) Release 2 solves a different problem: the exchange of whole documents, such as the discharge summary or the imaging report. A CDA document is written in XML and contains both a human-readable presentation and machine-processable data structures. HL7 states that future development for structured clinical documents is expected to focus mainly on FHIR [301].

8.2.3 openEHR and ISO 13606

Messaging standards describe how data travel. Two other specifications describe how data are kept in the record. openEHR is an open set of specifications, clinical models, and software that separates the technical part (how the data are stored, described by the reference model) from the clinical part (what the data say), described through archetypes [15,302]. Archetypes are reusable models of a clinical concept, such as blood pressure, built by health and health technology professionals. The international archetype catalog is hosted in the Clinical Knowledge Manager tool [302,303]. The specification covers the recording, storage, and querying of data [303]. The ISO 13606 standard is also based on a reference model and archetypes, but it is oriented toward the exchange of record extracts with their full meaning preserved, and it is recommended for regional, national, or international interoperability projects [303]. In 2023, a group of experts concluded that openEHR, ISO 13606, and FHIR should each be applied to the needs for which they were created, and that they can work together [303]. In practice, adoption remains uneven: in a review of 99 studies from 2021-2024, 57% used the OMOP model for research, 39% FHIR, and only 8% openEHR [304].

8.2.4 HL7 FHIR

FHIR (*Fast Healthcare Interoperability Resources*) was first presented by HL7 in 2011 to overcome the complexity of development based on the HL7 Reference Information Model (RIM), and it was designed to remain compatible with systems that use v2 or v3 messages and with CDA implementations [15]. The basic unit is the resource: all exchangeable content is defined as a

resource, with elements, data types, and a human-readable section. Resources are represented in JSON or XML and are linked to one another through references [305]. Five resources cover a good part of everyday clinical activity:

- **Patient**, demographic data and identifiers;
- **Encounter**, the patient's contact with health services (outpatient visit, admission);
- **Observation**, measurements and results (vital signs, laboratory tests, scores);
- **Condition**, problems and diagnoses;
- **MedicationRequest**, the prescription of a medication.

Access takes place through a REST interface, that is, through the ordinary requests of the web. `GET [base]/Patient/123` reads the patient with the identifier 123, `GET [base]/Observation?patient=123` searches for all of that patient's observations and returns them as a Bundle, and `POST [base]/Observation` creates a new observation (Figure 8.3). Below is a minimal, fictitious Patient resource:

```
{
  "resourceType": "Patient",
  "id": "123",
  "identifier": [{ "system": "urn:oid:1.2.3.4.5", "value": "1800115123456" }],
  "name": [{ "family": "Popescu", "given": ["Ion"] }],
  "gender": "male",
  "birthDate": "1980-01-15"
}
```

The same data as in the PID segment appear here with explicit field names. The identifier and its system (`system`) are fictitious.

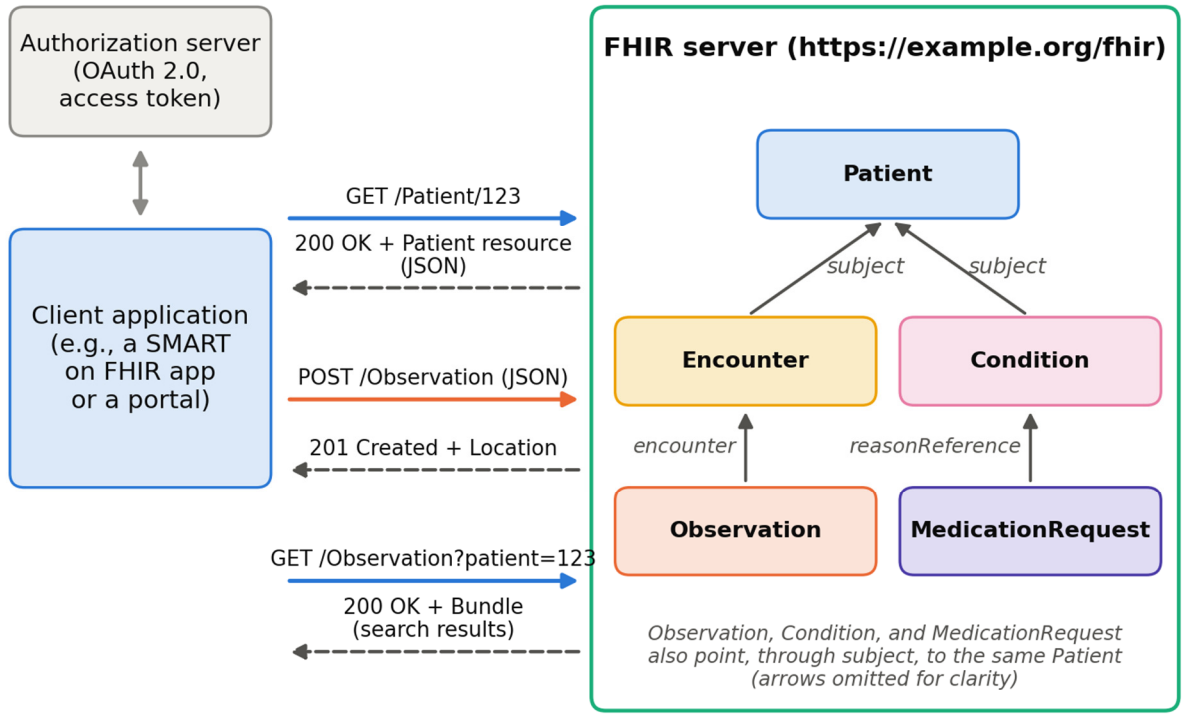


Figure 8.3. REST interaction between a client application and a FHIR server (read, create, search), and the graph of references between the Patient, Encounter, Observation, Condition, and MedicationRequest resources. The names of the linking elements are those of FHIR R4. Adapted from [305,306].

The base resources are general. A profile constrains a resource for a particular context, for example by fixing the mandatory terminologies or the fields that must be filled in, and an implementation guide (IG) brings together the profiles, rules, and examples for a use case, such as the patient summary in Section 8.3. The extension mechanism allows elements to be added without modifying the base resources [305].

FHIR is published in successive versions. Release 4 (4.0.1, October 30, 2019) brought the first normative, stable content, and most current implementations and guides rely on it. R4B (4.3.0, May 28, 2022) was an interim version, and R5 (5.0.0, March 26, 2023) is the latest published version. R6 had not been published as of September 27, 2026: the second full ballot (6.0.0-ballot5) and a snapshot labeled "stable" (6.0.0-snapshot1) appeared on July 17, 2026, both with draft status [307]. For a student, the practical rule is simple: the implementation guides cited in this chapter (IPS, SMART App Launch, CDS Hooks) and the public test server in Practical Session 8 use R4.

The review of 99 studies cited above observed that FHIR implementations frequently require local adaptations, which complicates the very interoperability they aim for [304].

8.2.5 SMART on FHIR and CDS Hooks

SMART App Launch is the HL7 guide that describes how an external application obtains authorization, authenticates, and integrates with a FHIR server. It is based on OAuth 2.0 and defines two scenarios: the EHR launch, in which the application receives the context (for example, the patient currently open), and the standalone launch, in which the user chooses the patient after startup. Permissions are expressed through scopes, for example the right to read and search Encounter resources. The current version is STU 2.2, published on April 30, 2024 [306]. On this basis, a group of researchers integrated two patient-reported outcome applications in 18 practices that used 3 different EHR platforms [308].

CDS Hooks is complementary. The specification describes how an EHR calls, via a REST API and JSON messages, an external decision support service at certain points in the workflow, called *hooks*, for example opening the record (`patient-view`) or signing a prescription (`order-sign`). The service responds with "cards" displayed to the clinician [309]. The current version is STU 2 (2.0.1, March 2025), and a normative version (R3) was balloted in December 2025 [309]. In a randomized trial with 70 clinicians in an emergency department, CDS Hooks cards suggesting a clinical calculator increased the use of the corresponding SMART application from 2.6% to 6.0% of interactions (OR 2.45) [310].

Table 8.2. The main standards for the electronic record and their role. Source: [300-303,305,306,309,311,312].

Standard	Organization	Level	What it is used for
HL7 v2	HL7 International	syntactic	messages between the systems of a hospital (admissions, laboratory)
CDA R2	HL7 International	syntactic	clinical documents in XML (discharge summary, report)
FHIR R4/R5	HL7 International	syntactic, with semantic bindings	resources and REST API for applications and exchange
openEHR	openEHR International	semantic (modeling)	storing the record on the basis of archetypes
ISO 13606	ISO/CEN	semantic (modeling)	communicating record extracts
ISO 27269	ISO/CEN	content	the data set of the International Patient Summary
ISO 17523	ISO/CEN	content	data requirements for an e-prescription
SMART App Launch	HL7 International	technical, security	authorizing applications on FHIR data
CDS Hooks	HL7 International	workflow integration	calling decision support services

8.3 The International Patient Summary and Cross-Border Services

8.3.1 IPS

A physician on call who receives an unconscious tourist needs a few pieces of information, not the entire record: active problems, allergies, medication. The International Patient Summary (IPS) answers this need. The ISO 27269 standard defines the IPS data set as minimal and non-exhaustive, intended for scheduled and unscheduled care across any organizational boundary, including between countries [311]. The HL7 guide that implements it also adopts from the standard the statement that the data set is specialty-agnostic and condition-independent [313]. Its first form was a European standard in 2019, subsequently promoted to an international standard in 2021. The second edition, ISO 27269:2025, remains an abstract definition from which implementable models are derived. It adds the "Alerts" and "Patient Story" sections, and the former "Cross-Border" and "Provenance" sections are merged into a single block, "Cross Organization" [311].

The main technical implementation is the HL7 FHIR IPS guide. The current version, 2.0.1 (STU 2), was published on June 19, 2026, and is based on FHIR R4 [313]. The guide divides the summary into sections. Three are mandatory: the problem list, allergies and intolerances, and the medication summary. Immunizations, results of investigations, procedures, and medical devices are recommended, and the rest (vital signs, social history, pregnancy, advance directives) are optional [313]. The summary can also travel through the patient: four implementations, from the United States, Canada, and the organization of the Hajj pilgrimage, used IPS together with QR codes and SMART Health Links so that patients could present their own history [314].

The limitation of the IPS is its very conciseness. An analysis carried out for two European heart failure projects found that the IPS guide can also serve research, but only after extensions (minor, in the authors' opinion) with specific elements [315].

8.3.2 E-Prescription

The ISO 17523 standard, in its second edition of 2025, sets out the requirements for the content of an electronic prescription, that is, the digital document issued by the prescribing physician and received by the pharmacist who dispenses the medication. It does not regulate the method of transmission, validation, or reimbursement, which are left to each country, and it relies on the ISO standards for the identification of medicinal products (IDMP) [312]. Two national systems that conform to ISO 17523 may therefore remain technically incompatible without a common exchange infrastructure.

8.3.3 MyHealth@EU

The common European infrastructure is called MyHealth@EU. It offers two services: e-prescription and e-dispensation, through which citizens can collect their medication from a pharmacy in another Member State, and the patient summary, which gives the physician in the country of treatment the basic information about the patient in the physician's own language [316]. Each state connects through a national contact point, and the data flow between these points (Figure 8.4) [23].

Actual use has grown slowly. An analysis of the official indicators from 2019 to the second quarter of 2022 counted eleven countries with at least one operational service, nine of which had reported transactions. The program was still used mainly for testing, and cross-border dispensing of medication was carried out mainly for Estonian and Finnish patients [317]. In a survey of

pharmacists in these two countries, 88% of the Estonians and 81.8% of the Finns considered that cross-border e-prescription had improved access to medicines. At the same time, 61% and 42.8%, respectively, had encountered ambiguities or errors in prescriptions [318].

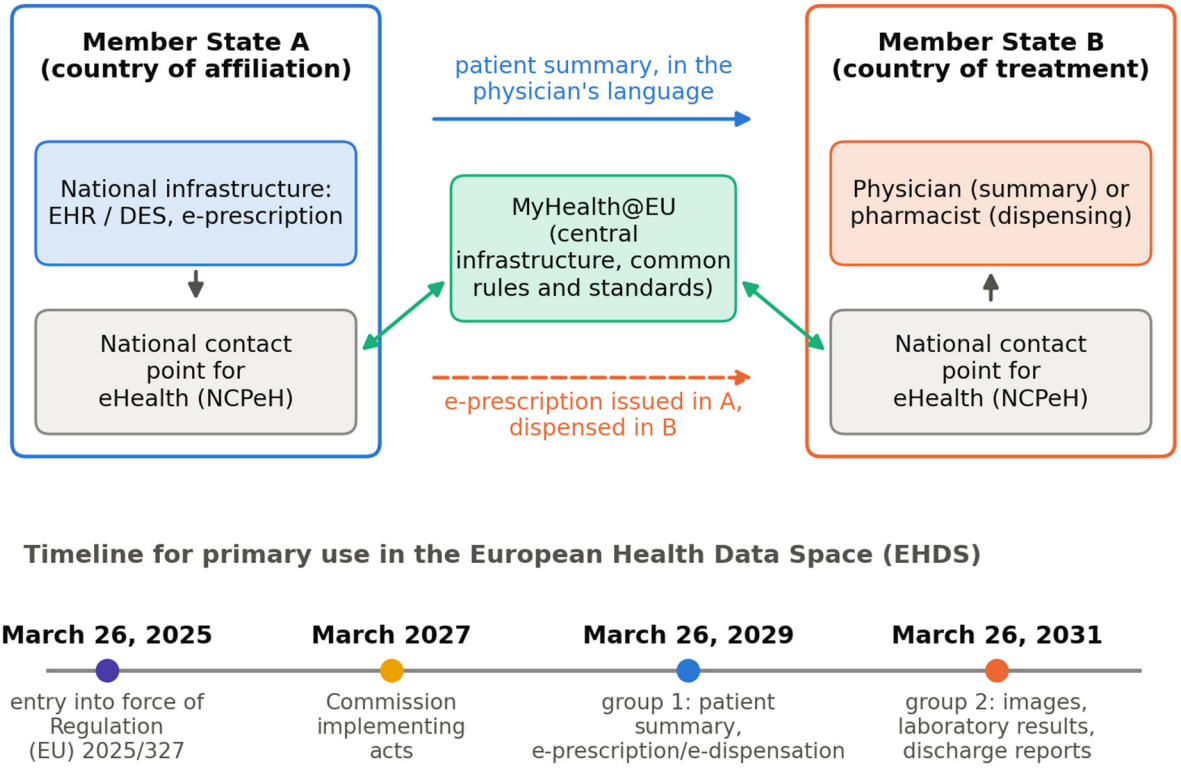


Figure 8.4. Cross-border exchange through MyHealth@EU: the patient summary and the e-prescription travel between the national contact points of the country of affiliation (A) and the country of treatment (B). Below, the timeline for primary use set by Regulation (EU) 2025/327. Adapted from [316,319].

8.4 The European Health Data Space

Regulation (EU) 2025/327 on the European Health Data Space (EHDS) was adopted on February 11, 2025, published in the Official Journal on March 5, 2025, and entered into force on March 26, 2025. It generally applies from March 26, 2027, but the specific obligations take effect in stages [23,319]. The regulation distinguishes two uses of data. Primary use means using the data for the care of the person concerned, while secondary use means reusing them for research, innovation, public policy, or regulatory activities [319,320].

8.4.1 Primary Use

For primary use, the regulation defines six priority categories of data: patient summaries, electronic prescriptions, electronic dispensations, medical imaging studies and the related reports, medical test results (including laboratory results) and the related reports, and discharge reports [23]. The timeline has two steps (Figure 8.4). From March 26, 2029, patients' rights and the obligations relating to the first three categories, as well as to the EHR systems that process them, apply in all Member States. From March 26, 2031, images, laboratory results, and discharge reports are added [23,319].

Patients receive explicit rights: access to their own data in the priority categories and the right to restrict health professionals' access to certain data. Member States may, in addition, provide for

a right to opt out of primary use [23]. Each state designates a national contact point for digital health, connected to MyHealth@EU, so that the cross-border exchange described in Section 8.3 becomes an obligation, not an option [23]. The common technical format, called the European Electronic Health Record Exchange Format (EEHRxF), relies on specifications such as the patient summary described above [315].

8.4.2 Requirements for EHR Systems

The EHDS also regulates software. An EHR system that processes priority categories must include two harmonized software components: the European interoperability component and the European logging component, which records who accessed the data. The system may be placed on the market only if it complies with the requirements of the chapter on EHR systems [23]. A hospital that buys software today must therefore check these requirements.

8.4.3 Secondary Use in Brief

The chapter on secondary use also applies from March 26, 2029, with some categories (for example, genetic and genomic data) only from 2031. Researchers obtain access on the basis of a permit issued by a health data access body, and cross-border exchange goes through the HealthData@EU infrastructure [23,319]. A rapid review of the literature on the ethical and social implications of the EHDS found that most of the issues discussed concern secondary use and that states that already have digitalization difficulties risk being disproportionately affected [320].

8.5 Romania in 2026

The description below reflects the official sources available on September 27, 2026. The infrastructure is undergoing rapid change.

8.5.1 The Electronic Health Record (DES)

The legal basis of the Romanian national electronic health record (Dosarul electronic de sănătate, DES) is Title IX^1 of Law No. 95/2006, introduced by Law No. 45/2019 (Articles 346^1-346^12). The law describes the DES as a set of electronic records consolidated at national level, with clinical, biological, diagnostic, and therapeutic data. Physicians access the record with the patient's consent. The exception is the "Sumar de urgență" (Emergency Summary) module, accessible without consent to physicians in emergency services (the on-call room, the emergency department or UPU, the emergency unit or CPU, and the ambulance service) and in primary care. Patients may expressly refuse the use of the DES, in which case their data are anonymized, without losing access to medical services [321]. In the terms of this chapter, the DES is a national EHR.

8.5.2 The Health Insurance IT Platform

The DES is part of the Health Insurance IT Platform (Platforma informatică din asigurările de sănătate, PIAS), together with the Single Integrated IT System (Sistemul informatic unic integrat, SIUI), the Electronic Prescription IT System (Sistemul informatic de prescripție electronică, SIPE), and the electronic health insurance card (cardul electronic de asigurări de sănătate, CEAS) [322]. The PNRR (National Recovery and Resilience Plan) project "Redimensionarea, standardizarea și optimizarea Platformei Informatice din Asigurările de Sănătate - PIAS" (Resizing, standardization, and optimization of the Health Insurance IT Platform - PIAS) was funded under Component 7 (Digital Transformation), investment I3.1, and was declared

completed on August 31, 2026. The approved value was 493,900 thousand lei excluding VAT, and the amount paid was 434,679 thousand lei excluding VAT. The completion announcement lists the new infrastructure, the migration of the existing platform, the e-Sănătatea Mea portal, the electronic medical forms, and the Register of Insured Persons [323]. The official wording is "PIAS modernizată" (modernized PIAS), that is, the same platform resized and optimized, not a system with a new name [323].

For patients, the visible result is the eSănătateaMea portal, launched on September 1, 2026. Access is through ROeID, the national digital identity system, and the portal displays insurance status, prescriptions issued in the last 5 years, health history for the last 10 years, reimbursement history, the status of the health card, and online appointments with providers [324]. In the terms of Section 8.1, the portal functions as a national PHR. The starting point is modest: in 2024, only 9.9% of individuals in Romania had accessed their health records online (Figure 8.1) [297].

8.5.3 The Health Card and the Electronic Prescription

The national health insurance card remains in use. The National Health Insurance House (CNAS) specified that, from September 1, 2026, the use of the electronic identity card does not become mandatory for providers, and insured persons' access to services is not conditional on it [325]. The electronic prescription works through SIPE, a component of PIAS, and issued prescriptions now also reach the patient portal [322,324]. The difference from the European model described in Section 8.3 lies at the organizational and semantic levels: for a Romanian prescription to be dispensed at a pharmacy in another state, SIPE must be connected to MyHealth@EU through the national contact point, an obligation that the EHDS sets for March 2029 [23]. One concrete semantic obstacle is terminology: as of September 27, 2026, Romania was not among the members of SNOMED International [151].

8.5.4 Hospital Digitalization through the PNRR

Also under Component 7, investment I3.3 funded IT systems and digital infrastructure for approximately 200 public health facilities, with a budget of 493,900,000 lei excluding VAT [326]. At the time of writing, the official sources consulted did not present the aggregate physical results of these projects.

Case Study

- On August 17, 2024, the "Pius Brînzeu" County Emergency Clinical Hospital of Timișoara (Spitalul Clinic Județean de Urgență "Pius Brînzeu" Timișoara) announced the project "Sănătatea, siguranța și confortul pacientului prin servicii integrate digitale la SCJU Timișoara" (Patient health, safety, and comfort through integrated digital services at SCJU Timișoara), code PNRR/C7/MS-733/I3.3, with a total value of 5,877,410.00 lei including VAT [327].
- The declared activities: storage equipment for data migration, terminals (desktops, medical tablets, barcode and QR code readers), mobile applications for physicians and nurses to facilitate the integration of data from the patient's electronic record, artificial intelligence software for the analysis of chest CT scans, applications for the flow of medical supplies and medicines, and staff training [327].
- The declared objective includes improving interoperability and data exchange within the hospital, with suppliers, with patients, and with other entities [327]. The project page, consulted on September 27, 2026, mentions neither the completion date nor the results obtained.
- Discussion questions: which activities belong to the technical level and which to the semantic level? What must be checked before 2029, when patient summaries become mandatory under the EHDS?

8.6 Usability, Safety, and Human Factors

An EHR can produce errors through the way it is built and used. An analysis of 2700 incident reports from more than 595 health facilities found 1508 medication errors associated with health information technology, half of which reached the patient. Almost all of them (97.3%) involved a usability problem, particularly data entry, workflow support, and alerts, and 80.5% were wrong doses [328]. A national survey of 1933 physicians attributed 38% of the variation in usability scores to differences between products and 51% to differences between hospitals using the same product [329]. The way a system is implemented seems to matter as much as the software itself.

8.6.1 Alert Fatigue

Decision support alerts are useful only if they are rare and relevant. Alert fatigue means ignoring alerts because of the large number of irrelevant warnings. At an academic medical center, 95.7% of 16,011 high-priority drug-drug interaction alerts were overridden, including 87.3% of those of the highest severity. Only 45.4% of the overrides were judged appropriate, and adverse events were more frequent after inappropriate overrides (9.4% versus 4.3%) [330]. Measurement of the phenomenon is not yet standardized: of 22 systematic reviews analyzed in 2026, only one defined alert fatigue operationally [331]. In the survey of family physicians cited earlier, the usefulness of alerts received the lowest rating [295].

8.6.2 Copy-Paste and Note Bloat

The ease of copying text from one note to another saves time, but it produces "note bloat" and can carry forward information that is no longer true, for example a "yesterday" copied into a note written on a different day [285]. An analysis of 104 million notes from a US academic health system showed that 50.1% of all text was duplicated from earlier notes about the same patient, and the proportion rose from 33.0% in 2015 to 54.2% in 2020. Almost half of the duplicated text came from a different author [332].

8.6.3 Safety Measures

The SAFER Guides, used by US hospitals to self-assess EHR safety, were revised in 2025: the number of recommended practices was reduced by 40%, and each recommendation was assigned a level of evidence [333]. Regardless of guidelines, any EHR can experience downtime, planned or unplanned, so every institution needs working procedures for the periods when the system is unavailable [285]. For the clinician, a few simple rules reduce the risk: verifying the patient's identity before each prescription, rereading copied text, and reporting useless alerts to the team that configures the system.

Common Pitfalls

- Believing that an "HL7-compatible" system automatically communicates with any other system. Without shared profiles and terminologies, messages can be formally correct and still be misunderstood.
- Confusing the EHR with the PHR. The EHR is managed by care providers, the PHR by the patient.
- Regarding FHIR as a database or a product. FHIR is a standard (resources and an API), implemented by various servers and applications.
- Treating the patient summary as a copy of the entire record. The IPS is intentionally minimal.
- Reflexively overriding all alerts. Among the irrelevant ones are alerts that prevent real adverse events.

Key Points

- The EHR is an active system, not an archive: it checks data, displays trends, and sends alerts.
- Structured data can be reused, while free text preserves reasoning. Most systems combine the two.
- Physicians spend a large part of their day in the EHR, and this time is increasing. Burnout is associated with this burden.
- Interoperability has four levels. The semantic level requires shared terminologies.
- HL7 v2 dominates in hospitals; FHIR dominates in new applications and in API-based exchange.
- The IPS has three mandatory sections: problems, allergies, medication.
- The EHDS requires Member States to exchange patient summaries and e-prescriptions from March 2029.

Self-Assessment Questions

1. (SA) Which level of interoperability is affected when two laboratories send the same test correctly formatted but with different LOINC codes?

a) technical
b) syntactic
c) semantic
d) organizational
e) physical

2. (SA) In the HL7 v2 segment `PID|1||123456^^^SPITAL_DEMO^MR||POPESCU^ION^^^^^L||19800115|M`, the value 19800115 represents:

a) the patient identifier
b) the date of admission
c) the date of birth
d) the ward code
e) the message number

3. (SA) Which of the following sections is mandatory in the HL7 FHIR IPS guide?

a) vital signs
b) social history
c) immunizations
d) allergies and intolerances
e) advance directives

4. (MA) Which of the following statements about HL7 FHIR are true?

a) the basic unit of exchange is the resource
b) resources can be represented in JSON or XML
c) searching for a patient's observations is done with a GET request
d) FHIR requires a single central database for all hospitals
e) profiles constrain resources for a particular use case

5. (MA) Regarding the documentation burden in the EHR, the studies cited in the chapter show that:

a) physicians spend approximately one third of their working day in the EHR
b) US clinicians spend more active time in the EHR than those outside the United States
c) artificial intelligence scribes completely eliminated after-hours work
d) after-hours EHR time was associated with higher odds of burnout
e) EHR time decreased steadily after 2019

CHAPTER 9

Chapter 9. Telemedicine: Forms, Evidence, Limits, and Legal Framework

Objectives

By the end of this chapter, the student will be able to:

1. define telemedicine, telehealth, and eHealth and assign a given service to one of the forms of telemedicine (teleconsultation, tele-expertise, telemonitoring, teleassistance), specifying the actors and the synchronous or asynchronous mode.
2. describe the technical components of a teleconsultation or telemonitoring system, including the security layers and the link with the electronic health record.
3. interpret the results of the main systematic reviews published in 2020-2026 for heart failure, hypertension, diabetes, psychiatry, dermatology, stroke, and radiology.
4. list the limits and risks of teleconsultation and recognize the situations that require an in-person examination or an emergency call.
5. explain the legal and reimbursement framework of telemedicine in Romania as of September 27, 2026, and the applicable European reference points.
6. write a complete teleconsultation note based on a template.

9.1 Definitions and Taxonomy

9.1.1 Telemedicine, Telehealth, eHealth

The World Health Organization (WHO) defines telemedicine as the delivery of healthcare services, where distance is a critical factor, by any healthcare professional using information and communication technologies (ICT) for the exchange of valid information for diagnosis, treatment, and prevention [334]. The definition has three elements: distance, the professional, and valid information. A message sent to a physician friend on a chat application does not automatically meet these conditions.

Telehealth is the broader notion. It includes telemedicine, but also the distance education of professionals and other applications in which electronic communications support health services [334]. eHealth covers any use of ICT in health: the electronic record, the e-prescription, registries, mobile applications. The relationship between the three terms is one of inclusion (Figure 9.1): every act of telemedicine is an act of eHealth, but the converse is not true.

9.1.2 Modes: Synchronous, Asynchronous, Continuous Monitoring

The main technical criterion is the moment at which the other participant responds. The 2022 WHO guide describes three modes [334]:

- **synchronous** (interactive, real-time): the patient and the physician, or two physicians, communicate simultaneously, by video or telephone.

- **asynchronous**, also called *store-and-forward*: data and images are sent and analyzed later by a specialist, usually in situations that are not urgent (dermatology, radiology, ophthalmology).
- **remote monitoring** (telemonitoring): connected medical devices transmit parameters, and the medical team receives alerts when values fall outside the set ranges.

The second criterion is the actors involved. Telemedicine can connect the patient to a professional (*client-to-provider*) or two professionals to each other (*provider-to-provider*). The WHO guide lists as frequent uses teleconsultation (remote clinical services for the patient), tele-expertise (a physician asks a specialist for a second opinion), teletriage (prioritizing emergency decisions based on transmitted parameters), and guided self-help, used in mental health [334]. Teleassistance refers to remote support for people living at home (alarm button, scheduled call, support for the caregiver). Its legal definition in Romania is discussed in Section 9.5.

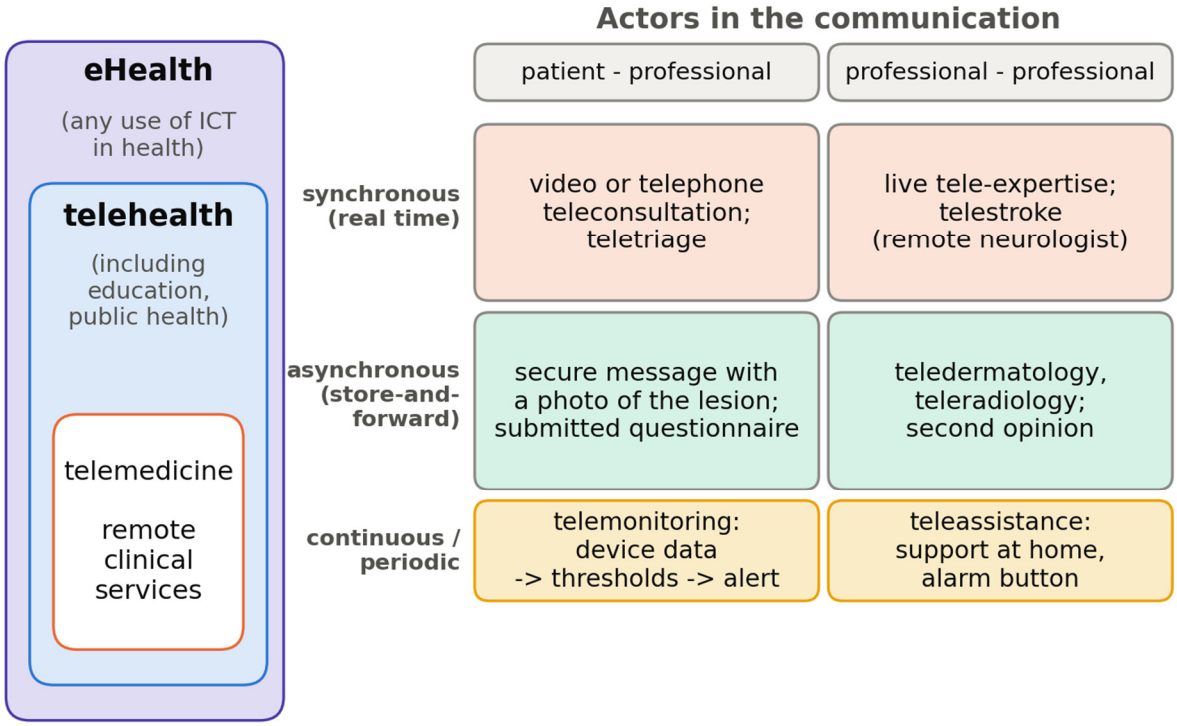


Figure 9.1. Taxonomy of telemedicine: the inclusion relationship between eHealth, telehealth, and telemedicine (left) and the forms of telemedicine by actors and mode (right). Based on [334,335].

Table 9.1. Forms of telemedicine, with examples. Source: [334,335].

Form	Actors	Mode	Example	WHO 2023 code
Teleconsultation	patient - physician	synchronous (video, telephone)	3-month follow-up for hypertension	2.4.1
Tele-expertise	physician - physician	synchronous or asynchronous	the family physician sends a dermoscopic image to the dermatologist	2.4.4
Telemonitoring	device - team	continuous or periodic	weight and blood pressure transmitted daily in heart failure	2.4.2
Transmission of medical data	patient or physician - specialist	asynchronous	radiograph sent for interpretation (teleradiology)	2.4.3
Teleassistance	patient, caregiver - service	synchronous, on demand	alarm button for an older person living alone	no separate code

9.1.3 The WHO Classification of Digital Interventions (2nd Edition, 2023)

To describe a project without ambiguity, in 2023 WHO published the second edition of the classification of digital interventions, services, and applications in health [335]. The classification has three axes: health system challenges (for example, "delayed provision of care"), digital interventions, and types of services and applications. The interventions are grouped by primary user: persons (group 1), healthcare providers (group 2), health management and support personnel (group 3), and data services (group 4).

Telemedicine has its own subgroup, 2.4, with four interventions: 2.4.1 consultations between a remote person and a healthcare provider, 2.4.2 remote monitoring of a person's health or diagnostic data by the provider, 2.4.3 transmission of medical data (images, notes, videos) to the healthcare provider, and 2.4.4 consultations for case management between healthcare providers [335]. For 2.4.1, the classification lists interactive and synchronous telemedicine as synonyms, and for 2.4.3, store-and-forward, that is, asynchronous telemedicine. A distinction that is useful for the student: self-monitoring carried out by the patient, without the data reaching the physician, is coded separately (1.4.2), in the persons group. A glucose meter that only displays values is not telemonitoring, even though it uses a digital device.

9.2 Technical Infrastructure

9.2.1 Videoconferencing and Bandwidth

Video teleconsultation requires a stable connection in both directions. The WHO guide points out that telemedicine usually requires more bandwidth than other digital interventions, so the digital divide is felt more strongly [334]. A study of 3,659,699 US veterans classified the connection according to the speeds reported to the regulatory authority: inadequate (download of at most 25 Mbps, upload of at most 3 Mbps), adequate, and optimal (at least 100/100 Mbps) [336]. After the onset of the pandemic, patients with an optimal connection used video mental health consultations more than those with an inadequate connection (incidence rate ratio 1.52, 95% CI 1.45-1.59) [336]. Consequently, every video service must have a fallback option (a telephone call) and a plan for when the connection drops.

9.2.2 Peripheral Devices

The remote physical examination relies on peripheral devices: a digital stethoscope, a dermatoscope attached to the phone, a video otoscope, and a connected blood pressure monitor and pulse oximeter. In an evaluation of 57 patients with heart murmurs, the agreement between remote auscultation with an electronic stethoscope and conventional auscultation for detecting murmurs was 90% [337]. In teledermatology, a meta-analysis of 155 studies showed that the use of the dermatoscope increased diagnostic concordance for skin cancers from 67% to 80% [338].

Not every clinical sign can be transmitted. A scoping review of 25 studies from the pandemic period identified three main types of virtual physical examination (musculoskeletal, head and neck, and chest), but also barriers: technical problems, doubts about effectiveness, the confidence of the nurse handling the device, and the patient's condition and level of health literacy [339].

9.2.3 Platforms and Security

A telemedicine platform brings together videoconferencing, secure messaging, document storage, scheduling, and, for telemonitoring, alert rules. Among the technical requirements, the WHO guide mentions user authentication and indicates that remote monitoring may require cloud

solutions with encrypted data [334]. The same guide specifies that audio or video recording of a consultation is done only with the patient's informed consent, and the archive must be protected [334].

The distinction between "hop-by-hop" encryption and end-to-end encryption (E2EE) is well illustrated by the documentation of Jitsi Meet, an open-source video platform. All audio and video traffic is encrypted on the network (DTLS-SRTP), but this encryption is removed on the server that relays the streams, where the packets pass through memory without being stored [340]. End-to-end encryption is available only in certain browsers and covers audio, video, and screen sharing, but not chat messages [340]. The practical consequence: physicians must know what the platform protects and what it does not, and must not send sensitive data through the unprotected channel.

Data concerning health are a special category of personal data (Article 9 of the General Data Protection Regulation, GDPR), and Article 32 requires the controller to implement technical and organizational measures appropriate to the risk, including encryption [118]. A review of telehealth interventions during the pandemic grouped 28 reported challenges into technical (14), non-technical (10), and privacy and policy (4) challenges, the most critical technical problems being those of networking, reliability, and system compatibility [341].

9.2.4 Interoperability with the Electronic Record

An isolated telemedicine service produces information that does not reach the patient's electronic health record (EHR). The WHO guide recommends avoiding "siloed" implementations by reusing common components: unique identifiers for patients, professionals, and facilities, terminology services, registries, and shared records [334]. The HL7 FHIR standard (Fast Healthcare Interoperability Resources, version R5 of 2023) provides the resources through which a platform can write observations (blood pressure, weight), documents (the consultation note), and appointments into the EHR [305]. Integration is feasible on a large scale: at Mayo Clinic, home video consultations in 2020 ran through a Zoom application integrated into the Epic electronic record [342].

Figure 9.2 summarizes the architecture. The devices at home transmit, usually via Bluetooth, to a *gateway* (the patient's phone or a hub), which sends the data over the internet to the platform. The platform applies the alert rules and writes the data into the EHR, and the clinician sees a dashboard, decides, and responds to the patient. Security is not a separate component but a layer that covers the entire chain.

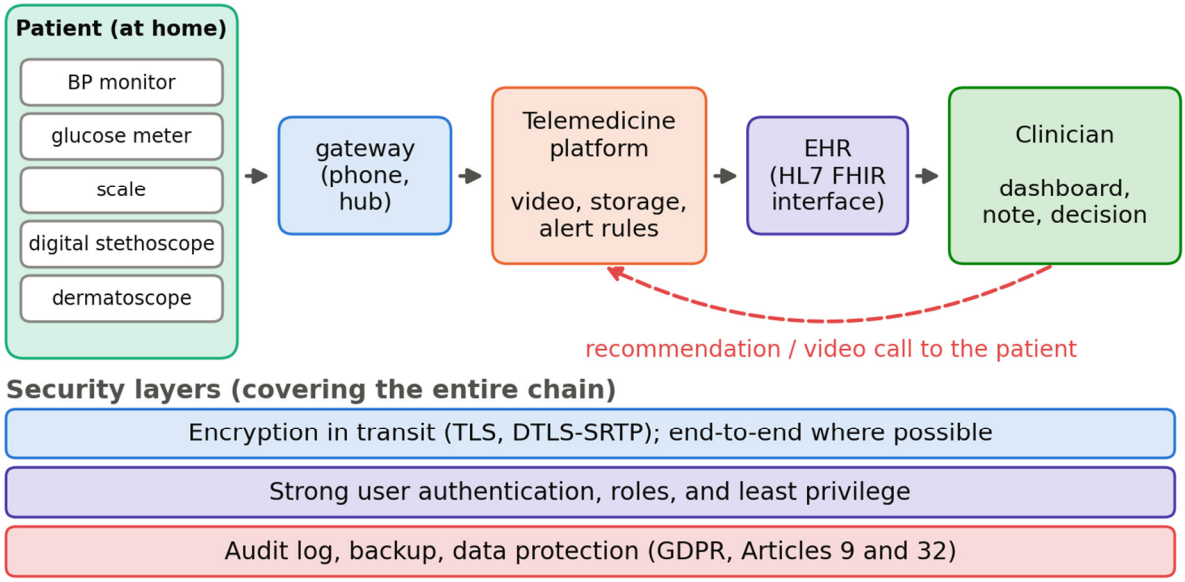


Figure 9.2. Architecture of a teleconsultation and telemonitoring system, with the security layers covering the entire chain. Teaching diagram based on [118,305,334].

9.3 Evidence of Effectiveness, 2020-2026

9.3.1 The COVID-19 Pandemic and Accelerated Adoption

The pandemic changed the scale of telemedicine in the first months of 2020. In a database of 16.7 million insured persons in the United States, 30.1% of all outpatient visits between March 18 and June 16, 2020, took place through telemedicine, and the weekly number of teleconsultations increased 23-fold compared with the prepandemic period [343]. The differences between specialties were large: 68% of endocrinologists used telemedicine, compared with 9% of ophthalmologists. By condition, the proportion of teleconsultations ranged from 53% for depression to 3% for glaucoma [343]. The difference has an obvious clinical explanation: depression is assessed by interview, glaucoma by instrumental measurements.

In Europe, the PRICOV-19 study included 5,065 family practices from 38 countries. Fewer than half (47.5%) used video consultations during the pandemic [344]. Use was highest in the United Kingdom, Luxembourg, the Nordic countries, and France (82.6-94.4%) and lowest in Portugal, Spain, Serbia, Bosnia and Herzegovina, Switzerland, and Czechia (11.1-23.1%). At the country level, the only statistically significant associated factor was access to affordable internet [344].

After the end of the public health emergency, some services were withdrawn. A longitudinal study of 1,001 US mental health centers found a decrease in the provision of telehealth from 81.6% to 79.0%, with more access restrictions for therapy and medication management [345].

The health crisis also reorganized the information flows in the surgical specialties. A narrative review of 23 studies on elective surgery for colorectal cancer during the pandemic, coauthored by the author of this textbook, reported decreases in the number of elective procedures ranging from 14% to 70% in various countries, shorter hospital stays, and longer waiting times until admission. In a surgical department in Timişoara, elective procedures for colon cancer decreased by 42% in

the first year of the pandemic [116]. Among the lessons formulated for a future crisis are the online organization of multidisciplinary meetings and the routine use of telemedicine for preoperative assessment and postoperative follow-up [116]. The study does not measure the effect of telemedicine on clinical outcomes. Its value for this chapter lies elsewhere: it shows that, when physical contact becomes risky, the flow of information (triage, prioritization scores, remote multidisciplinary boards) becomes part of the medical act.

9.3.2 Evidence by Specialty

Heart failure (HF) is the field with the most evidence. A meta-analysis of 92 studies (65 with noninvasive telemonitoring, 27 with implanted devices), with 36,549 patients followed for a mean of 11.5 months, found a 16% reduction in all-cause mortality (OR 0.84, 95% CI 0.77-0.93) and a 19% reduction in first HF hospitalization (OR 0.81) [346]. The authors note, however, that telemonitoring methods remain very diverse. A second meta-analysis, restricted to cardiovascular disease, found that the combination of remote monitoring plus remote consultation in HF reduced cardiovascular mortality (RR 0.83) and cardiovascular hospitalizations (RR 0.71), with no effect on all-cause hospitalizations and mortality. Remote consultation alone brought no benefit [347]. Figure 9.3 places the results side by side.

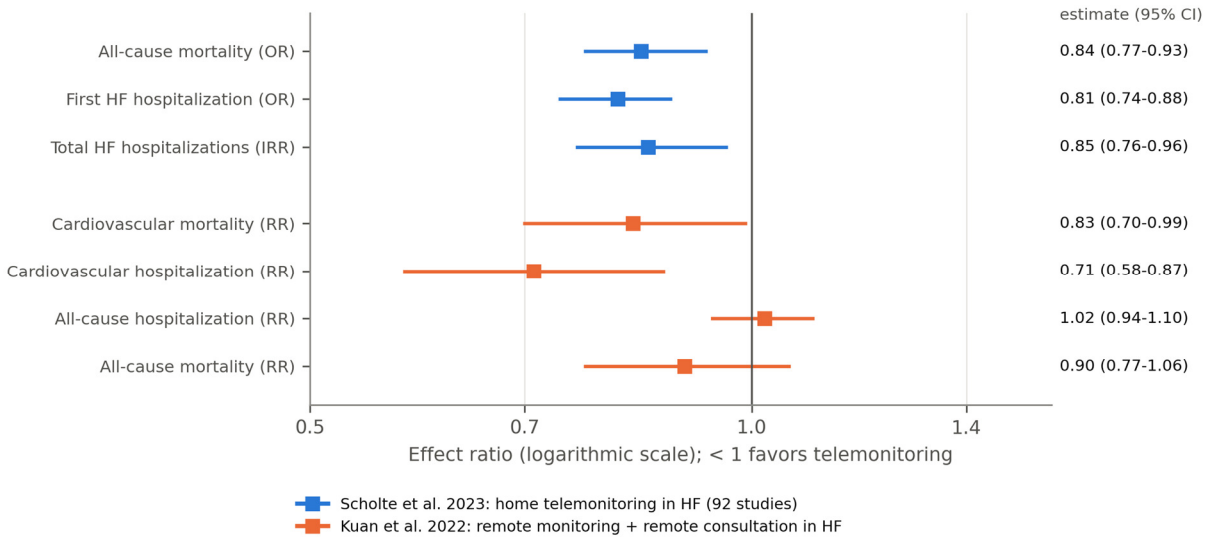


Figure 9.3. Pooled effects of telemonitoring in heart failure, reported in two meta-analyses. Values below 1 favor telemonitoring. OR, odds ratio; IRR, incidence rate ratio; RR, relative risk; CI, confidence interval. Data from [346,347].

Hypertension. A meta-analysis of 23 randomized trials with automatic transmission of home blood pressure readings reported a 6.0 mmHg decrease in office systolic blood pressure compared with usual care [348]. A network meta-analysis of 24 RCTs estimated for telemonitoring a 3.69 mmHg reduction in systolic blood pressure (95% CI 1.57-5.82) compared with usual care, but found no difference compared with home self-measurement without transmission [349]. The result suggests that part of the benefit comes from regular measurement, not only from transmission.

Type 2 diabetes mellitus. A review of 246 randomized trials, 168 of which allowed the effect to be calculated, found a mean decrease in HbA1c of 0.42 percentage points, with very high heterogeneity (I^2 93%). Interventions that included telemonitoring had larger effects [350]. Synchronous teleconsultation, analyzed separately in 27 studies, reduced HbA1c by 0.35 percentage points, with no effect on body mass index, blood pressure, or LDL cholesterol [351].

Chronic obstructive pulmonary disease. The 2021 Cochrane review (29 studies, 5,654 participants) is more cautious: remote monitoring added to usual care appears not to change the number of exacerbations or quality of life, with low or very low certainty, although COPD readmissions at 26 weeks were probably reduced (HR 0.42, moderate certainty) [352]. Pulmonary telerehabilitation, by contrast, gave results similar to center-based rehabilitation for the 6-minute walk distance (mean difference 0.06 m), with a completion rate of 93%, compared with 70% for conventional programs [353].

Psychiatry. A meta-analysis of 32 randomized trials (3,592 participants, 11 conditions) found no overall differences between telepsychiatry and face-to-face treatment. By diagnosis, telepsychiatry was superior in depressive disorders, and face-to-face care in eating disorders [354].

Dermatology. The meta-analysis cited above estimated a diagnostic concordance of 76% between teledermatology and in-person examination for all skin conditions and of 73% for skin cancers, with no differences between the synchronous and asynchronous modes [338].

Stroke. A meta-analysis of 12,540 patients found no differences between telestroke and conventional treatment for door-to-needle time, functional recovery (modified Rankin score 0-2, OR 1.06), 90-day mortality, or symptomatic intracerebral hemorrhage [355]. A German noninferiority study from 2026, in 15 hospitals from telestroke networks, compared video neurological ward rounds with in-person rounds in the same patients: full adherence to the quality indicators was 92% in teleneurology, compared with 54% for the bedside visit [356]. The study was not randomized, and the comparison concerns the quality of documentation, not clinical outcomes.

Radiology. In a prospective study of 7,209 on-call CT reports, clinically important discrepancies between the preliminary and the final report occurred in 6.1% of the in-hospital reports and in 5.8% of the teleradiology reports [357].

Table 9.2. Summary of the evidence by specialty (reviews and studies, 2020-2026).

Field	Source	Type, size	Main result
Heart failure	[346]	meta-analysis, 92 studies, 36,549 patients	mortality OR 0.84
Heart failure	[347]	meta-analysis, 34 studies in the quantitative analysis	CV mortality RR 0.83, consultation alone without effect
Hypertension	[348,349]	meta-analyses, 23 and 24 RCTs	systolic BP -6.0 and -3.69 mmHg
Type 2 diabetes	[350,351]	meta-analyses, 168 and 27 studies	HbA1c -0.42 and -0.35 percentage points
COPD	[352]	Cochrane, 29 RCTs	exacerbations unchanged, readmissions HR 0.42
Psychiatry	[354]	meta-analysis, 32 RCTs	no overall difference
Dermatology	[338]	meta-analysis, 155 studies	concordance 76%
Acute stroke	[355]	meta-analysis, 12,540 patients	similar outcomes
Radiology	[357]	prospective study, 7,209 reports	discrepancies 5.8% vs. 6.1%

9.3.3 What Works and What Does Not

Several regularities emerge from these data. Telemedicine has the best results when data transmission is linked to a team that reacts: in HF, monitoring followed by consultation reduced cardiovascular events, but consultation alone did not [347], and in diabetes the monitoring component amplified the effect [350]. The effects on intermediate parameters (blood pressure, HbA1c) are real but modest, on the order of a few mmHg or a few tenths of a percentage point. Heterogeneity is high [350], and the benefits in HF come mainly from studies with short follow-up [347]. Finally, equivalence with face-to-face care is more likely where the diagnosis relies on interview or images (psychiatry, dermatology, radiology) and less likely where palpation, fine auscultation, or instrumental measurements are needed.

9.4 Limits and Risks

9.4.1 Limited Physical Examination and Diagnostic Errors

The most direct limit is the lack of a complete physical examination. At Mayo Clinic, the provisional diagnosis made at a video consultation was confirmed by the subsequent in-person examination in 86.9% of the 2,393 new cases [342]. Concordance ranged from 77.3% in otorhinolaryngology to 96.0% in psychiatry, and by ICD-10 chapter it fell to 64.7% for diseases of the ear [342]. In other words, roughly one case in eight received a different diagnosis at the in-person examination. The WHO guide adds a limit that is harder to measure: the video channel reduces nonverbal cues, which affects trust and the checking of understanding [334].

9.4.2 Continuity of Care

The same technology can produce care of different quality, depending on who uses it. In a study of 27,686 children with acute respiratory infections, antibiotics were prescribed in 28.9% of the teleconsultations carried out by the patient's own family physician and in 37.2% of those offered by direct-to-consumer telemedicine companies [358]. Return visits within the next 1-2 days were also less frequent when the teleconsultation was carried out by the physician who knew the patient (5.0% vs. 8.0%) [358]. Telemedicine integrated into an existing physician-patient relationship is not the same thing as an episodic consultation with an unknown physician.

9.4.3 The Digital Divide

The patients with the greatest care needs are often the least prepared for telemedicine. An analysis of a representative sample of 4,525 older Americans estimated that 38% (approximately 13 million people) were not ready for a video consultation and 20% not even for a telephone one, because of inexperience with technology, hearing and communication deficits, or dementia [359]. Among people over 85 years of age, the proportion of those not ready for video reached 72%, and among those who rated their health as poor, 77% [359]. Telemedicine use was lower in communities with higher poverty [343], and a poor internet connection limited access to video consultations [336]. The WHO guide therefore calls for multiple access routes (several communication channels, access points such as kiosks, support from community health workers) and for an assessment of the effect on equity before implementation [334].

9.4.4 Confidentiality and Liability

A review of the ethical and legal literature from the pandemic found that the most frequently discussed issues were informed consent and autonomy (87% of articles), privacy (78%), data

protection and security (74%), and professional liability (70%) [360]. The authors conclude that telemedicine services can be used, for the time being, as tools complementary to traditional care. The WHO guide recommends clarifying, before launch, the licensing framework for professionals, medical liability, and the way services are audited [334]. Confidentiality risks also arise outside the platform: a patient speaking from an open-plan office, a family member listening in, a shared screen with another window open.

9.4.5 Fraud

In July 2022, the Office of Inspector General of the US Department of Health and Human Services issued a special fraud alert about fraudulent practices of some telemedicine companies, in the context of several fraud cases worth billions of dollars pursued by the authorities [361]. An analysis of recent US federal studies nevertheless concluded that there is little evidence that telehealth is associated with higher levels of fraud and abuse than conventional care [362]. For the student, the lesson is practical: the teleconsultation note must reflect exactly what was done, how long it lasted, and through which channel, because billing and auditing are based on it.

9.5 Legal and Reimbursement Framework

9.5.1 Romania, the Situation as of September 27, 2026

Telemedicine entered Romanian legislation through Government Emergency Ordinance (Ordonanța de urgență a Guvernului, OUG) No. 196/2020 (November 18, 2020; Official Gazette of Romania [Monitorul Oficial] No. 1108 of November 19, 2020), which introduced Articles 30^1-30^11 into Law No. 95/2006 on healthcare reform (Legea nr. 95/2006 privind reforma în domeniul sănătății) [363]. The law defines telemedicine as the totality of medical services provided at a distance, without the simultaneous physical presence of the medical staff and the patient, for diagnosis, treatment, monitoring, or prevention, through information technology and electronic communications, under safe conditions. The regulated forms are teleconsultation, tele-expertise, medical teleassistance, teleradiology, telepathology, and telemonitoring [363,364].

The methodological norms were approved by Government Decision (Hotărârea Guvernului, HG) No. 1133 of September 14, 2022 (Official Gazette No. 909 of September 15, 2022) [364]. Several provisions of Annex 2 to the norms have a direct effect on practice:

- Except in emergency medicine, accessing the platform or providing the service through telemedicine requires the patient's express consent, proven by any means of evidence, and the patient may refuse the service (Article 4).
- Telemedicine is not used when the physician and the patient do not speak the same language well enough (Article 9(1)).
- Services are scheduled in the same way as face-to-face services, so that the physician has access to the patient's record at the time of the consultation (Article 9(3)).
- The report of the service is entered in the consultation register or in the patient's record, with the means of communication used, the date and time interval, the prescriptions and documents issued, the identity of the participating medical staff, and any technical incidents (Article 10).

A terminological detail is worth remembering. In the norms, teleassistance means the real-time integration of data or opinions from remote professionals for the purpose of performing medical procedures [364]. The Romanian legal meaning is therefore close to an expert assisting with a procedure, not to the alarm button of social care described in Section 9.1.

Reimbursement from the National Single Fund for Social Health Insurance (Fondul național unic de asigurări sociale de sănătate) follows the framework contract (Contractul-cadru) approved by Government Decision No. 521/2023 and its implementing norms, Order No. 1857/441/2023 of the Ministry of Health and the National Health Insurance House (Casa Națională de Asigurări de Sănătate, CNAS) [365,366]. In the consolidated version of the norms published by CNAS in August 2025, the family physician may provide remote consultations, "by any means of communication," for diseases with endemic-epidemic potential that require isolation and for chronic conditions, under the same conditions as in the office, recording them in the register and in the patient's record together with the means of communication and the time interval [366]. In specialized outpatient care, remote consultations for chronic conditions are provided and reimbursed in the same way as office consultations, without a referral, and psychiatric therapies and genetic counseling may also be provided remotely [366,367]. Amendments to the norms made after August 2025 were not verified for this chapter, and the reader should consult the version in force on the date of the consultation.

Romania in 2026

- Funding for the digitalization of healthcare under the National Recovery and Resilience Plan (Planul Național de Redresare și Reziliență, PNRR) falls under Component 7, Digital Transformation, investment I3, "Development of the eHealth and telemedicine system," not under the component dedicated to health [368].
- Sub-investment I3.3 funded information systems and digital infrastructure for 200 public healthcare facilities, with an allocation of 493,900,000 lei excluding VAT [326].
- The project to modernize the CNAS information platform (I3.1) was declared completed on August 31, 2026 [323].
- As of September 27, 2026, no verified official sources were available on the physical targets achieved in telemedicine through the PNRR.

9.5.2 European Reference Points

The organization of health services remains a competence of the Member States, but several acts of the European Union apply directly to telemedicine. Directive 2011/24/EU on the application of patients' rights in cross-border healthcare, which is in force, establishes in Article 3(d) that, in the case of telemedicine, healthcare is considered to be provided in the Member State where the healthcare provider is established [369]. This rule has consequences for the applicable law and for liability. The directive also allows the Member State of affiliation to apply to care received through telemedicine the same reimbursement conditions as in its own system (Article 7(7)) [369].

Regulation (EU) 2025/327 on the European Health Data Space (EHDS), published in March 2025, recognizes telemedicine as an increasingly important tool and organizes the cross-border exchange of health data through the MyHealth@EU infrastructure [23]. The GDPR remains applicable to all forms of telemedicine [118].

9.6 Documenting the Teleconsultation

An undocumented teleconsultation is, from a medical and legal point of view, a teleconsultation that did not take place. In Romania, the minimum content of the record is set by Article 10 of Annex 2 to the telemedicine norms (Section 9.5) [364]. The WHO guide requires

mechanisms for the mutual identification of the physician and the patient and for tracking previous exchanges and decisions [334]. The same guide considers that, in most countries, consent is implied when the patient requests the service, and that only a teleconsultation initiated by the professional requires documented consent, with explicit communication of the limits of telemedicine [334]. The Romanian norms are stricter: outside emergencies, consent must be express and proven, regardless of who initiates the consultation [364]. The Australian practice guidelines for teledermatology, updated in 2025, include separate chapters on patient selection and consent, image standards, privacy, documentation, and image retention [370].

A frequent weak point is the recording of safety-netting advice, that is, the instructions about which signs should prompt the patient to come back or to call 112. In an analysis of 295 recorded consultations in family medicine, safety-netting advice was given verbally in two-thirds of the consultations but was documented in only one-third, covering 20.3% of the problems assessed [371]. In telemedicine, where the examination is incomplete, these instructions matter more than in the office.

A teleconsultation note contains at least the following:

1. **identification**: the patient (name, date of birth, or another verified identifier), the physician (name, specialty, physician's seal code, i.e., the Romanian parafă), the date, and the start and end times.
2. **the patient's location** at the time of the consultation and a telephone number for calling back.
3. **the mode**: video, telephone, asynchronous message, the platform used, and the persons present.
4. **consent** to the remote consultation and, separately, to recording, if any.
5. **clinical data**: the reason, the history, and what could be observed or measured (including values transmitted by devices).
6. **limitations**: what could not be examined and the impact on the conclusion.
7. **assessment and plan**: presumptive diagnosis, recommendations, prescriptions, investigations, and referrals.
8. **safety**: the warning signs explained and the time frame for reassessment or for an in-person examination.
9. **the physician's signature** (electronic, where the system allows it).

Example

Teleconsultation note (fictitious patient, invented data). Date: October 14, 2026, 10:05-10:22. Physician: Dr. A.B., internal medicine, seal code (parafă) [fictitious]. Patient: I.P., 67 years old, identity verified by date of birth and the personal numeric code (CNP) on file. Location: home, Timiș County, backup telephone number confirmed. Mode: video on the institution's platform, no recording, patient alone in the room. Express consent to the remote consultation, given verbally and recorded. Technical incidents: none. Reason: hypertension follow-up. Values transmitted over the last 7 days: mean 148/88 mmHg. No headache, dyspnea, or chest pain. Observed on video: fluent speech, no visible ankle edema (limited examination). Limitations: auscultation and blood pressure measurement by the physician could not be performed. Plan: the dose of the antihypertensive is increased according to the regimen; self-measurement twice a day continues. E-prescription issued. Safety: the patient was told to call 112 in case of chest pain, weakness of a limb, or speech disturbances. Video reassessment in 2 weeks, in-person examination with ECG in one month. Electronic signature: Dr. A.B.

Common Pitfalls

- Considering any application that displays health data to be telemedicine. Without a professional who receives and uses the data, it is self-monitoring.
- Confusing encryption in transit with end-to-end encryption. Some platforms, for example Jitsi Meet without E2EE enabled, decrypt the stream on the server.
- Reading a favorable meta-analysis as a guarantee for any product. The effects depend on the monitoring component, on the team that reacts, and on the population.
- Leaving out of the note the limitations of the examination and the safety-netting advice.
- Assuming that all patients can use video. Older adults and people with hearing deficits or without a stable internet connection need alternatives.
- Continuing the teleconsultation when warning signs appear, instead of directing the patient immediately to emergency care.
- Using the term teleassistance in its social care sense when discussing Romanian legislation, in which medical teleassistance means real-time support from an expert for performing a procedure.

Key Points

- Telemedicine is the provision of medical services at a distance through ICT, synchronously, asynchronously, or through monitoring.
- The 2023 WHO classification codes telemedicine in subgroup 2.4, with four interventions.
- The clearest benefits appear in telemonitoring for heart failure, when the data reach a team that reacts.
- In hypertension and diabetes, the effects are real but modest.
- In psychiatry, dermatology, and radiology, the results are close to those of face-to-face care.
- Roughly one diagnosis in eight made by video changed at the in-person examination in a cohort of more than 2,000 patients.
- The teleconsultation note records the mode, consent, the limitations, and the warning signs explained.
- In Romania, telemedicine is regulated by OUG No. 196/2020 and by the norms approved by HG No. 1133/2022, which require express consent outside emergencies.

Self-Assessment Questions

1. (SA) A family physician photographs a pigmented lesion with a dermatoscope and sends the image to a dermatologist, who replies the next day. The form of telemedicine is:

a) synchronous teleconsultation
b) asynchronous tele-expertise
c) telemonitoring
d) teleassistance
e) self-monitoring

2. (SA) In the 2023 WHO classification, remote monitoring of the patient's data by the healthcare provider has the code:

a) 1.4.2
b) 2.4.1
c) 2.4.2
d) 2.4.4
e) 4.1.2

3. (SA) In the meta-analysis by Scholte et al. (2023), home telemonitoring in heart failure was associated with:

a) an increase in all-cause mortality
b) a reduction in all-cause mortality of approximately 16%
c) no difference for first hospitalization for heart failure
d) effects only for implanted devices
e) a 50% reduction in mortality

4. (MA) Which of the following items must be recorded in a teleconsultation note?

a) the mode (video, telephone) and the platform used
b) the limitations of the remote examination
c) the safety-netting advice and the warning signs explained
d) the password of the patient's account on the platform
e) consent to the remote consultation

5. (MA) Which of the statements about the limits of telemedicine are supported by the data presented in the chapter?

a) the provisional diagnosis made by video was confirmed at the in-person examination in approximately 87% of cases, in a Mayo Clinic cohort
b) approximately 38% of older Americans were not ready for a video consultation
c) encryption in transit guarantees that the server cannot access the stream
d) pediatric teleconsultations offered by direct-to-consumer companies were associated with more antibiotic prescriptions than those carried out by the family physician
e) fraud is proven to be more frequent in telemedicine than in conventional care

CHAPTER 10

Chapter 10. Mobile Health, Wearable Devices, and Language Models in the Consultation

Objectives

By the end of this chapter, the student will be able to:

- distinguish a wellness application from an application that is a medical device, starting from the intended purpose stated by the manufacturer
- interpret the results of a meta-analysis on the effectiveness of mobile applications (standardized mean difference, NNT, certainty of evidence)
- list the sensors of a wearable device and judge, for each of them, what level of accuracy has been documented
- apply the steps of the MDCG 2019-11 guidance and Rule 11 of the MDR to assign an application to a risk class
- describe the path of a measurement from the sensor to the electronic record and the standards that make it possible (ISO/IEEE 11073, Bluetooth, FHIR)
- use a large language model responsibly: verify statements and references, protect patient data, and disclose the use

The phone in the patient's pocket has become, without anyone having planned it, the most widespread measuring instrument in medicine. It counts steps, receives the pulse from the watch, displays the glucose level from the sensor on the arm and, since the end of 2022, can hold a fluent conversation about symptoms. The chapter follows this sequence: the application, the sensor, the regulatory framework, the transfer standards and, finally, the language models used in the consultation. The methods of artificial intelligence proper (machine learning, evaluation, explainability, the AI Act) are covered in Chapter 13. Telemonitoring organized by a medical service was described in Chapter 9.

10.1 Mobile Health

10.1.1 Definitions and Categories of Applications

The global strategy of the World Health Organization (WHO) defines digital health as the field of knowledge and practice associated with the development and use of digital technologies to improve health. The document specifies that it expands the concept of eHealth to include digital consumers and a wider range of smart and connected devices [36]. Mobile health (mHealth) is the part of this field that uses phones, tablets, and wearable devices.

For the physician, the most useful classification is not the technical one but the one based on the intended purpose. The European guidance MDCG 2019-11, revised in June 2025, states that software intended exclusively for non-medical purposes, including wellness or fitness applications, does not qualify as medical device software [372]. The same guidance gives

examples of software that does qualify: a smartwatch application that sends alarms when it recognizes irregular heartbeats, a program that calculates the insulin dose and controls the pump, and an application that supports the treatment of schizophrenia through symptom monitoring and personalized interventions [372]. For applications without a medical purpose, the Regulation on the European Health Data Space (EHDS) introduced its own definition, that of a wellness application, that is, software, alone or in combination with hardware, intended by the manufacturer to be used by a natural person for the processing of electronic health data for purposes other than the provision of healthcare [23].

Three practical categories follow. The first comprises lifestyle applications (steps, sleep, nutrition, meditation) with no declared medical purpose. The second brings together applications that are medical devices, with CE marking, from glucose diaries with dose calculation to prescribed digital therapeutics. The third, narrower category includes digital medical devices reimbursed by a public insurance system, the most studied example of which is the German DiGA model (Section 10.3.3).

10.1.2 What the Evidence from 2020-2026 Shows

The number of randomized trials has grown so much that synthesis is now done at the level of reviews of reviews (umbrella reviews). The largest, published in 2026, included 78 systematic reviews with meta-analysis, covering 496 primary studies and 177,373 adult participants with, or at risk of, noncommunicable diseases [373]. Compared with control groups, the applications reduced systolic blood pressure (standardized mean difference, SMD -0.444) and diastolic blood pressure (SMD -0.414), glycated hemoglobin (SMD -0.587), weight, and sedentary time. The effects on anxiety (SMD -0.215) and depression (SMD -0.097) were small. No effect was found on cardiovascular mortality or on blood lipids, and heterogeneity was very high (I^2 of 91-96% for blood pressure and the glycemic indicators) [373].

A second synthesis, with 34 meta-analyses and 235 randomized trials from 52 countries, applied stricter criteria for the strength of evidence. The only "convincing" effect was the reduction in glycated hemoglobin in type 2 diabetes (d = 0.44), and 42% of the effects analyzed were not statistically significant. The authors also noted that adverse effects are rarely reported [374].

By disease, the picture is as follows:

- **Type 2 diabetes.** In a meta-analysis of randomized trials, app-based interventions reduced HbA1c by 0.42 percentage points compared with control, SMS-based interventions by 0.37, and web-based interventions had no statistically significant effect (-0.09) [375].
- **Hypertension.** A review of 74 studies and 92,686 participants found decreases in systolic and diastolic blood pressure for mHealth interventions, and medication adherence improved only in the mHealth group [376].
- **Smoking cessation.** A 2026 meta-analysis (31 studies, 12,802 participants) estimated that an application used alone might increase continuous abstinence at 6 months compared with no support or minimal support (RR 2.85), but the certainty of the evidence was low, and the comparison was based on only 4 studies [377].
- **Depression and anxiety.** In 176 randomized trials, the applications had small effects on depressive symptoms (g = 0.28, NNT 11.5) and on generalized anxiety symptoms (g = 0.26, NNT 12.4). The effects were larger when the application contained elements of cognitive behavioral therapy [378].

What do these figures mean for a student? An NNT of 11.5 means that, on average, 11-12 people must use the application for one additional person to have a clinical response compared with control. This is a real but modest effect, obtained under the controlled conditions of a trial.

10.1.3 Application Quality and User Retention

In app stores, quality cannot be inferred from the number of stars. The most widely used instrument for assessing it is the Mobile App Rating Scale (MARS), which scores four dimensions: user engagement, functionality, aesthetics, and information quality. A validation study of 1,299 applications from 15 health domains confirmed the structure of the scale, with good to excellent reliability (omega coefficient 0.79-0.93) and high inter-rater agreement (ICC 0.82) [379]. A 2024 umbrella review found MARS to be the most frequently used framework, but observed that most frameworks do not cover all the criteria identified. Data protection and security were the most frequently mentioned criterion, in 13 of 15 reviews [380].

The second problem is attrition. A systematic review of 99 studies on applications for noncommunicable diseases calculated a mean adherence score of 56.0%, relative to the use intended by the manufacturer [381]. The factors associated with better adherence were personalized content, individualized notifications, a stable and easy-to-use design, and complementary human support. A 2026 meta-analysis of randomized trials of applications for depression and anxiety found an initial uptake of 92.4%, but an adherence of only 61.8% and a loss to follow-up of 28.4% at the follow-up assessment. The losses were smaller where there were reminders and human contact [382]. These figures come from studies in which participants are selected, motivated, and actively followed up. For this reason, they should not be read as estimates of real-world use.

10.2 Wearable Devices and Sensors

10.2.1 What a Wearable Device Measures

A wearable device is a sensor worn on the body, usually on the wrist, the finger, the chest, or the arm, that transmits data to a phone. The complete path of a measurement, from the sensor to the physician, is shown in Figure 10.1. The points marked on the diagram (consent, encryption, access control) are the places where it is decided whether the data can be used clinically.

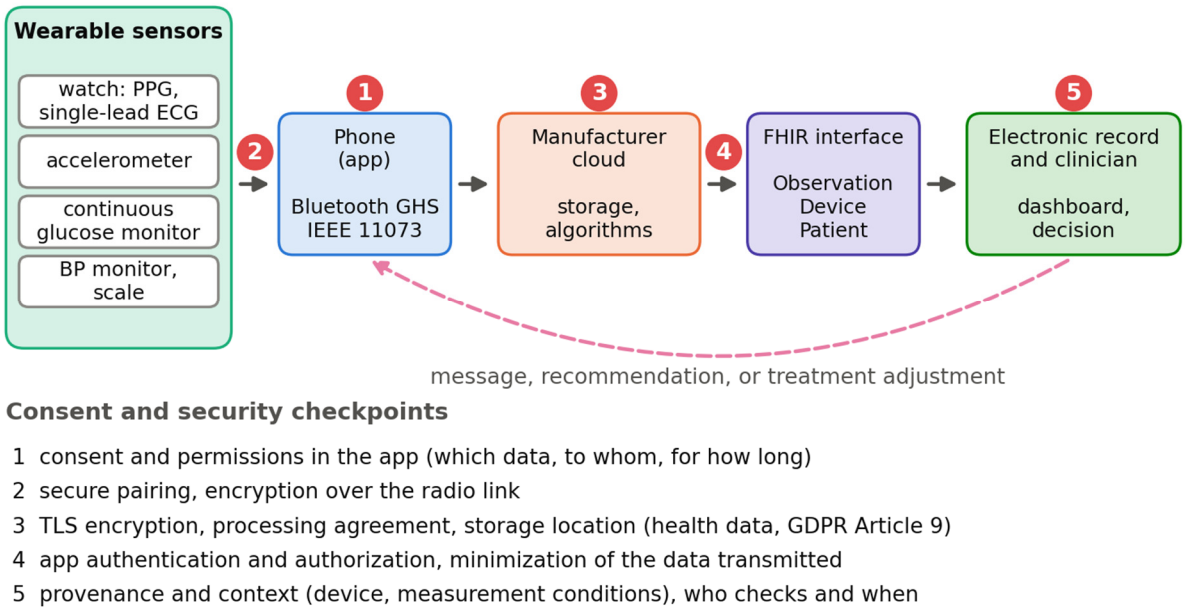


Figure 10.1. The flow of patient-generated data: sensor, phone, manufacturer's platform, electronic record, and clinician, with the consent and security checkpoints. Teaching diagram based on [118,383-385].

The main sensors are the following:

- **The accelerometer** (often combined with the gyroscope) records movement along three axes. The step count, the intensity of activity, falls and, indirectly, sleep are derived from it.
- **Photoplethysmography** (PPG) emits green or red light into the skin and measures the variation of the reflected light with each pulse wave. Heart rate, rhythm irregularity and, with different wavelengths, oxygen saturation (SpO2) are calculated from the signal.
- **The single-lead electrocardiogram** is obtained when the user touches the electrode on the watch with a finger and closes the circuit between the two arms. The result is a short tracing similar to lead I.
- **The temperature sensor** measures skin temperature, not core temperature. Its usefulness lies mainly in trends.
- **The continuous glucose monitor** (CGM) has a filament inserted into the subcutaneous tissue and measures interstitial glucose every few minutes. By its intended purpose, it is a medical device.

Table 10.1. Sensors of wearable devices, documented accuracy, and evidence-supported clinical uses. Source: [238,242,243,386-388].

Sensor	What it measures	Reported accuracy	Evidence-based use
Accelerometer	steps, activity, sleep	steps underestimated (error -9% to 12%), sleep overestimated	recovery after surgery, physical activity
PPG	pulse, rhythm, SpO2	pulse with a mean error of ±3%	AF detection (sensitivity 95%, specificity 97%)
Single-lead ECG	tracing similar to lead I	for AF, comparable to PPG	checking an irregular rhythm
Skin temperature	skin temperature, trends	not core temperature	postoperative recovery, in combination
Optical SpO2	oxygen saturation	mean difference of up to 2%	indicative, not diagnostic
Cuffless blood pressure	BP estimated from PPG	non-standardized validation	not yet recommended clinically
CGM	interstitial glucose	medical device, effectiveness in RCTs	reduced HbA1c in type 1 and type 2 diabetes

10.2.2 Accuracy and Validation

A "living" umbrella review (continuously updated) brought together 24 systematic reviews and 249 validation studies with more than 430,000 participants [386]. Only about 11% of the devices on the market had been validated for at least one measurement. For heart rate, the mean error was ±3%, and for arrhythmia detection the pooled sensitivity and specificity were 100% and 95%, respectively. Steps were generally underestimated, as was energy expenditure; maximal oxygen consumption was overestimated, and total sleep time was overestimated, with errors typically above 10%. For SpO2, the mean absolute difference reached 2.0% [386].

Skin color matters for optical sensors. A meta-analysis of 23 pulse oximetry studies (197,353 SpO2-arterial blood pairs) showed that the root mean square error exceeded the regulatory threshold of 3% in all pigmentation groups (3.96%, 4.71%, and 4.15% for light, medium, and dark skin) [242]. For the pulse measured by wearable devices, the same analysis found no systematic error related to pigmentation, although the limits of agreement were wide.

For blood pressure measured without a cuff, the evidence is weaker. In a review of 16 studies and 15 devices, only 8 devices had a mean error below 5 mmHg. Validation according to standard protocols and in real life remains the condition for their clinical use [387].

At the opposite end lies CGM, evaluated in dozens of randomized trials. A meta-analysis of 64 randomized trials found for real-time CGM, compared with self-monitoring with a glucose meter, a reduction in HbA1c of 0.24 percentage points in type 1 diabetes and of 0.40 in type 2 diabetes. Time in the target range (70-180 mg/dL) increased by approximately 5.5 percentage points in both types of diabetes [388].

10.2.3 Clinical Uses: Arrhythmias and Postoperative Monitoring

Atrial fibrillation (AF) is the best-documented application. A 2025 meta-analysis of 26 studies and 17,349 patients reported for smartwatches a pooled sensitivity of 95% and a specificity of 97%, with comparable performance between PPG-based and ECG-based models [238]. The Fitbit Heart Study, a clinical trial conducted entirely remotely, enrolled 455,699 adults. The irregular rhythm notification occurred in 1% of participants and in 4% of those over 65 years of age. Of those who subsequently wore an analyzable ECG patch, 32.2% had AF, and the positive predictive value of a notification occurring during patch monitoring itself was 98.2% [235]. The figures also show the limit: in approximately two-thirds of the notified participants who wore the patch, it did not record AF during the week of monitoring. The diagnosis is confirmed on an ECG, and the decision to anticoagulate belongs to the physician.

The second, more recent use is monitoring recovery after surgery. A systematic review by the author of this book analyzed continuous monitoring with wearable sensors after colorectal surgery [243]. Nine studies with 778 patients were included: one randomized trial, seven prospective cohorts, and one retrospective analysis. Five studies used only the step count, and four combined it with heart rate or skin temperature, over a median wear time of 6 days. At least 1,000 steps on the first postoperative day were associated with a shorter length of stay (OR 0.63), and an ERAS (Enhanced Recovery After Surgery) protocol supplemented with smart wristbands shortened the protocol-defined hospital stay by 1.1 days. Preoperative inactivity (fewer than 5,000 steps per day) and incomplete return to the baseline activity level predicted complications and 30-day readmission. The review concluded that such monitoring is feasible and its data clinically useful, but that standardized outcome definitions, open signal-processing algorithms, and multicenter validation are lacking [243].

10.3 Software as a Medical Device

10.3.1 Qualification: When an Application Becomes a Medical Device

Regulation (EU) 2017/745 on medical devices (MDR) also applies to software [284]. The practical question, "is my application a medical device?", is answered in the MDCG 2019-11 rev.1 guidance [372]. The central criterion is the medical intended purpose stated by the manufacturer. The risk of causing harm is not, in itself, a qualification criterion. The guidance proposes five decision steps:

1. Is the product software according to the definition of the guidance?
2. Is it an MDR Annex XVI device, an accessory for a medical device, or software driving or influencing the use of a (hardware) medical device? If so, it follows the regime of that device.
3. Does it perform an action on data different from storage, archival, communication, simple search, or lossless compression?
4. Is the action for the benefit of an individual patient? Software that only aggregates population data, reproduces general guidelines, or serves epidemiological registries does not meet this condition.
5. Does the software fall within the definition of a medical device (diagnosis, prevention, monitoring, prediction, prognosis, treatment, or alleviation of a disease)?

Figure 10.2 translates these steps and Rule 11 into a decision tree.

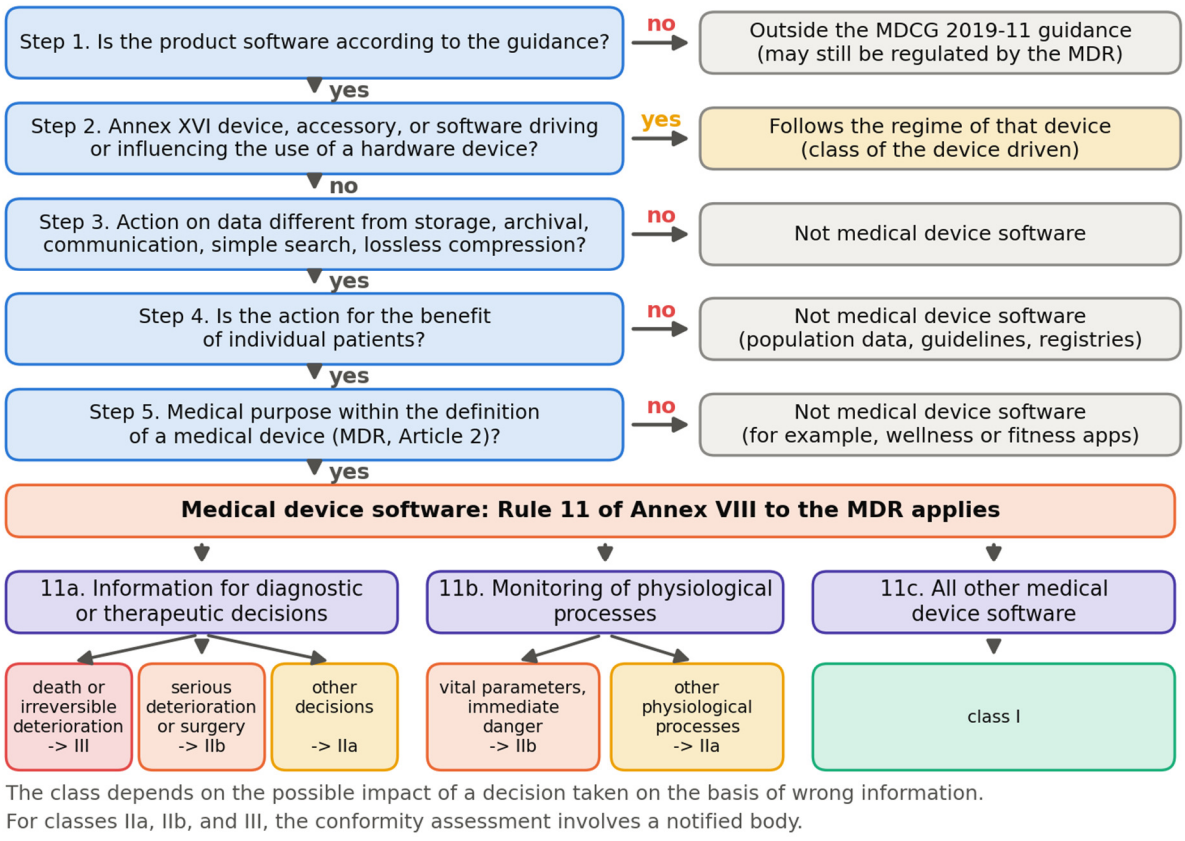


Figure 10.2. Decision tree for the qualification and classification of an application as medical device software, following Steps 1-5 of the MDCG 2019-11 rev.1 guidance and Rule 11 of Annex VIII to the MDR. The diagram is for teaching purposes and does not replace the conformity assessment. Based on [284,372].

10.3.2 Classification under Rule 11

Rule 11 of Annex VIII to the MDR has three sub-rules [372]. Sub-rule 11a concerns software intended to provide information which is used to take decisions with diagnosis or therapeutic purposes. Such software is in class IIa, but moves up to class IIb if a wrong decision may cause a serious deterioration of a person's state of health or a surgical intervention, and to class III if it may cause death or an irreversible deterioration. Sub-rule 11b concerns software intended to monitor physiological processes, classified in class IIa, except for the monitoring of vital physiological parameters whose variations could result in immediate danger to the patient (class IIb). The vital parameters listed by the guidance are respiration, heart rate, cerebral functions, blood gases, blood pressure, and body temperature. Sub-rule 11c places all other medical device software in class I.

Table 10.2. Examples of the classification of medical device software from Annex IV of the MDCG 2019-11 rev.1 guidance. Source: [372].

Intended purpose	Class	Rule
image analysis for therapeutic decisions in acute stroke	III	11a
application that analyzes heartbeats, detects abnormalities, and informs the physician for diagnosis	IIb	11a
diagnostic score for depression from the symptoms entered	IIb	11a
continuous surveillance of vital functions in anesthesia or intensive care	IIb	11b
home monitoring of certain vital signs in routine checks	IIa	11b
ranking of chemotherapy options for the physician	IIa	11a
fertility application with a red, green, or yellow indicator	I	11c
application that converts symbols into speech for communication disorders	I	11c

The class determines the route to CE marking. For classes IIa and IIb, the conformity assessment procedure necessarily involves a notified body, which checks development, production, and post-market surveillance [389]. The same medical purpose can lead to different classes if the context changes: the guidance specifies that its examples are indicative and that any change in the intended purpose or in the clinical situation may change the classification [372]. When the application uses machine learning, the obligations of the AI Act are added. After the 2026 amendment by Regulation (EU) 2026/1744 (Digital Omnibus), the obligations for high-risk AI systems integrated into products, and therefore also into medical devices, apply from August 2, 2028. The details are in Chapter 13 [22,197].

10.3.3 Reimbursed Applications: The German DiGA Model

Germany was the first country to include physician-prescribed health applications in statutory health insurance reimbursement, starting in September 2020 [390]. Digital health applications (Digitale Gesundheitsanwendungen, DiGA) are class I or IIa medical devices with CE marking, assessed by the Federal Institute for Drugs and Medical Devices (BfArM) in a fast-track procedure lasting at most three months. They can be prescribed by physicians or psychotherapists and are reimbursed for approximately 73 million insured persons [391]. The manufacturer must demonstrate a positive effect on care: a medical benefit or a patient-relevant improvement in the structure and processes of care [392]. Since March 26, 2024, under the German Digital Act (DigiG), class IIb devices can also become DiGA. For these, provisional listing in a testing phase is not provided for, so the medical benefit must be proven at the time of submission [389].

Independent analyses show both the progress and the limits of the system. On July 1, 2024, there were 56 listed DiGA, 35 permanently and 21 provisionally. The median initial price for a three-month prescription was 514 euros, reduced after negotiation to a median of 221 euros, and the number of prescriptions per reporting period increased from approximately 41,000 to 209,000 [392]. All 35 permanently listed applications had proven their effect through randomized trials [392]. The quality of these trials has, however, been criticized. For six DiGA in the neurological and mental health fields, the risk of bias was high, mainly due to lack of blinding, and in most trials attrition was higher in the intervention groups [390].

10.3.4 Wellness Applications in the EHDS

The EHDS introduces a new obligation for wellness applications that claim interoperability with an electronic health record system. They must bear a label issued by the manufacturer, indicating the categories of data for which compliance with the essential requirements and common specifications has been confirmed. The label is valid for at most three years and, for applications consisting only of software, is displayed in digital format in the application (Article 47) [23]. Under Article 105, these provisions apply from March 26, 2029, for the first priority categories of data. The label concerns interoperability and is not equivalent to the CE marking of a medical device.

10.4 Device Interoperability

10.4.1 The ISO/IEEE 11073 Family

The data from a device are useful only if the system that receives them knows what each value means. The ISO/IEEE 11073 family of standards solves this problem through a common nomenclature, defined in the base standard ISO/IEEE 11073-10101 and extended by other parts of the family. A well-defined example is Part 10103, republished by ISO in 2025, which extends the nomenclature for implantable cardiac devices: pacemakers, defibrillators, cardiac resynchronization devices, and implantable monitors [225]. The standard defines the terms needed to transmit a clinically relevant summary of a device interrogation, with identifiers that begin with the prefix `MDC_IDC_`. Its stated purpose is data exchange between the manufacturers' proprietary interrogation systems and the clinic's electronic record, and the nomenclature can be used together with the HL7 standards [225]. In practice, a remote follow-up report of a defibrillator can reach the patient's record as discrete data, not as a scanned document.

10.4.2 Bluetooth Profiles for Health

Many personal-use devices (blood pressure monitors, glucose meters, scales, watches) communicate with the phone via Bluetooth. In June 2023, the organization that manages the standard, the Bluetooth SIG, adopted the Generic Health Sensor Profile (GHSP), intended for various types of personal health devices, which transmits health observations using the ACOM content model defined in IEEE 11073-10206 [383]. The stated advantage is that new types of sensors can be added without changing the infrastructure that receives the data.

10.4.3 FHIR and Patient-Generated Data

Patient-generated health data (PGHD) are data collected by patients or their devices outside of contact with the health service. To enter the electronic record, they must be translated into the HL7 FHIR resources described in Chapter 8. The HL7 implementation guide for personal health devices (Personal Health Device FHIR IG), version 2.0.0 of April 2026, is based on the ACOM

model of IEEE 11073-10206 and accepts data arriving via Bluetooth GHS [384]. A case study with a Garmin watch examined the mapping of pulse, steps, sleep, and activity level to the Observation, Device, and Patient resources, as well as the feasibility of transferring them to FHIR servers [385].

Technical mapping alone does not solve the problem of trust. Clinicians want to know in what context a measurement was made and with what device, that is, its provenance. An ontology-based model proposed capturing this information at the very moment of collection and transmitting it, together with the value, in FHIR format [393]. Security is the second layer. A team from Ștefan cel Mare University of Suceava described in 2026 a *zero trust* architecture, with device-bound certificates, mutual TLS, and a layer of adapters that transforms HL7 v2, IEEE 11073, and proprietary message formats into FHIR resources. The authors state that the results represent an initial validation, not a full penetration test [394]. The legal basis for processing remains the one described in Chapter 5, because data concerning health are a special category of personal data [118].

10.5 Large Language Models in the Consultation

10.5.1 What a Large Language Model Is

A large language model (LLM) is a neural network trained on very large volumes of text to predict the next word fragment in a sequence. From this simple task comes the ability to answer questions, summarize, and draft text. The model does not consult a database of verified facts and does not "know" when it is wrong. The architecture, training, and evaluation methods are covered in Chapter 13. What matters here is its use in the consultation, by the physician and by the patient.

10.5.2 Performance: From Examinations to Clinical Cases

A meta-analysis of 45 studies from several countries estimated, on medical licensing examinations, an accuracy of 81% for GPT-4 and of 58% for GPT-3.5. GPT-4 passed the examination in 26 of the 29 situations analyzed [395]. Med-PaLM 2 reached 86.5% on the MedQA question set, and the physician raters preferred its answers to those given by other physicians on eight of nine clinical axes [396]. A 2025 review of 39 benchmarks put these results into context. On knowledge tests, leading models reach 84-90%, but on tests that simulate clinical practice performance drops to 45-69%, and on safety evaluation to 40-50% [397].

Studies that placed the model next to the physician give more nuanced results. A 2023 cross-sectional study compared 195 answers given by physicians on a public forum with ChatGPT's answers to the same questions. The evaluators preferred the chatbot's answer in 78.6% of evaluations, rated it as good or very good quality 3.6 times more often, and as empathetic 9.8 times more often [398]. The physicians' answers, however, averaged 52 words and the chatbot's 211, and the setting was a forum, not a medical office.

In a randomized trial with 50 physicians, access to an LLM in addition to the usual resources did not improve diagnostic reasoning: the scores were 76% and 74%, a difference that was not statistically significant. The model used alone scored 16 percentage points higher than the group with conventional resources [399]. By contrast, in a randomized trial with 92 physicians and therapeutic management cases, the GPT-4 group obtained scores 6.5 percentage points higher, but spent on average almost two minutes longer per case [400]. An experimental conversational system had better diagnostic accuracy than 20 primary care physicians in 159 scenarios with simulated patients. The consultation, however, took place through written chat, a format unusual

in practice, and the authors call for caution in interpretation [401]. On 2,400 real cases of acute abdominal pathology, the models tested diagnosed less accurately than physicians, did not follow diagnostic and treatment guidelines, and did not interpret laboratory tests correctly. The results also depended on the order in which they received the information [402].

10.5.3 The Model Used by the Patient

The most important result for practice in 2026 comes from a preregistered randomized study published in *Nature Medicine* [403]. A total of 1,298 adults in the United Kingdom received ten medical scenarios and had to identify the likely condition and choose the appropriate course of action, from "stay at home" to "call an ambulance". They used either one of three LLMs or any source they would normally use. Figure 10.3 summarizes the results.

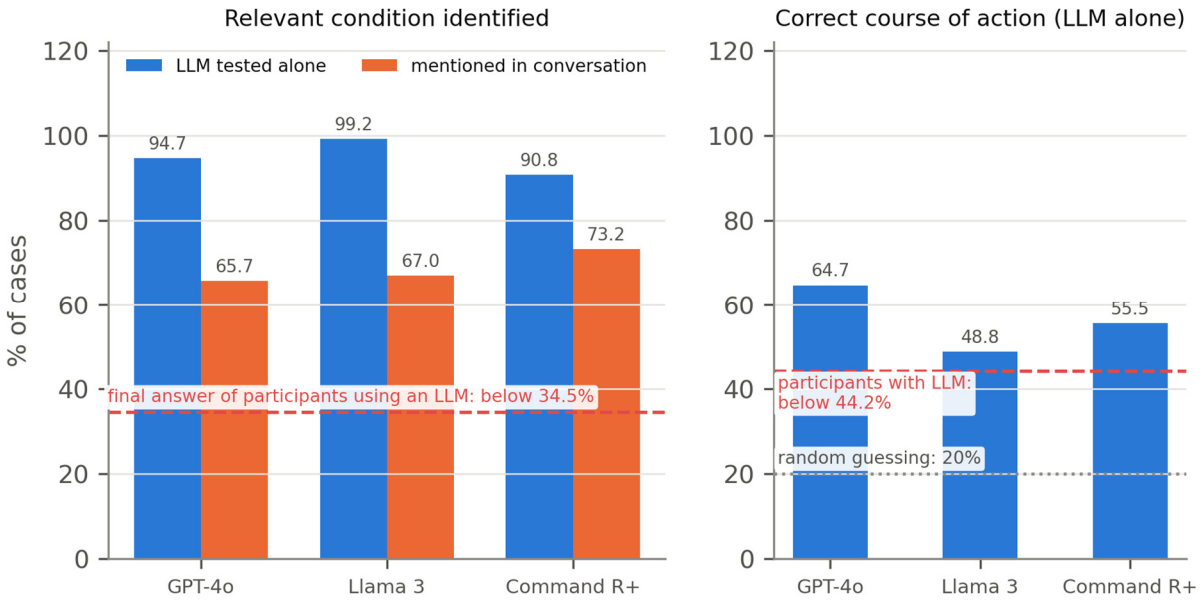


Figure 10.3. Performance of three language models used alone and in conversation with members of the public, in ten medical scenarios. The bars show the percentages reported for each model. The horizontal lines mark the reported upper limits for the final answers of participants who used an LLM (below 34.5% and below 44.2%) and the level of random guessing among five courses of action (20%). Data from [403].

Tested alone, the models identified a relevant condition in 94.9% of cases. The participants who used them identified one in fewer than 34.5% of cases and chose the correct course of action in fewer than 44.2%. The results were no better than in the control group, which even had 1.76 times higher odds of naming a relevant condition [403]. The analysis of the conversations showed where the difference came from. In 16 of the 30 interactions analyzed, the user's first message contained only part of the information, and the model's correct suggestions were not always retained in the final answer [403]. A structured test of a consumer triage tool launched in January 2026 found undertriage in 52% of the reference emergencies: cases of diabetic ketoacidosis or impending respiratory failure were directed to an evaluation within 24-48 hours [404]. An audit of five popular chatbots on topics prone to misinformation found problematic answers in 49.6% of cases. The chatbots refused to answer only 2 times out of 250 [405].

10.5.4 Hallucinations, Sycophancy, and Fabricated References

A hallucination is a statement presented as fact but not supported by the source or by reality. In a framework for evaluating clinical notes generated automatically from consultation transcripts, 12,999 sentences from the notes were checked by clinicians. The hallucination rate was 1.47% of

the note sentences, and the omission rate was 3.45% of the transcript sentences [406]. Of the sentences with hallucinations, 44% contained major errors that could have changed the diagnosis or management if left uncorrected. By type, the hallucinations were fabrications (43%), negations (30%), contextual errors (17%), and causality errors (10%), and fabrications occurred mainly in the treatment plan section [406].

A second vulnerability is sycophancy: the model tends to accept the user's false premise. When asked to write texts based on a false relationship between two equivalent drugs, five leading models complied in up to 100% of cases [407]. The third problem concerns the bibliography. A study of 636 references found that 55% of those generated by GPT-3.5 and 18% of those generated by GPT-4 did not exist [408]. The newer version erred less often, but it did not remove the need for verification in PubMed either.

10.5.5 What a Chatbot Cannot Do

The model works only with the text it receives. It does not see the patient, does not palpate the abdomen, does not notice pallor or respiratory effort, and cannot order a test. The study on acute abdomen cases showed that results deteriorate when information is incomplete or arrives in a different order [402], and the study with members of the public showed that they often leave out information needed for a correct recommendation [403]. Warning signs (red flags) are uncovered through the history taken by the physician and through the clinical examination.

10.5.6 Confidentiality of the Data Entered

Text entered into a public chatbot ends up on the provider's servers. If the text contains data that allow the patient to be identified (name, date of birth, personal numeric code (CNP), address, the combination of a rare disease and a small town), the physician is processing data concerning health, a special category protected by Article 9 of the GDPR [118]. Such processing requires a legal basis and, in the case of an external provider, contractual safeguards. The physician must check whether the tool used provides them, and an account created for private use does not guarantee them. Confidentiality consistently appears among the ethical concerns identified in a systematic review of the use of LLMs in medicine [409]. The practical rule is simple: only fictitious cases or truly anonymized data are entered into public tools, and tools approved by the institution are used according to its procedures.

10.5.7 Guidance, Responsibility, and Disclosure

In January 2024, WHO published guidance on the ethics and governance of large multimodal models, with more than 40 recommendations for governments, technology companies, and healthcare providers [410]. The guidance lists as risks false, inaccurate, or incomplete statements, biased training data, and automation bias, that is, the tendency of professionals and patients to stop noticing the system's errors [410,411]. The systematic review of LLM ethics found the same recurring requirement: human oversight [409].

Responsibility for the clinical decision remains with the physician, regardless of the tool consulted. In scientific writing, the ICMJE recommendations updated in January 2026 require authors to disclose the use of AI tools and to describe how they were used. Writing assistance is mentioned in the acknowledgments section, whereas use for data collection or analysis or for figures is described in the methods, with the tool, the version and, where applicable, the prompts given [202]. Chatbots cannot be authors, because they cannot be accountable for the accuracy and integrity of the work. AI-generated material cannot be cited as a primary source, and references must be verified in a bibliographic source such as PubMed [202].

Table 10.3. Checklist for the responsible use of an LLM by a physician or student. Source: [118,202,406,408,411].

Stage	Checkpoint question
Before	Have I removed any data that could identify the patient?
Before	Is the tool approved by the institution for clinical data?
The question	Have I given the complete context (age, sex, symptoms, duration, medication)?
The answer	Have I checked every statement in a guideline or a primary source?
The answer	Have I looked for omissions, especially warning signs?
The references	Have I found every reference in PubMed and read the abstract?
The decision	Is the decision mine, and can I justify it without the chatbot?
After	Have I disclosed the use (tool, version, purpose)?

10.6 Romania in 2026

Eurostat data show a slower adoption of digital health services. In 2025, 42.7% of people in Romania searched for health information on the internet, compared with 60.3% on average in the EU-27. In 2024, 9.1% made an appointment with a physician through a website, compared with 39.8% in the EU-27 [297]. The trend in online appointments and in access to the personal record is presented in Chapter 8.

On September 1, 2026, the National Health Insurance House (Casa Națională de Asigurări de Sănătate, CNAS) launched the eSănătateaMea portal. Access is through the ROeID digital identity, and insured persons can view their prescriptions from the last 5 years, their health history from the last 10 years, and the providers with whom they can book appointments online [324]. The press release does not yet mention a link with wearable devices or with patients' applications.

Applications and wearable devices that are medical devices follow the same MDR rules in Romania as in the rest of the Union. Incidents involving medical devices are reported to the National Agency for Medicines and Medical Devices of Romania (Agenția Națională a Medicamentului și a Dispozitivelor Medicale din România, ANMDMR), whose dedicated page has a section for incident reporting [412]. The EHDS Regulation applies directly, without a transposition law, including the provisions on the labeling of wellness applications [23].

Common Pitfalls

- Believing that an application is safe because it has many downloads or stars. Quality is assessed with validated instruments, such as MARS.
- Treating the watch's irregular rhythm notification as a diagnosis of atrial fibrillation. It is a signal that requires confirmation on an ECG.
- Reading the SpO2 from the watch as clinical pulse oximetry, without taking into account the errors and skin pigmentation.
- Assuming that every health application is a medical device or, conversely, that none is. The intended purpose stated by the manufacturer decides.
- Confusing the EHDS interoperability label of a wellness application with the CE marking of a medical device.
- Extrapolating from LLM scores on examinations to their safety in the consultation.
- Entering a real case into a public chatbot "with initials only".
- Copying a reference suggested by a chatbot without searching for it in PubMed.

Key Points

- Mobile applications have real but modest effects, the most consistent being on HbA1c and blood pressure.
- Attrition is the rule: mean adherence to applications for chronic diseases was approximately 56%.
- Accuracy varies greatly from sensor to sensor. Pulse and atrial fibrillation detection are well validated; cuffless blood pressure is not.
- The intended purpose stated by the manufacturer decides whether an application is a medical device, and Rule 11 decides the class.
- ISO/IEEE 11073, Bluetooth GHS, and HL7 FHIR form the chain through which a measurement reaches the electronic record.
- An LLM that performs very well on examinations may be of little help to a patient who uses it alone.
- Verify every statement and every reference, protect patient data, and disclose the use of AI.

Self-Assessment Questions

1. (SA) According to the MDCG 2019-11 rev.1 guidance, a fitness application that counts steps and calories, with no declared medical purpose:

a) is a class I medical device
b) is a class IIa medical device, because it monitors a physiological process
c) does not qualify as medical device software
d) becomes a medical device if it has more than one million users
e) is regulated as an accessory of the mobile phone

2. (SA) Under Rule 11 of the MDR, software intended for the continuous surveillance of vital functions in intensive care falls into class:

a) I
b) IIa
c) IIb
d) III
e) it is not a medical device

3. (SA) In the randomized study by Bean et al. (2026), participants from the general population who used an LLM:

a) identified the relevant condition in more than 90% of cases
b) performed better than the control group in choosing the course of action
c) identified the relevant condition in fewer than 34.5% of cases, with no advantage over the control group
d) obtained the same results as the models tested alone
e) were excluded from the final analysis

4. (MA) Which of the statements about wearable devices are supported by the data in the chapter?

a) in a 2025 meta-analysis, smartwatches had a pooled sensitivity of 95% for atrial fibrillation
b) only about 11% of the wearable devices on the market have been validated for at least one measurement
c) blood pressure measured without a cuff is recommended for the diagnosis of hypertension
d) pulse oximeters exceeded the 3% error threshold in all skin pigmentation groups
e) continuous glucose monitoring did not change HbA1c in randomized trials

5. (MA) According to the ICMJE recommendations updated in January 2026:

a) the use of AI for writing assistance is disclosed in the acknowledgments section
b) a chatbot can be a coauthor if it wrote a substantial part of the text
c) the use of AI for data analysis is described in the methods section
d) authors are responsible for material produced with the help of AI
e) AI-generated material can be cited as a primary source

CHAPTER 11

Chapter 11. The Information System, Information Security, and the Electronic Signature

Objectives

By the end of this chapter, the student will be able to:

- distinguish the information system of a healthcare facility from the computer-based information system that supports it and list their components
- describe the strategic, tactical, and operational levels of information system management and the stages in the life cycle of a computer-based information system
- estimate the level of a security risk with a probability x impact matrix, starting from the asset, the threat, and the vulnerability
- explain the difference between symmetric encryption, asymmetric encryption, and the hash function, and follow the steps by which a digital signature is created and verified
- classify an electronic signature as simple, advanced, or qualified under the eIDAS Regulation and recognize its applications in the Romanian healthcare system
- identify the standardization bodies in medical informatics and check the current status of a standard on iso.org

11.1 The Information System and the Computer-Based Information System

11.1.1 Two Notions That Are Often Confused

In everyday language, "the hospital's IT system" means the software running on its computers. Medical informatics, however, distinguishes two notions. A system is a set of persons, things, events, and the relationships between them that form a whole, and a system made up of human and technical components is called socio-technical [413]. The information system of a healthcare facility comprises all data, information, and knowledge processing in that facility, together with the human or technical actors who carry it out [413]. The computer-supported part is the computer-based information system. The rest remains non-computerized: the telephone, the printed form, the discussion at the morning report after the on-call shift [413].

A hospital has an information system from its first day of operation, even without a single computer. The question is how well it works and how it can be improved [413].

11.1.2 The Components

The information system can be viewed through four components (Figure 11.1).

- **People**: physicians, nurses, patients, relatives, administrative staff, IT specialists. They appear in the system in the role of information processors, following established rules [413].

- **Processes**: functions describe what is done (admission, discharge), with no defined beginning or end, whereas work processes describe how it is done, as a sequence of activities with their conditions [413].
- **Data**: the types of entities about which data are kept, such as "patient", "episode of care", or "medical letter" [413].
- **Technology**: application systems (a software product installed and configured for a particular facility) and physical processing systems, from servers and tablets to paper forms and telephones [413].

The purpose of the whole can be summed up as information logistics: the right information, at the right time, in the right place, for the right people, and in the right form, so that they can make correct decisions [413].

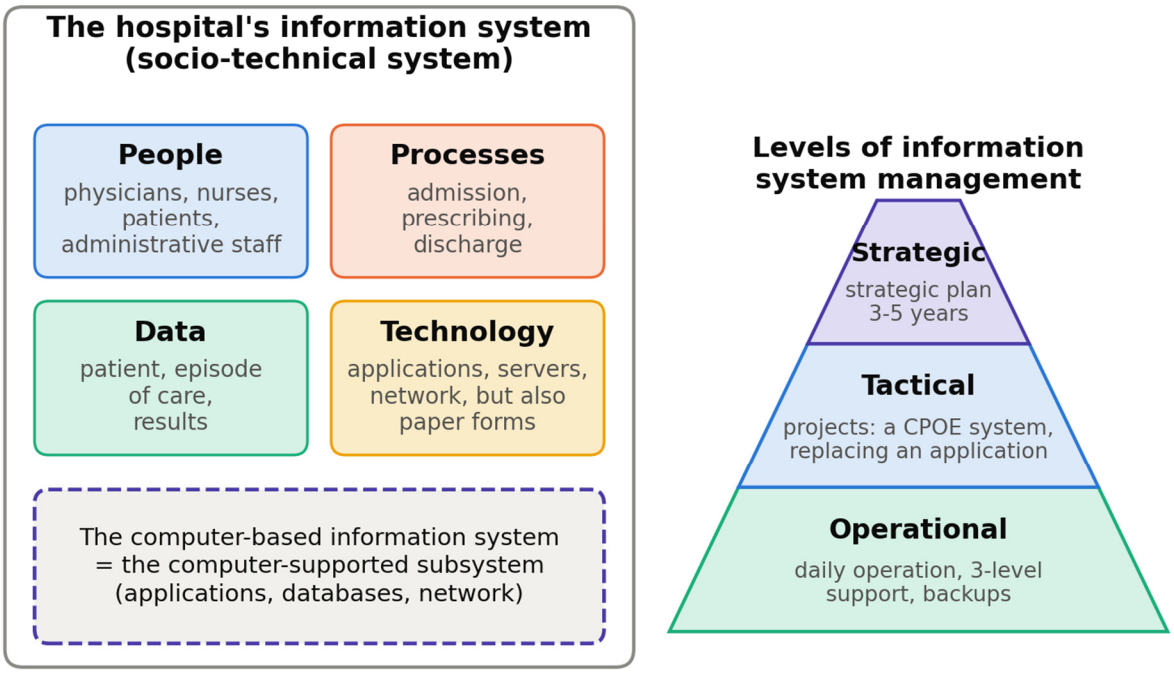


Figure 11.1. Left: the components of a hospital's information system, in which the computer-based information system is only the computer-supported subsystem. Right: the levels of information system management, with examples. Diagram based on [413,414].

11.1.3 The Levels: Strategic, Tactical, Operational

Information system management means planning the information system and its architecture, directing its construction and operation on the basis of these plans, and monitoring compliance with them [413]. It is divided into three levels [414].

Strategic management looks at the system as a whole and produces the strategic information management plan. This plan describes the goals of the facility, the current state, the desired state, and the steps for moving from one to the other. It is usually valid for 3-5 years, with an annual project portfolio that is updated more often [414]. Tactical management deals with specific components that are introduced, changed, or retired, almost always through projects. Typical examples are the introduction of a computerized provider order entry (CPOE) system or the evaluation of user satisfaction in an anesthesia and intensive care (ATI) unit [414]. Operational management ensures day-to-day operation. It includes the user help desk, three-level support (the first level receives the incident, the second is provided by the hospital's IT staff, and the third, for example, by the application vendor's specialists), and backups [414].

11.1.4 The Life Cycle of a Computer-Based Information System

A clinical information system is not something one "buys" once. For teaching purposes, its life cycle can be divided into six stages.

1. **Planning**: the project starts from the strategic plan, with its own plan that specifies the problem, objectives, tasks, and resources [414].
2. **Requirements analysis**: it is established what the system must do (functional requirements, for example "prescribes and checks interactions") and how it must do it (non-functional requirements: availability, response time, security). For a genetic diagnostic laboratory system, the analysis combined a literature review, user interviews, and a market analysis, a process also required by the European regulation on in vitro diagnostic medical devices and by the IEC 62304 standard [415].
3. **Selection or development** of the software product and its customization for the facility.
4. **Implementation**: data migration, training, go-live. At the end, the project hands the system over to operational management through a signed handover record [414].
5. **Operation and maintenance**: incidents, problems, and technical changes are managed according to the ITIL library, which recommends a committee that coordinates and monitors changes, because any modification can have side effects [414].
6. **Evaluation and retirement**: evaluation shows when the system needs to be replaced, but patient data must sometimes be kept for up to 30 years, longer than any application lasts, so archiving is entrusted to a dedicated system [131].

11.1.5 Change Management

The term has two meanings. In the technical sense (ITIL), change management controls the modifications made to the components of the system [414]. In the organizational sense, it concerns the people who have to work differently once the system is introduced. A review of 90 studies on the implementation of the electronic health record (EHR) grouped the factors into compliance processes, collaboration processes, competence development processes, and costs [416]. A synthesis of 41 qualitative or mixed-methods studies showed that a clear presentation of the purpose and early involvement of users prepare the implementation, whereas tailored training and on-demand technical support sustain use after go-live [417].

Romania in 2026

- The health insurance IT platform (Platforma informatică a asigurărilor de sănătate, PIAS) brings together the single integrated information system (Sistemul informatic unic integrat, SIUI), the electronic prescription system (Sistemul informatic de prescripție electronică, SIPE), the health insurance card, and the national electronic health record (Dosarul electronic de sănătate, DES) [322].
- The project to modernize PIAS, funded under the National Recovery and Resilience Plan (PNRR), was completed on August 31, 2026. Among its results, the National Health Insurance House (Casa Națională de Asigurări de Sănătate, CNAS) lists the migration of the existing platform, the eSănătateaMea portal, and electronic medical forms [323].
- Since September 1, 2026, the electronic medical forms component has been in use (referral slips, medical letters, sick leave certificates). For the forms to be visible in the eSănătateaMea portal, providers' applications must be adapted, and CNAS has advised providers to update them according to the SIUI interface specifications [418].

11.2 Information Security Management

11.2.1 The Information Security Management System

Data protection, discussed from a legal and ethical perspective in Chapter 5, requires internal organization. The ISO/IEC 27001:2022 standard (3rd edition, with the 2024 amendment) sets out the requirements for an information security management system (ISMS), which an organization establishes, implements, maintains, and continually improves, with a risk management process suited to its size [419]. The concrete measures (the controls) are described in ISO/IEC 27002:2022.

For healthcare, ISO 27799:2025, the 3rd edition published in December 2025, adapts these controls to health organizations. Its scope includes, in addition to ordinary IT equipment, EHR systems, medical devices with software, and even infection control or building management equipment [189]. The draft standard (ISO/DIS 27799:2025) keeps the ISO/IEC 27002 structure of four themes: organizational controls, people controls, physical controls, and technological controls [190]. What is specific to healthcare comes from medical devices, from the balance between clinical safety and security, from a highly diverse user population, and from services that run 24 hours a day [190]. Responsibility for security lies with information system management, and the facility's leadership must provide the financial resources [131].

11.2.2 Risk Assessment

Controls are not picked from a catalog. They are chosen after a risk assessment with a clearly defined scope [190]. The ISO/IEC 27005:2022 guidance describes this process for information security [420]. The basic vocabulary consists of five terms:

- **the asset** is what must be protected: the database with patient records, the PACS server, the laptop of the physician on call, but also the hospital's reputation.
- **the threat** is the event that can cause harm (a phishing e-mail, ransomware, a flood).
- **the vulnerability** is the weakness through which the threat acts (a reused password, an operating system without updates, the lack of backups).
- **the impact** is the consequence for patients and for operations.
- **the probability** is the likelihood that the threat will exploit the vulnerability within a given period.

Risk combines probability with impact. In the teaching version in Figure 11.2, each is scored from 1 to 5, and their product places the risk in a color zone. Decisions on the identified risks are made according to the risk acceptance criteria and the treatment options set by the organization [190]. In practice, risks in the critical zone require immediate measures (treatment), whereas those in the low zone may be knowingly accepted. A review written for non-specialists found that the protection strategies used by US hospitals are often weak [421].

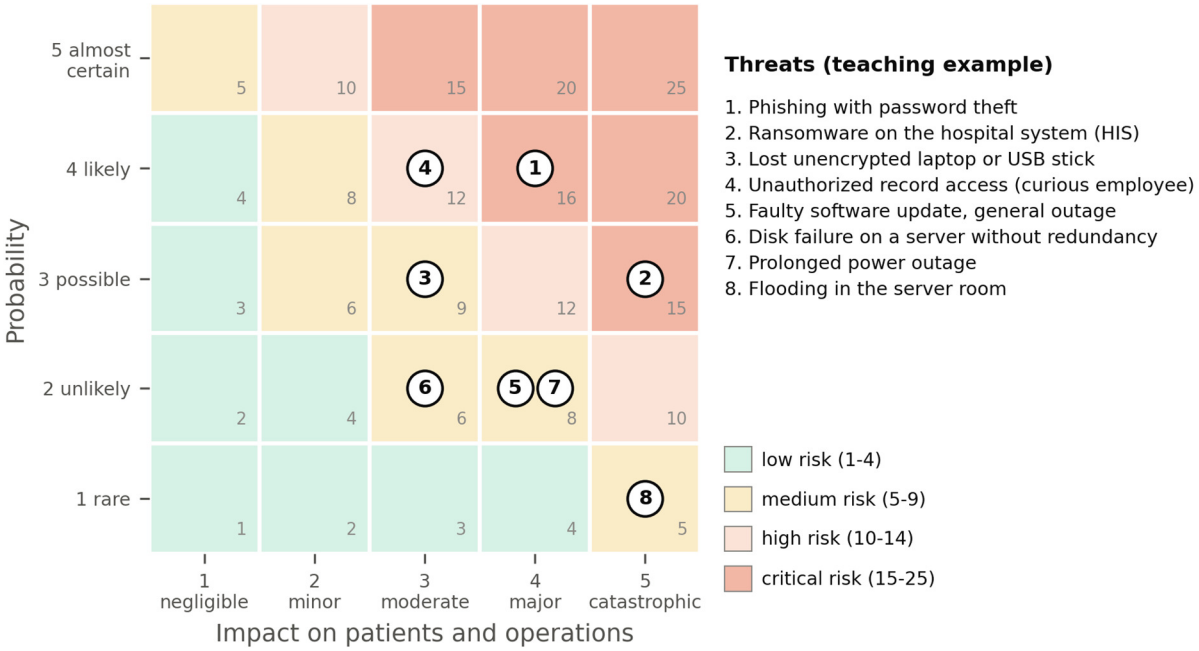


Figure 11.2. Probability x impact risk matrix for eight threats typical of a hospital. The positions are illustrative, chosen for teaching purposes, and do not come from a real assessment. Each hospital sets its own scores.

11.2.3 Policies, Information Classification, and Access Control

The security policy is the document approved by management that states what is permitted and who is accountable. It relies on information classification. The draft ISO 27799 recommends that personal health information be classified as at least "confidential" and allows for different levels: a fractured leg does not have the same sensitivity as a sexually transmitted disease, a psychiatric episode from ten years ago differs from a current one, and genetic data also concern the patient's relatives [190].

Role-based access control (RBAC) grants rights to the role (ward physician, nurse, pharmacist, registration clerk), not to the person. A user should access a patient's data only when part of the team caring for that patient or when acting on the patient's behalf, and emergencies require an authorized procedure for bypassing the usual rules, called "break the glass" [190]. Overly strict control can make it impossible, in an emergency, for a physician from another ward to prescribe a medication. For this reason, after a risk analysis, some restrictions may be relaxed for medical staff, with increased individual accountability [414].

11.2.4 Audit Logs

Every access to the electronic record leaves a trace in the audit log: who, what, when, and from which workstation. A systematic review of 85 studies showed that these logs, originally designed for access control, are now also used to study clinical activity [422]. In malpractice litigation, the audit log can show who viewed a result and when [423], which is why the log itself must be protected [190]. The European logging component required of EHR systems by the EHDS Regulation is presented in Chapter 5.

11.2.5 Business Continuity and Recovery

A hospital must keep working even when its computer-based information system is down. A review identified 43 downtime events in 166 US hospitals in 2012-2018, totaling 701 days, almost half of them (48.8%) related to cyberattacks [424]. Not every outage is an attack. On July 19, 2024, a faulty update to a security software product was followed by a loss of response to internet scanning in 759 of 2,232 US hospitals analyzed (34.0%) [425].

The continuity plan identifies vital processes and fallback procedures (paper forms, spreadsheets). The draft ISO 27799 requires the plans to be tested in stages, from tabletop exercises to full-scale rehearsals [190]. The recovery plan relies on backups that are themselves protected against attacks and on data centers with at least two redundant locations [131]. Incident reporting under the NIS2 Directive and Government Emergency Ordinance (OUG) No. 155/2024, as well as the 2024 ransomware attack on Romanian hospitals, are analyzed in Chapter 5.

11.3 Cryptography Basics for Physicians

11.3.1 Symmetric Encryption

Encryption transforms plaintext into ciphertext, which can be read back only with the right key. In symmetric encryption, the sender and the recipient use the same secret key. The current standard is AES (Advanced Encryption Standard), which processes data in 128-bit blocks, with keys of 128, 192, or 256 bits [426]. Symmetric encryption requires little computation and is suited to large volumes: a laptop's disk, the PACS archive, backups. Its problem is key distribution. How does the key reach the recipient without being intercepted?

11.3.2 Asymmetric Encryption

Public key cryptography uses a pair of mathematically related keys. The public key can be known by anyone, while the private key stays with its holder alone. Whatever is encrypted with the recipient's public key can be decrypted only with the recipient's private key. Conversely, only the holder of the private key can generate a digital signature, and anyone can verify it with the public key [427]. The US Digital Signature Standard (FIPS 186-5) approves three families of algorithms (RSA, ECDSA, and EdDSA) and requires a modulus of at least 2048 bits for RSA [427]. In practice, encryption is hybrid: the document is encrypted symmetrically, and only the short symmetric key is encrypted asymmetrically.

11.3.3 The Hash Function

A hash function computes, from data of any length, a value of fixed length [428]. The same input always gives the same digest, and changing a single comma produces a completely different digest. A hash function is not encryption: it has no key, and the document cannot be reconstructed from the digest. Its role is to verify integrity: when signing, what is usually signed is the digest, the condensed version of the document, not the whole document [427].

11.3.4 Digital Certificates and the Public Key Infrastructure

How does the pharmacist know that a public key belongs to Dr. X? Through the digital certificate, an electronic document by which a certification authority (CA) binds the public key to the identity of its holder. Authorities are organized hierarchically, from the root CA to intermediate CAs, and the whole forms the public key infrastructure (PKI). For healthcare, ISO 17090-1:2021 describes identity certificates, attribute certificates (for example, for the professional role), and CA hierarchies [429]. ISO 17090-4:2026, published in January 2026, sets out the minimum requirements and formats for generating and verifying digital signatures on healthcare documents [430]. The draft version of this part describes the European formats CAdES, XAdES, PAdES, and JAdES and their levels, including the basic signature, the signature with a time stamp, and the archival signature, which allows verification many years later. The physician's private key is kept on a smart card, a USB token, or a software token [428].

11.3.5 TLS, HTTPS, and End-to-End Encryption

When a portal's address begins with https, the connection is protected by the TLS (Transport Layer Security) protocol. TLS allows applications to communicate over the internet in a way designed to prevent eavesdropping, tampering, and message forgery. Version 1.3 was republished in July 2026 as RFC 9846, which obsoletes RFC 8446 [431]. The server presents its certificate, the browser verifies the certification chain, and then a symmetric session key is established. TLS protects data only in transit.

In end-to-end encryption, the message is decrypted only on the recipient's device, and the intermediate server cannot read it. No technique helps, however, if the data leave through personal e-mail, a habit also documented in Chapter 5.

The RSA and ECDSA algorithms are not expected to withstand an attack by a large-scale quantum computer [427]. In August 2024, NIST published the post-quantum standard FIPS 203, which specifies the ML-KEM mechanism for key establishment, currently believed to be secure even against an adversary with a quantum computer [432].

Example

A teaching example with small numbers (real RSA uses numbers of more than 600 decimal digits). The physician has the public key (n = 55, e = 3) and the private key d = 27.

- The fictitious digest of the medical letter is h = 9.
- Signing: s = h^d mod n = 9^27 mod 55 = 4. The physician sends the letter and the signature s = 4.
- Verification by the recipient: s^e mod n = 4^3 mod 55 = 64 mod 55 = 9. The recipient recomputes the digest of the received letter, obtains 9, and the values match: the letter is unmodified and was signed with the matching private key.
- If someone changes the dose in the letter, the recomputed digest becomes, for example, 10. Verification still gives 9, and the signature is rejected.

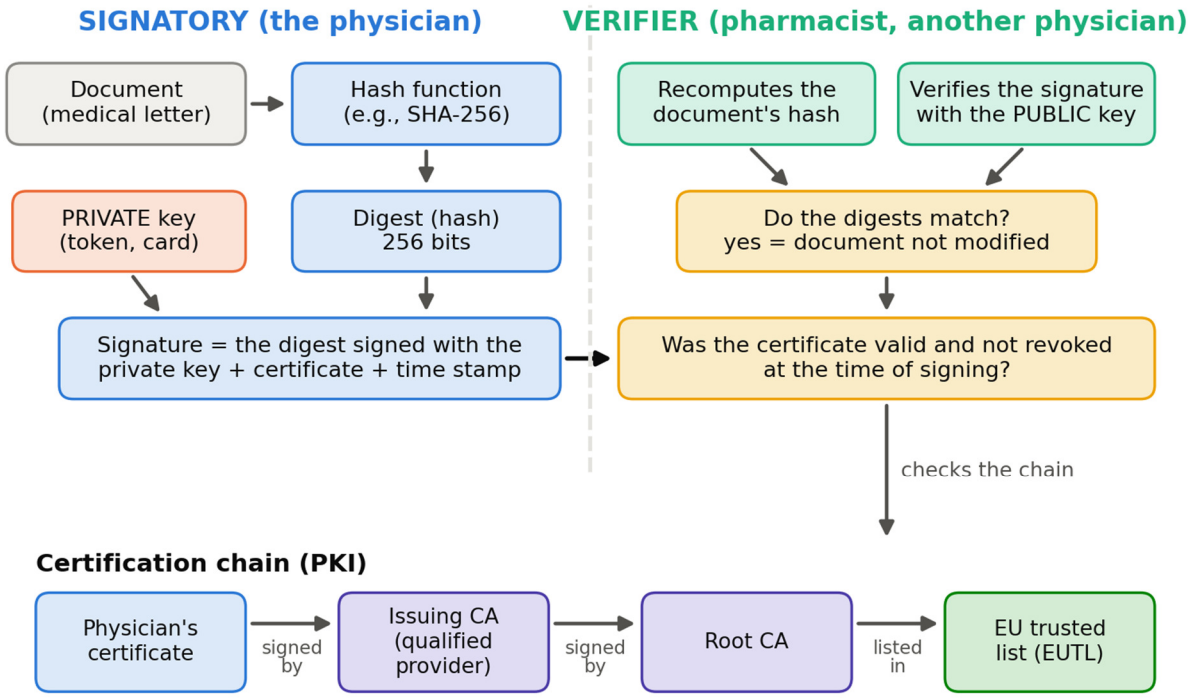


Figure 11.3. Creating and verifying a digital signature: signing the digest (hash) with the private key, verifying with the public key from the certificate, and following the certification chain up to the EU trusted list. Diagram based on [427,428,433].

11.4 The Electronic Signature

11.4.1 Types of Signature under eIDAS

Regulation (EU) No 910/2014 (eIDAS) defines three levels (Table 11.1). An electronic signature (simple) means data in electronic form which is attached to or logically associated with other data in electronic form and which is used by the signatory to sign [433]. An advanced electronic signature must meet four requirements: it is uniquely linked to the signatory, it is capable of identifying the signatory, it is created using electronic signature creation data that the signatory can, with a high level of confidence, use under his sole control, and it is linked to the data signed therewith in such a way that any subsequent change in the data is detectable [433]. A qualified electronic signature is an advanced electronic signature that is created by a qualified electronic signature creation device and is based on a qualified certificate for electronic signatures [433].

The legal effects differ. An electronic signature cannot be denied legal effect solely on the grounds that it is in electronic form or that it is not qualified, but only a qualified electronic signature has the equivalent legal effect of a handwritten signature [433]. Following the 2024 amendment, the recognition of qualified electronic signatures in all Member States is provided for in the new Article 24a [434]. Legal persons (a hospital, a health insurance house) use the electronic seal, which ensures the origin and integrity of the data [433].

Table 11.1. Types of electronic signature under the eIDAS Regulation. Source: [433,434].

Type	Main requirements	Legal effect	Examples
Simple	electronic data used by the signatory to sign	cannot be denied legal effect solely because it is electronic	typed name, scanned image of a signature
Advanced	uniquely linked to the signatory, identifies the signatory, under the signatory's sole control, detects changes	as for the simple signature, cannot be denied legal effect solely because it is electronic	signature with a non-qualified certificate
Qualified	advanced + qualified certificate + qualified signature creation device	equivalent to a handwritten signature, recognized in all EU Member States	physician's signature with a token or card from a qualified trust service provider

11.4.2 The Time Stamp and Validation

An electronic time stamp binds a set of data to a particular point in time and proves that the data existed at that time. A qualified electronic time stamp enjoys the presumption of the accuracy of the date and time it indicates and of the integrity of the data, and its time source is linked to Coordinated Universal Time (UTC) [433]. It matters because a certificate may expire or be revoked, and the signature remains valid only if it can be proved that it was created before that happened.

The validation of a qualified electronic signature confirms, among other things, that the certificate was a qualified certificate, valid at the time of signing and issued by a qualified trust service provider, and that the integrity of the signed data has not been compromised [433]. Each Member State publishes a trusted list of the qualified trust service providers it supervises [433], and the European Commission offers an online tool for browsing these lists (Practical Session 11).

In the long term, qualified electronic signatures can guarantee the integrity of archived documents only if the archiving system can renew the signatures and replace outdated hash algorithms [131].

11.4.3 Applications in the Romanian Healthcare System

Since February 1, 2014, medicines with and without patient co-payment in outpatient treatment have been prescribed on the electronic prescription form, online or offline, with one part completed by the physician and one by the pharmacist (Order of the Ministry of Health and CNAS No. 674/252/2012). When the physician has an extended electronic signature (a term from the legislation in force before 2024), the document bears a statement that it was generated and signed electronically. The physician signs by hand only if he or she has no extended electronic signature, cannot use it, or is working offline [435]. In the version consolidated on July 29, 2026, this statement still refers to Law No. 455/2001 [435], although that law was repealed by Law No. 214/2024 on the use of electronic signatures, time stamps, and the provision of trust services, in force since October 2024. The new law designates the Romanian Authority for Digitalization (Autoritatea pentru Digitalizarea României, ADR) as the supervisory body for trust services [436].

Since September 1, 2026, the CNAS electronic medical forms component has been in use. Its forms include referral slips to specialists or for hospital admission, referral slips for paraclinical investigations, the medical letter, recommendations for home care, and the sick leave certificate. Patients can see these forms in the eSănătateaMea portal if the provider's application has been adapted [418]. In the portal, patients authenticate through ROeID [324]. Once archived, the discharge letter remains verifiable in the long term if the archiving system uses qualified electronic signatures and renews them in time [131].

11.4.4 eIDAS 2 and the European Digital Identity Wallet

Regulation (EU) 2024/1183 amended eIDAS and entered into force on May 20, 2024. Each Member State must provide at least one European Digital Identity Wallet, an application in which citizens keep their person identification data and electronic attestations of attributes, and the source code of the application's components must be licensed under an open-source license [434]. The European Commission states that the wallets will be available by the end of 2026 [437].

11.5 Standardization in Medical Informatics

11.5.1 Why Standards

Interoperability standards reduce the cost of interfaces between applications, whereas conversion (mapping) between formats usually loses information or meaning [131]. HL7 version 2 is the most widely implemented communication standard in hospitals, DICOM covers imaging, and IHE (Integrating the Healthcare Enterprise) is not a standardization body. IHE selects and constrains existing standards for specific clinical workflows, through profiles tested at Connectathon events [131]. The current version of HL7 FHIR is R5 (5.0.0, March 2023), and R6 had not been published as of September 27, 2026 [307]. DICOM also exists as an ISO standard, ISO 12052:2026 [438]. Chapter 8 covers these standards in detail.

11.5.2 ISO/TC 215 and CEN/TC 251

The ISO/TC 215 *Health informatics* technical committee was established in 1998, and its secretariat is held by ANSI (US). On September 27, 2026, its web page listed 251 published ISO standards (212 under the direct responsibility of the committee, figures that include revisions) and

63 under development, with 39 participating members and 32 observing members [439]. It has two subcommittees (genomics; pharmacy and medicines) and working groups such as those for architecture (WG 1), systems and device interoperability (WG 2), semantic content (WG 3), security, safety, and privacy (WG 4), and telehealth (WG 12), plus a joint working group with ISO/IEC JTC 1/SC 42 on artificial intelligence [439].

ISO open data allow a view over time (Figure 11.4). From 2001 to September 2, 2026, ISO/TC 215 and its subcommittees published 392 documents (excluding amendments). On September 27, 2026, 246 were in force (118 International Standards, 82 Technical Specifications, 45 Technical Reports, and one Publicly Available Specification), and 146 had been withdrawn or replaced [438]. In each year of the 2021-2025 period, the committee published 22-25 documents.

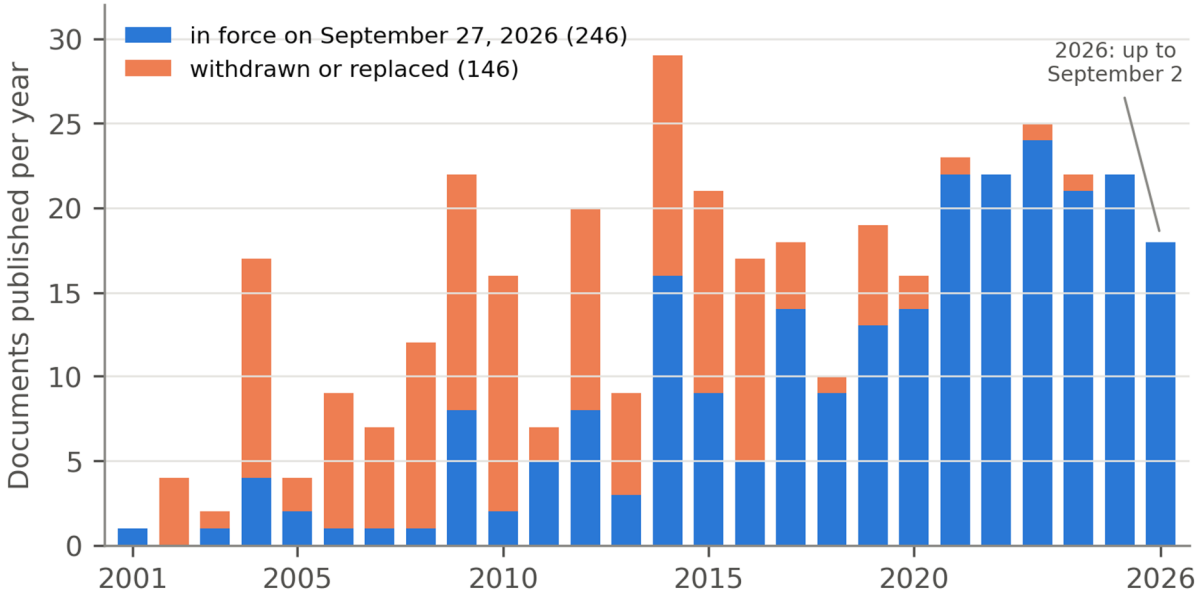


Figure 11.4. Documents published by ISO/TC 215 and its subcommittees, by year of publication, according to their status on September 27, 2026. The year 2026 is incomplete. Own calculation on the iso_deliverables_metadata data set, downloaded on September 27, 2026 (ODC-By 1.0 license), excluding amendments and corrigenda [438].

Its European counterpart is CEN/TC 251 *Health informatics*. Under the Vienna Agreement, wherever possible, a single standard is approved for each subject, adopted by both ISO and CEN, and CEN/TC 251 is a member of the Joint Initiative Council for Global Health Informatics Standardization (ten members in 2022) [440]. The International Patient Summary standard started as an initiative led by CEN/TC 251, with funding from the European Commission, and is among the CEN/TC 251 standards adopted by ISO/TC 215 [440]. The current edition is ISO 27269:2025 [438].

11.5.3 The National Body

In Romania, European and international standards are adopted through the Romanian Standards Association (Asociația de Standardizare din România, ASRO). Medical informatics standards are managed by technical committee CT 319 "Informatică medicală" (Medical Informatics), which has published, for example, SR EN ISO/IEEE 11073-10103:2025 on the nomenclature for implantable cardiac devices [441]. The author of this book works as an expert in this technical committee.

11.5.4 Standards Change: Examples

Table 11.2 brings together several standards related to the topics of this chapter. The "status" column shows why a standard should always be checked on iso.org before it is cited. The ISO 21549 series (patient healthcard data) had new editions in 2023 and 2024, which were withdrawn in April 2025. Conversely, the drafts ISO/DIS 17090-4, ISO/DIS 22532, and ISO/DIS 16791 became published standards in 2026 [438].

Table 11.2. Medical informatics standards related to security, signatures, and identification. Status checked on iso.org on September 27, 2026. Source: [189,419,430,438].

Standard	Subject	Status on September 27, 2026
ISO/IEC 27001:2022	information security management systems: requirements	in force (2024 amendment)
ISO 27799:2025	information security controls in health	in force, 3rd edition
ISO 17090-4:2026	digital signatures for healthcare documents	in force, 3rd edition
ISO 12052:2026	DICOM, including workflow and data management	in force, 3rd edition
ISO 17523:2025	requirements for electronic prescriptions	in force, 2nd edition
ISO 27269:2025	international patient summary	in force, 2nd edition
ISO 11239:2023	identification of medicinal products: pharmaceutical dose forms, routes of administration, packaging	in force, 2nd edition
ISO 22532:2026	identification of medicinal products (IDMP): core vocabulary	in force, 1st edition
ISO 16791:2026	machine-readable coding of medicinal product package identifiers	in force, 3rd edition
IEC/TS 81001-2-2:2025	communication of security needs, risks, and controls for health software	in force, 1st edition
ISO/TR 11636:2009	dynamic on-demand virtual private network for health information infrastructure	in force
ISO 21549-5:2023 and 21549-7:2024	identification data and medication data on the patient healthcard	withdrawn in 2025

11.6 The Human Factor

11.6.1 Phishing

Phishing is a message that imitates a trusted source in order to obtain a password, a click, or the opening of an attachment. A systematic review of 70 articles on the human factor in the security of healthcare organizations found a steady growth in the literature on phishing and pointed to the need for digital hygiene among medical staff on social networks [442].

Studies that measure actual behavior are more instructive. In a hospital in Italy with more than 6,000 employees, three simulated campaigns showed that a personalized message tied to the local context is far more effective: in the first campaign, 55% of recipients clicked the personalized link, compared with 7% for the standard message (Figure 11.5). The authors also describe the obstacles: the reaction of the trade unions, the difficulty of coordination, and the ethical problems of fake messages [443]. Another study, discussed in Chapter 5, associated clicking with workload rather than with stated intention [181]. In a Norwegian hospital, the click rate was 53% in the

neutral group (396 of 753), and a brief cognitive dissonance intervention did not produce a statistically significant difference (P = 0.22) [444].

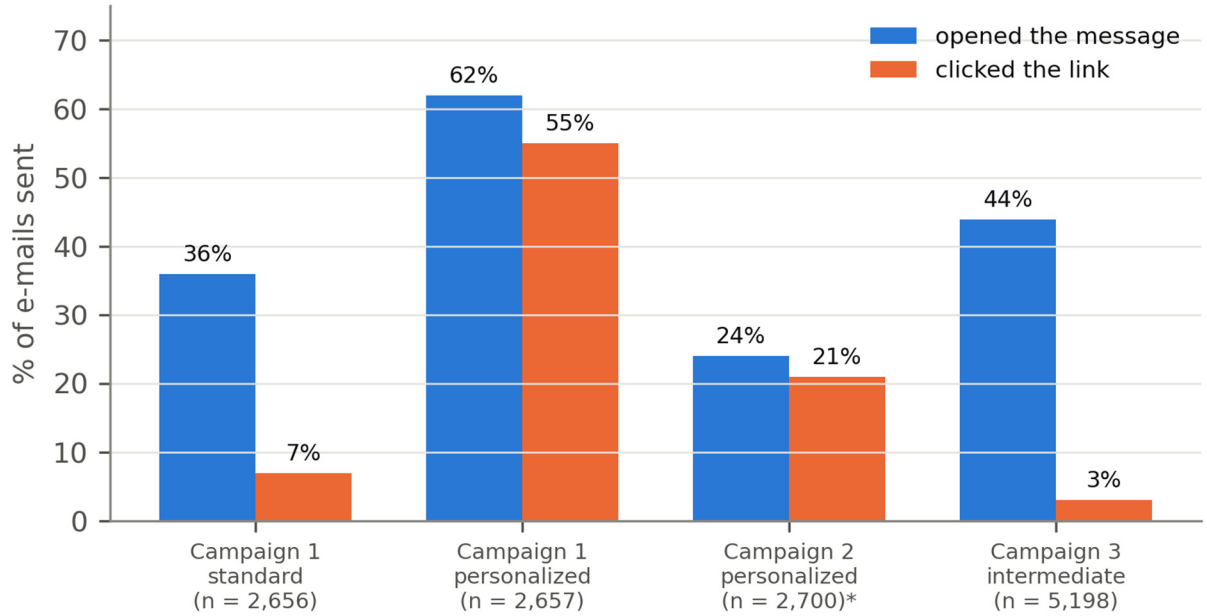


Figure 11.5. Results of three simulated phishing campaigns in an Italian hospital with more than 6,000 employees: the percentage of messages opened and of links clicked, relative to the e-mails sent. *In Campaign 2, 1,600 messages ended up in the spam folder. Data from [443].

11.6.2 Passwords and Multi-Factor Authentication

The NIST SP 800-63B-4 guideline (2025) changed several rules learned by generations of users. A password used as a single factor must be at least 15 characters long, and one used in multi-factor authentication at least 8. Mandatory combinations of character types are no longer imposed, nor are periodic password changes, but a password must be changed when there is evidence of compromise, and it is checked against a list of compromised or commonly used values [445]. Applications assessed at authenticator assurance level AAL2 must also offer a phishing-resistant option, and a code sent by SMS remains a restricted authenticator [445].

Multi-factor authentication (MFA) combines something you know (the password), something you have (the phone, the token), and something you are (the fingerprint). Patients access eSănătateaMea through their ROeID account, with their e-mail address, password, and the code generated by the authenticator app [324].

11.6.3 Digital Hygiene in the Hospital

The 2025 ENISA guide for the health sector groups the measures into eight domains. They include MFA for systems accessible online, the principle of least privilege, at least one ransomware-resistant or offline backup, the clean desk and clear screen policy, and phishing simulations [446].

For students, a few rules apply from the first day of clinical training. An open session is closed before the user walks away. A password is not lent to a colleague on call, because the audit log will attribute the actions to the account holder. Results are not photographed with a personal phone. A suspicious e-mail is reported, not silently deleted. Cybersecurity incidents can be reported to the Romanian National Cyber Security Directorate (Directoratul Național de Securitate Cibernetică, DNSC) through the PNRISC platform or by calling 1911 [447].

Common Pitfalls

- Believing that change management means only installing a new version. Without training and user involvement, a good system can fail.
- Regarding hashing as a form of encryption. A hash function has no key and cannot be reversed, so it hides nothing; it only verifies integrity.
- Believing that a scanned signature pasted into a PDF is a qualified signature. It is a simple electronic signature.
- Sending one's private key or token PIN to a colleague "so that they can sign the prescriptions". The signature is attributed to the key holder.
- Reading the padlock in the browser as a guarantee that the website is legitimate. HTTPS protects the connection; it does not verify the honesty of whoever obtained the certificate.
- Citing a standard from a draft (DIS) or from a withdrawn edition without checking it on iso.org.

Key Points

- The computer-based information system is the computerized subsystem of the information system, which also includes people, processes, and paper-based tools.
- Strategic management produces the 3-5-year plan, tactical management runs projects, and operational management ensures day-to-day operation.
- Risk is estimated as probability x impact for each asset-threat-vulnerability combination.
- Symmetric encryption uses a single key, asymmetric encryption uses a key pair, and a hash function produces a fixed-length digest.
- The digest is signed with the private key, and verification is done with the public key from the certificate.
- Only a qualified electronic signature has a legal effect equivalent to that of a handwritten signature.
- On September 27, 2026, ISO/TC 215 had 246 documents in force, and the status of a standard is checked on iso.org.

Self-Assessment Questions

1. (SA) Which of the following is NOT a requirement for an advanced electronic signature under Article 26 of the eIDAS Regulation?

a) it is uniquely linked to the signatory
b) it is capable of identifying the signatory
c) it is created using data under the sole control of the signatory
d) it allows any subsequent change in the signed data to be detected
e) it is created by a qualified signature creation device

2. (SA) A physician's signing certificate was revoked on March 10. Which element makes it possible to prove that a medical letter signed by this physician is valid, if it was signed on January 15?

a) the hash function of the document
b) the time stamp applied at signing
c) the recipient's public key
d) encryption of the document with AES
e) the certificate of the HTTPS server

3. (SA) In the RSA teaching example with the public key (n = 55, e = 3), the recipient receives the signature s = 4 and recomputes the digest of the received document: h = 10. Which conclusion is correct?

a) the signature is valid, because 4^3 mod 55 = 9 is smaller than 10
b) the signature does not match the received document, because 4^3 mod 55 = 9 differs from 10
c) verification requires the physician's private key
d) the digest must first be decrypted with the symmetric key
e) the result cannot be interpreted without the time stamp

4. (MA) Regarding the levels of management of a hospital's information system, which statements are correct?

a) strategic management produces a plan that is usually valid for 3-5 years
b) the introduction of a CPOE system is a typical tactical management project
c) the help desk and three-level support belong to operational management
d) tactical management deals exclusively with hardware equipment
e) the strategic plan, once approved, is no longer updated

5. (MA) According to the NIST SP 800-63B-4 guideline, which statements are correct?

a) a password used as a single factor is at least 15 characters long
b) users must be required to change their password every 90 days
c) composition rules such as "one uppercase letter, one digit, one symbol" are not imposed
d) a newly chosen password is checked against a list of compromised or commonly used values
e) a code received by SMS is the most phishing-resistant of all authenticators

CHAPTER 12

Chapter 12. Hospital Information Systems and the Single Integrated Information System of Romanian Health Insurance

Objectives

By the end of this chapter, the student will be able to:

- list the modules of a hospital information system and state which data each one produces and which data it consumes;
- compare the monolithic architecture with the modular architecture integrated through an integration engine, calculating the number of interfaces required in each case;
- interpret the evidence published after 2020 on electronic prescribing, clinical decision support systems, and alert fatigue;
- describe the path of a patient's data from the family physician to the health insurance house, including the role of SIUI, SIPE, DES, and the health insurance card;
- calculate the average length of stay, the bed occupancy rate, the case-mix index, and the readmission rate, and then interpret them on a dashboard;
- explain what downtime operation means for a hospital when its computer-based information system becomes unavailable.

12.1 The Hospital's Information System: Components and Modules

12.1.1 Information System and Computer-Based Information System

Every hospital has an information system from its first day of operation. Winter and colleagues define the information system of a healthcare facility as the socio-technical subsystem that comprises all data, information, and knowledge processing, together with the human and technical actors who carry it out [448]. The temperature chart on the bedside table, the on-call register, and the telephone between the ward and the laboratory are all part of it. The part that uses computers, which Winter calls the computer-based information subsystem, is what everyday language calls the hospital information system (HIS). The right question is therefore not whether the hospital "has a system", but how its functioning can be improved, for example through systematic management or through modern tools [448].

12.1.2 The Modules

An HIS consists of specialized applications (application components), each supporting one hospital function. Table 12.1 summarizes the most common ones.

The patient administration system, usually called the ADT (admission, discharge, transfer) module, records the patient's identity, admission, transfers between wards, and discharge. It supplies correct and up-to-date administrative data to all the other components, which is why

Winter describes it as the center of the hospital's administrative memory [448]. An identification error made at admission then propagates to the laboratory, to imaging, and to the invoice.

The electronic general clinical observation chart (foaia de observație clinică generală, FOCG), the Romanian inpatient medical record, brings together the history, the physical examination, the progress notes, the discharge summary, and the medical letter. Nursing documentation is kept in the same place. Computerized provider order entry (CPOE) for medications and investigations sends orders to the pharmacy, the laboratory, or radiology and can check doses, interactions, and allergies [448]. The laboratory information system (LIS) receives requests, tracks specimens, takes the results from the analyzers, and validates them. The radiology information system (RIS) schedules examinations and manages reports, and the image archive (PACS) stores DICOM images (Chapter 7).

The operating theater has its own module for scheduling rooms and for the operative report. Intensive care uses a patient data management system (PDMS), which automatically acquires parameters from the bedside equipment and calculates severity scores such as TISS or SAPS II [448]. The administrative modules (finance and accounting, human resources, procurement, and pharmacy with stock management) make up the enterprise resource planning (ERP) system. The data warehouse periodically extracts data from all of these for reports and dashboards.

Table 12.1. The main modules of a hospital information system. Source: adapted from [448].

Module	Main functions	Typical data	Usual data exchange
ADT (patient administration)	identification, admission, transfer, discharge	CNP, demographic data, ward, bed, referral diagnosis	HL7 v2 ADT, FHIR Patient/Encounter
Electronic FOCG and nursing	history, progress notes, discharge summary, care plan	clinical notes, ICD-coded diagnoses	structured clinical documents (Chapter 8)
CPOE with decision support	prescriptions, requests for investigations, order sets	drug, dose, route, frequency	order and result messages
Pharmacy	validation, dispensing, stock management	consumption per patient and per ward	integration with CPOE and ERP
LIS	requests, specimen collection, results, validation	numerical values, units, reference intervals	HL7 v2, LOINC codes
RIS/PACS	scheduling, images, reports	DICOM studies, text reports	DICOM, HL7 (Chapter 7)
Operating theater and ICU (PDMS)	room scheduling, anesthesia record, monitoring, scores	vital signs every few seconds or minutes, infusions, fluid balance	medical device gateway
Finance, DRG coding, ERP	reimbursement, costs, payroll, procurement	diagnosis and procedure codes, consumption	reporting to the national platforms

12.1.3 Architectures: Monolithic or Modular

The way the modules are connected to one another defines the architecture. If all applications write to a single central database, the architecture is of the DB1 type. If each application has its own database, it is of the DBn type, and the same data (the patient, the admission episode) exist in several copies that must be kept consistent [448]. In practice, almost all large hospitals have a mixed architecture. A second criterion is the number of applications: a monolithic architecture uses a single component for all functions, whereas a modular one uses several. The third criterion

is the number of vendors. The all-in-one solution, with everything from a single vendor, is the opposite of the best-of-breed combination, the best product for each function, which usually has multiple databases [448].

In a modular architecture, the applications must exchange messages. If every pair of applications has a direct link, the number of interfaces for n applications is n(n-1)/2: 10 applications require 45 interfaces, each of which has to be maintained and tested at every update. The literature calls this a "spaghetti" architecture. The classic solution is a communication server, also called an integration engine, placed at the center of a star architecture (Figure 12.1). Each application communicates only with the server, so only n links are needed [448]. The server receives messages asynchronously, keeps them in a queue if the recipient is down, transforms them from one format into another, and can distribute them to several recipients at the same time. A typical example: the ADT module sends an HL7 admission message, and the server forwards it to all the applications that will need the patient's data (for example, CPOE and radiology), each of which updates its local copy of the patient's data [448]. Messaging standards (HL7 v2, FHIR) and the electronic health record are covered in Chapter 8.

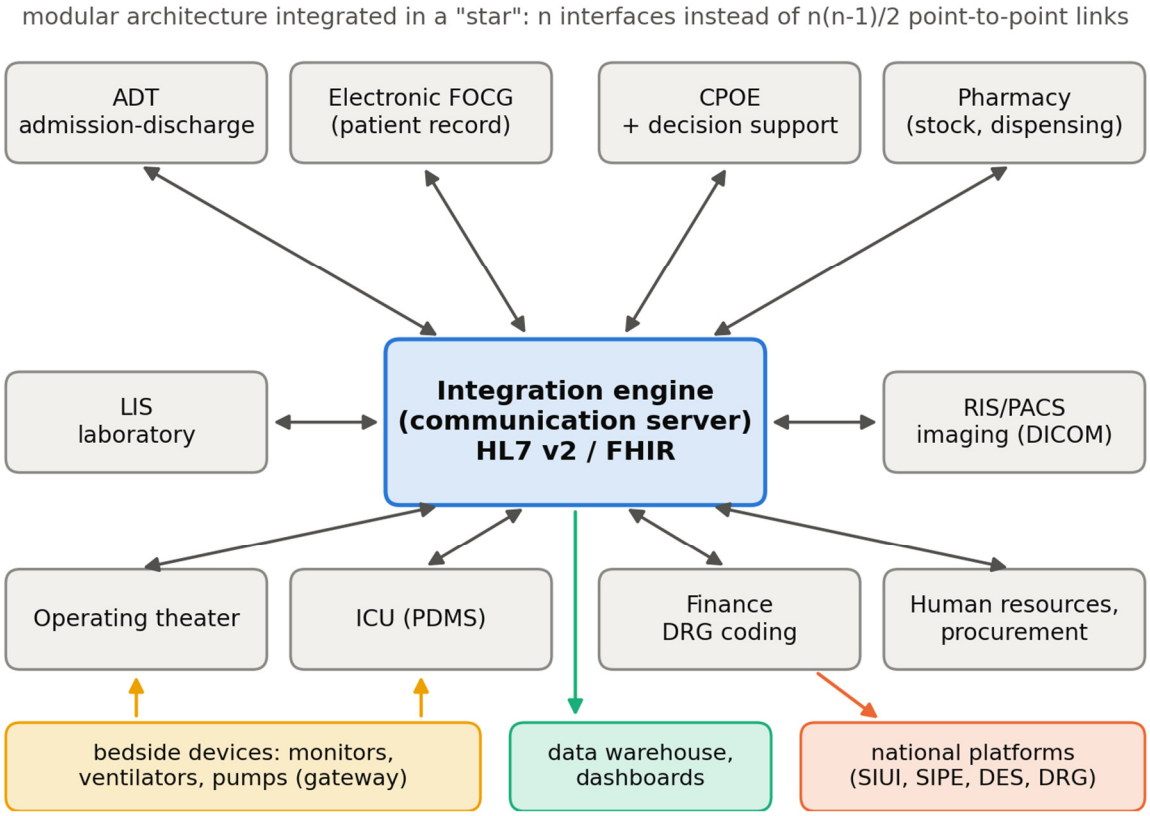


Figure 12.1. Modular architecture of a hospital information system integrated through an integration engine (communication server). Bedside equipment connects through a gateway to the operating theater and ICU modules. Aggregated data flow to the data warehouse and to the national platforms. Teaching diagram, adapted from [448].

The choice of architecture is not a purely technical matter. In an analysis of more than 1,600 German hospitals from the DigitalRadar project, greater use of a single vendor's modules and a higher integration ratio were associated with higher digital maturity scores, on a scale from 0 to 100 [449]. The authors recommend strategic HIS procurement, targeted system improvements, and incentives from public policy.

12.2 Electronic Prescribing in the Hospital and Clinical Decision Support

12.2.1 What CPOE Brings

On paper, a prescription passes through several hands: the physician writes it, the nurse transcribes it onto the treatment chart, and the pharmacy reads it and dispenses the drug. Every transcription is an opportunity for error. CPOE eliminates transcription, enforces mandatory fields (dose, route, frequency), and allows automatic checks. Clinical decision support (CDS) systems add dose calculation, checks for interactions, allergies, and the compatibility of the drug with laboratory results, as well as predefined order sets for common situations [448].

The evidence from 2020-2026 confirms the benefit, with some nuances. A systematic review of 35 hospital studies compared prescribing with health information technology against paper orders. The technology was associated with lower prescribing error rates, with a median odds ratio of 0.24 (interquartile range 0.03-0.57) [450]. The same review, however, found errors generated by the technology itself in 31 of the 35 studies, related to data entry, duplicate prescriptions, and alert fatigue. For chemotherapy, where a dosing error can be fatal, a meta-analysis of eight studies obtained a relative risk of 0.19 (95% CI 0.08-0.44) after the introduction of CPOE, with very high heterogeneity ($I^2 = 99\%$) and with one study in which errors increased [451]. A meta-analysis of 45 studies on CDS for prescribing reported a positive but small overall effect (standardized mean difference 0.084, 95% CI 0.067-0.102), and in some subgroups the difference from usual practice was not statistically significant [452]. Figure 12.2 places these values side by side.

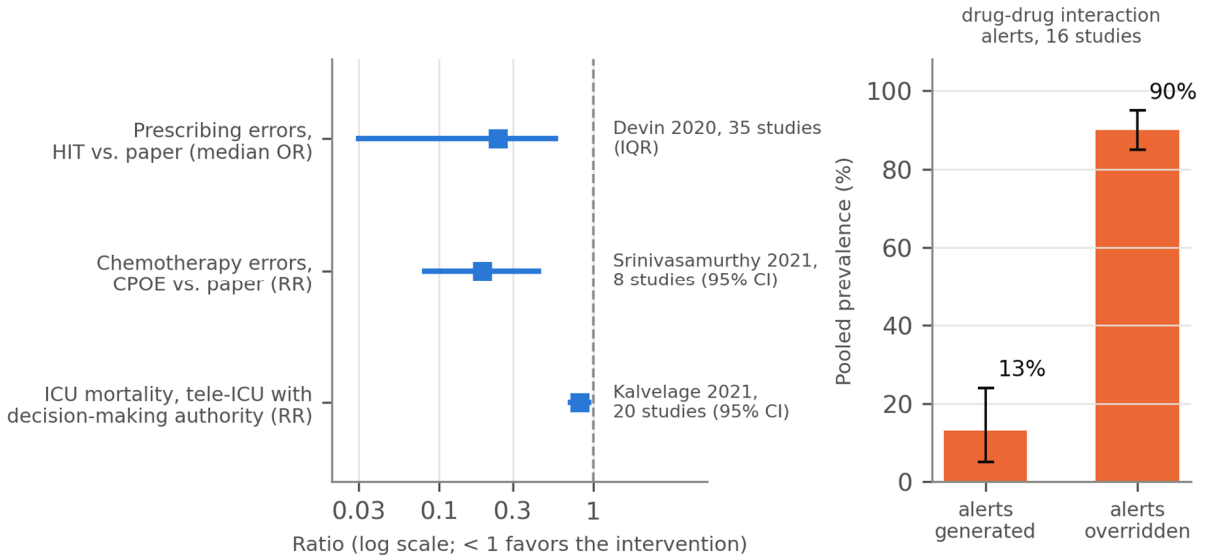


Figure 12.2. Published estimates for information technology in prescribing and in intensive care. Left: prescribing errors in hospitals (median OR, IQR, 35 studies) [450]; medication errors in chemotherapy after CPOE (RR, 95% CI, 8 studies) [451]; ICU mortality with a tele-ICU team that has decision-making authority (RR, 95% CI, 20 studies) [453]. Right: pooled prevalence of drug-drug interaction alerts generated and of those overridden by physicians (95% CI, 16 studies) [454]. Values taken exactly from the cited abstracts.

12.2.2 Alert Fatigue

A useful alert should appear rarely, at the right moment, and call for a clear action. Reality is different. A meta-analysis of 16 studies on drug-drug interaction alerts estimated an overall prevalence of generated alerts of 13% (95% CI 5-24%), and physicians override them in 90% of cases (95% CI 85-95%) [454]. A high override rate does not automatically mean negligence. Many alerts are clinically irrelevant, and a physician who sees dozens a day learns to close them without reading them. This phenomenon is called alert fatigue.

Its measurement is surprisingly imprecise. A review of 22 systematic reviews found a single operational definition of alert fatigue, and the most frequently used indicators were the number of alerts, the override rate, and the acceptance rate [331]. The authors propose measuring alert fatigue as a sustained decrease, relative to a baseline, in the rate of appropriate responses to alerts. On the remediation side, all the hospitals that published how they optimize their alerts used a multidisciplinary committee, and five of six evaluations reported a decrease in the number of alerts after the changes [455]. The lesson for the future physician: an alert is configured, measured, and withdrawn if it does not do its job.

12.3 Case Study: Computerizing Intensive Care

12.3.1 Why the ICU

A mechanically ventilated patient generates thousands of values a day: heart rate, invasive blood pressure, oxygen saturation, ventilator parameters, the flow rate of every infusion pump. On paper, the nurse copies some of them once an hour. A PDMS acquires them automatically from the equipment, displays warnings, calculates scores, and, at transfer, generates a summary for the ward that takes over the patient [448]. This abundance of signals also has a cost. False or clinically irrelevant alarms lead to desensitization of the staff, and a systematic review of 69 publications found evidence that informatics solutions can reduce alarm fatigue, with most of the work targeting the elimination of technically false alarms [456]. In a meta-analysis of 6,908 ICU nurses, alarm fatigue was at a moderate level, and higher levels were positively correlated with the tendency toward medical errors and with burnout [457].

12.3.2 A Romanian Set of Tender Specifications

The Ministry of Health published the tender specifications for the project "Sistem informatic pentru evidența clinică a secțiilor ATI" (Information System for the Clinical Records of ICU Departments, SIEC-ATI), which started in May 2021 and was funded by the Competitiveness Operational Programme 2014-2020 (SMIS code 131065) [458]. The document is useful for teaching because it shows how the requirements for a clinical system are formulated before any vendor is chosen.

Case Study

- Scope: 18 hospitals and institutes in Bucharest, Cluj-Napoca, Iași, Târgu Mureș, and Timișoara, with 649 ICU beds and 275 anesthesia workstations in operating rooms, plus a central component at the ministry [244].
- Device integration: automatic recording, at intervals of no more than 90 seconds, of all parameters from the monitor, the ventilator, volumetric pumps, syringe pumps, and the hemofiltration machine; data collection must not depend on the device model, so that replacing a piece of equipment does not affect the system [244].
- Medication safety: allergy alerts with configurable severity levels, prescriptions that require a double signature, and unusual doses flagged separately and possibly blocked [244].
- Scores and coding: automatic calculation of the Glasgow, APACHE II, SAPS II, and TISS scores; ICD-10 codes and export of the DRG data to the hospital system, together with ventilation time and the APACHE score [244].
- Interoperability: XML/JSON, HL7/FHIR, and web service interfaces with the HIS of each hospital [244].
- National level: a dashboard at the ministry showing free and occupied ICU beds in each hospital, updated within no more than 60 seconds, with the option of reserving a bed [244].
- Continuity: the ability to create a temporary patient record if the hospital's HIS is unavailable [244].

It is worth noting how the requirements are written. They are verifiable ("maximum 90 seconds", "maximum 60 seconds"), vendor-neutral (every brand mentioned carries the note "or equivalent"), and they separate clinical functions from infrastructure functions [244]. Vague tender specifications ("a modern, easy-to-use system") cannot be verified at acceptance.

12.3.3 What the Studies Show

A retrospective before-and-after study carried out by a team from Timișoara in a trauma ICU compared 43 patients from the period of paper documentation (2021-2022) with 65 patients from 2025, after a PDMS had been introduced. The median ICU length of stay decreased from 13 to 6 days ($p = 0.02$), but in the adjusted analyses the differences in mortality, infections, and length of stay were not statistically significant [459]. The result illustrates a general rule: a small before-and-after study suggests, but does not demonstrate, a causal effect. Also in Timișoara, an abstract on the digitalization of ICU services at the "Pius Brânzeu" County Emergency Clinical Hospital discussed the benefits of advanced monitoring for patient safety, together with the risks related to system reliability, cybersecurity, and staff training [460]. Staff preparedness is not uniform. In a survey of 161 professionals from an ICU in Romania, nurses stated more often than physicians that they could not assess the usefulness of digital tools, and training simulators had been used by 1.8% of nurses, compared with 47.2% of physicians [461].

The international literature adds three further directions. The first is telemedicine: a meta-analysis of 20 studies with 477,637 patients found lower ICU mortality when the tele-ICU team had decision-making authority (RR 0.82, 95% CI 0.71-0.94), but not when it only provided consultation (Figure 12.2) [453]. The second is automated checking of guideline adherence: in five university hospitals with different information systems, a system that applied digitally encoded recommendations to standardized data achieved an accuracy of 97.0%, compared with 86.6% for human assessors [462]. The third is data quality. A review of 29 studies in intensive care reported that medication errors related to the electronic record accounted for 34% of medication errors in the ICU, and copied-and-pasted text appeared in up to 82% of residents'

progress notes [463]. An overview of ICU digitalization also discusses the potential of machine learning models to outperform classic scores in predicting mortality and length of stay [464].

12.4 The Romanian Health Insurance Information System

12.4.1 The Components of the Platform

The hospital does not work in isolation. The services it provides to insured persons are reported, validated, and reimbursed through a national platform. Law No. 95/2006, as amended by Law No. 45/2019, provides that the National Health Insurance House (Casa Națională de Asigurări de Sănătate, CNAS) organizes and administers the health insurance IT platform (Platforma informatică din asigurările de sănătate, PIAS). It comprises the single integrated information system (Sistemul informatic unic integrat, SIUI), the national social health insurance card system, the national electronic prescription system (Sistemul informatic de prescripție electronică, SIPE), and the electronic health record system (Dosarul electronic de sănătate, DES) [321,322].

Each component has a distinct role. In SIUI, providers of medical and pharmaceutical services electronically report the services they have provided. SIPE introduced the electronic prescribing of reimbursed and free prescriptions. For the electronic health insurance card (cardul electronic de asigurări de sănătate, CEAS), CNAS publishes software libraries that are also compatible with the electronic identity card (cartea electronică de identitate, CEI) [322]. DES brings together, at the national level, electronic records with clinical, biological, diagnostic, and therapeutic data accumulated throughout a person's life. The record is created when the patient's first medical document is transmitted, and it has five modules: "Emergency Summary", "Medical History", "Patient-Reported History", "Medical Documents", and "Personal Data" [321]. Physicians see the data only with the patient's consent. The exception is the emergency summary, which is accessible without consent to physicians in emergency services and in primary care, strictly for the medical act, and the patient can see the complete history of accesses [321]. The structure of the record and its standards are discussed in Chapter 8. Table 12.2 summarizes the role of each component and its status on September 27, 2026.

Table 12.2. The components of the health insurance information system: role and status on September 27, 2026. Source: [321-325,465].

System	Who uses it	Role	Status on September 27, 2026
SIUI	all providers under contract with the health insurance houses	reporting of services provided, basis for reimbursement	operational, migrated to the modernized PIAS infrastructure
SIPE	prescribing physicians, pharmacies	reimbursed and free e-prescriptions	operational
Insurance card (CEAS) and CEI	patient, providers	identification of the insured person when a service is provided	the CEI has not become mandatory as of September 1, 2026; the national card remains in use
DES	physicians, patient	consolidated medical history, emergency summary	separately, the eSănătateaMea portal (September 1, 2026) shows the insured person a 10-year health history
DRG reporting (INMSS)	hospitals	minimum data set per patient, RO DRG grouping	RO DRG v1 grouper online, DRGNational application 1.40 (March 31, 2026)

12.4.2 The Path of a Patient's Data

Let us follow a fictitious patient (Figure 12.3). The family physician examines the patient, finds a cardiac condition, and issues a referral slip. The outpatient cardiologist performs the consultation and an echocardiogram, then reports the services in SIUI. The patient is admitted to the hospital. At discharge, the hospital issues the medical letter and the prescription, and the patient-level minimum data set reaches the National Institute for Health Services Management (Institutul Național de Management al Serviciilor de Sănătate, INMSS). The institute collects the data from all hospitals, for both inpatient and day care, and provides the RO DRG grouper [465]. Diagnoses are coded in ICD-10-AM, third edition, and the classification in force is RO DRG v.1.1, from March 2024 [139]. The details of grouping are in Chapter 4. The e-prescription goes through SIPE, the pharmacy dispenses it, and the county health insurance house (CAS) reimburses the services on the basis of the reports.

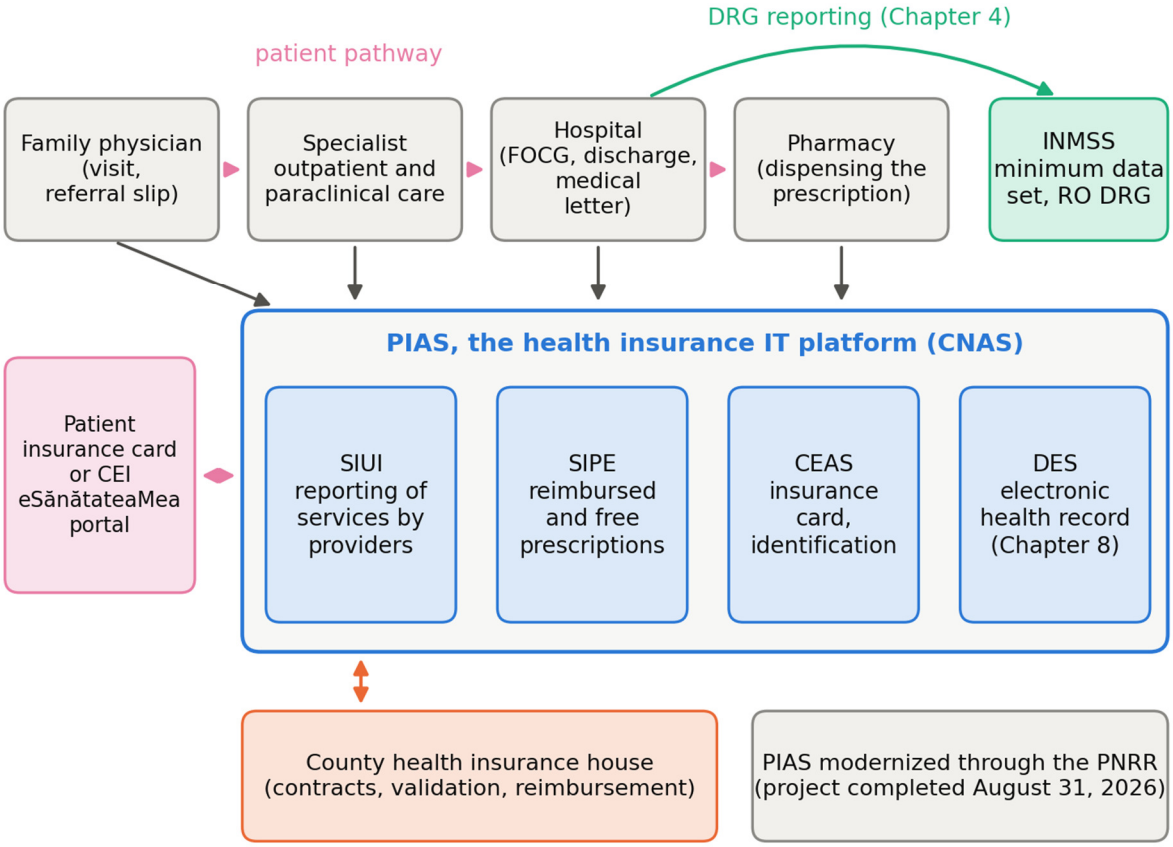


Figure 12.3. Data flow between providers and the health insurance IT platform (PIAS): the patient pathway, reporting to SIUI, prescribing through SIPE, identification with the insurance card or the CEI, the electronic health record, and DRG reporting to INMSS. Teaching diagram based on [321-323,465].

The diagram shows why administrative errors in the hospital have effects outside it. An incorrect diagnosis code in the FOCG changes the DRG grouping, and therefore the funding. An incorrect personal numeric code (CNP) at admission means that the service cannot be validated. For the resident physician, correct coding is part of the medical act.

12.4.3 The Modernized Platform, a Moving Target

In October 2025, CNAS announced that PIAS would be replaced "in 2026" with a modern information system funded through the National Recovery and Resilience Plan (Planul Național de Redresare și Reziliență, PNRR) [466]. The final press release of August 31, 2026, however,

uses different wording. The project "Redimensionarea, standardizarea și optimizarea Platformei Informatice din Asigurările de Sănătate - PIAS" (Resizing, Standardization, and Optimization of the Health Insurance IT Platform - PIAS), funded under PNRR Component C7, Investment I3.1, and carried out between March 26, 2024, and August 31, 2026, built a new technology infrastructure and migrated the existing platform onto it, together with its data and services [323]. New modules were added: the e-Sănătatea Mea portal, electronic medical forms, the Register of Insured Persons, and international relations. The approved budget was 493,900 thousand lei excluding VAT, of which 434,679 thousand lei were paid [323]. The document speaks of a "modernized PIAS", not of a new system with a different name that would replace SIUI.

Romania in 2026

- On September 1, 2026, CNAS launched the eSănătateaMea portal, with authentication through ROeID. Insured persons can see their insurance status, their prescriptions from the last 5 years, their health history from the last 10 years, reimbursements, and the status of their card, and they can book appointments online [324].
- As of the same date, the use of the CEI did not become mandatory for providers, and insured persons' access to services is not conditional on the CEI [325].

12.5 The Digital Transformation of Hospitals and Digital Maturity

12.5.1 PNRR Investments

The digitalization of healthcare in the PNRR is funded under Component 7, "Digital Transformation", through Investment I3, dedicated to the e-health and telemedicine system, and not under Component 12, which mainly funds buildings and equipment [368]. In addition to the modernization of PIAS (I3.1), Investment I3.2 allocated 496,380,000 lei excluding VAT for the digitalization of 13 institutions subordinated to the Ministry of Health, such as the public health directorates and the ambulance services [467]. Investment I3.3 allocated 493,900,000 lei excluding VAT for information systems and digital infrastructure in 200 public healthcare facilities, with a call opened on July 7, 2023, and closed, after an extension, on September 15, 2023 [326]. The European deadline for meeting the milestones and targets was August 31, 2026 [468]. As of September 27, 2026, no official summary of the results achieved by the hospitals funded under I3.3 had been identified.

12.5.2 How Digital Maturity Is Measured

A hospital that has bought computers is not, by that fact alone, a digital hospital. Maturity models describe the stages an organization goes through. The best known, the Electronic Medical Record Adoption Model (EMRAM) of the HIMSS organization, has eight stages, from 0 (no departmental systems) to 7 (a dynamic record and advanced data analytics), passing through ancillary systems, a clinical data repository, electronic documentation, electronic ordering with governance, data integration, and advanced data exchange [469]. Among 1,026 US hospitals, those at stages 6-7 had 3.25 times higher odds of obtaining a better safety grade than those at stage 0 (95% CI 2.33-4.55) [470]. The study is cross-sectional, so it shows an association, not a causal relationship.

EMRAM is not the only instrument. A systematic review identified 27 digital maturity models for hospitals and synthesized them into seven dimensions: strategy, technological capability,

interoperability, governance, patient-centered care, people and skills, and data analytics [471]. Germany offers a useful lesson for Romania. The national DigitalRadar program assessed about 90% of German hospitals [449], and the association between the process digitalization score and quality indicators was largely not statistically significant or difficult to interpret [472]. A panel of 49 German hospitals showed that, by 2025, clear progress had appeared only in medication management and discharge, while many funded projects were still under way [473]. Funding alone does not guarantee a change in the way people work.

12.6 Indicators and Dashboards

12.6.1 Activity Indicators

Data from the ADT and coding modules make it possible to calculate several standard indicators every month. The average length of stay (ALOS) is the ratio between the total number of inpatient days and the number of discharges [474]. In 2023, in most EU Member States, the ALOS for inpatients ranged between 5.2 and 9.4 days [474]. The bed occupancy rate is calculated as the inpatient days delivered divided by the available bed-days (beds x calendar days), multiplied by 100. In 2024, Romania had 731 hospital beds per 100,000 inhabitants, the third highest level in the EU, compared with a European average of 507 [475]. A large number of beds with low occupancy indicates inefficiently used capacity.

The case-mix index (CMI) is the mean of the relative weights of the cases discharged in a given period. A CMI of 1.25 shows that the hospital treated cases with an expected resource consumption 25% above the average reference case (the DRG system is explained in Chapter 4). In a database of 715 US hospitals, the mean CMI rose from 1.70 at the end of 2019 to 1.79 in the first quarter of 2024 [476].

The 30-day readmission rate is the proportion of discharges followed by a new admission within 30 days. It seems simple, but its interpretation requires care. In a neurosurgery department with 8,878 patients, the rate was 2.9%, and the authors concluded that the indicator has no uniform definition, depends on the health system, and requires risk adjustment [477]. A planned readmission (for example, for the second cycle of chemotherapy) is not a failure.

12.6.2 The Dashboard

A dashboard displays the key indicators on a single screen, with thresholds and signaling. Figure 12.4 shows an example built with fictitious data. A systematic review of 54 articles listed the functional requirements of hospital dashboards (reporting, alerts, customization, tracking of performance indicators) and the non-functional requirements (speed, security, integration with other systems, a data warehouse) [478]. As for their effects, a review of 70 studies found that, of the 43 outcomes reported for length of stay, 28 indicated a reduction, whereas mortality remained unchanged in 15 of 20 outcomes [479]. Activity data collected in the hospital also serve to evaluate a clinical service. For example, a retrospective analysis of 867 patients operated on for colon cancer at the "Pius Brânzeu" County Emergency Clinical Hospital in Timișoara, between February 2019 and February 2023, showed fewer elective operations, more emergency operations, and a shorter average length of stay in the years of the COVID-19 pandemic than in the preceding year [115].

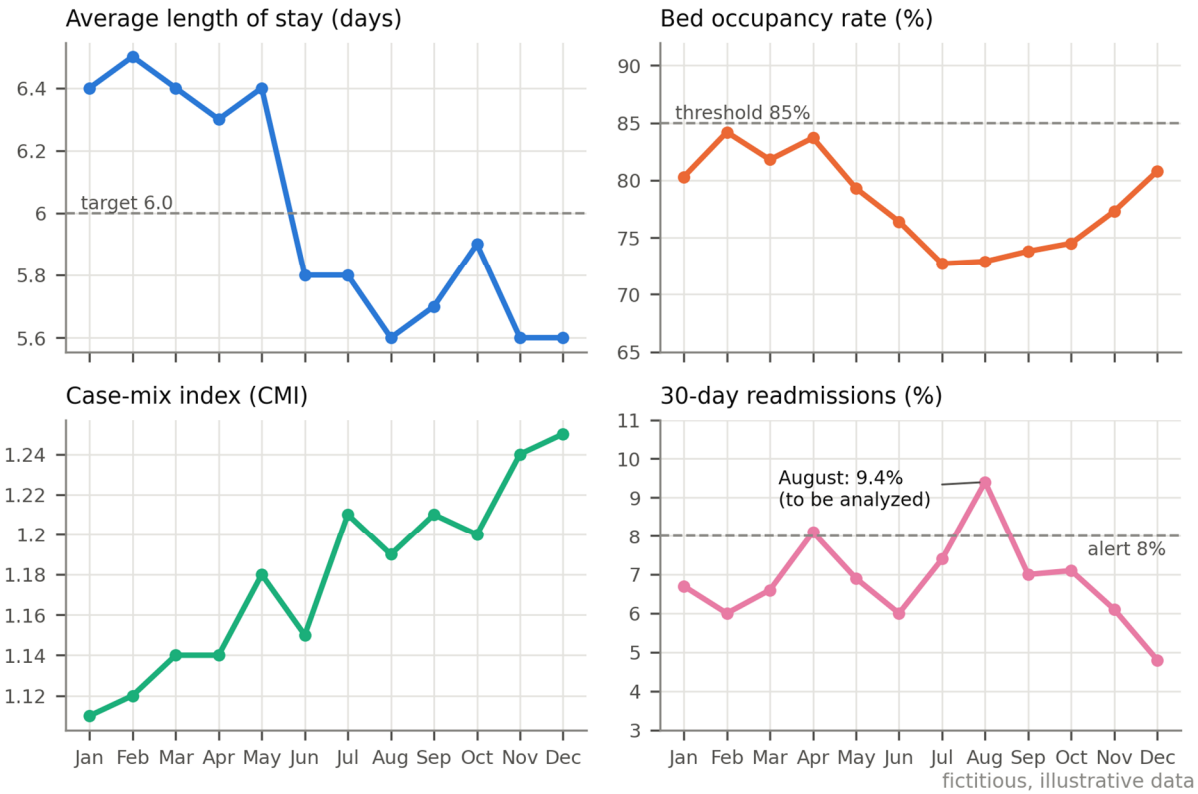


Figure 12.4. Teaching dashboard with four monthly indicators of a hospital: average length of stay, bed occupancy rate, case-mix index, and 30-day readmission rate, with reference thresholds. Fictitious, illustrative data.

Reading the dashboard in Figure 12.4 correctly requires linking the indicators to one another. The ALOS decreases after May, while the CMI increases. The hospital is treating more complex cases in fewer days, which looks like an efficiency gain. The peak in readmissions in August (9.4%), however, may be a sign of premature discharges. A good dashboard does not give the conclusion; it shows where to look for it.

12.7 Resilience: When the System Stops

12.7.1 Cyberattacks

According to press releases from the Romanian National Cyber Security Directorate (Directoratul Național de Securitate Cibernetică, DNSC), between February 10 and 12, 2024, the Backmydata ransomware, from the Phobos family, encrypted the data on the servers of several Romanian hospitals that used the Hipocrate hospital information system. DNSC initially confirmed 21 affected hospitals, and another 79 healthcare facilities were disconnected from the internet for investigation. Also according to DNSC, most of the affected hospitals had backups that were one to three days old, which could make data restoration easier [186]. The preamble of Government Emergency Ordinance (OUG) No. 155/2024 mentions an incident in the first quarter of 2024 that "affected 26 hospitals nationwide through a managed service provider" (our translation) [187]. For system architecture, the lesson is direct: a vendor shared by dozens of hospitals becomes a single point of vulnerability. The course of the attack and the legal framework for cybersecurity are presented in Chapter 5.

The problem is international. Between 2016 and 2021, 374 ransomware attacks on US healthcare providers exposed the data of almost 42 million patients. The annual number doubled, from 43 to 91, and 44.4% of the attacks disrupted care, most often by shutting down electronic

systems [183]. The effects extend beyond the hospital under attack. During a month-long attack on four hospitals in San Diego, the two emergency departments of a neighboring, unaffected institution received an average of 251.4 patients per day, compared with 218.4 before the attack [184]. In Ireland, the May 2021 attack shut down all the information systems of a university cancer center: radiotherapy was interrupted for 113 patients, and the backlog of reporting and data storage was cleared only on day 210 [480].

12.7.2 Downtime Operation

Downtime procedures describe how the hospital works without the system. Winter recommends a business continuity plan that shows how the hospital's functions are sustained even during a major failure [448]. A published example of a complete procedure includes periodic drills, a "backup viewer" with the essential data of inpatients, paper forms prepared in advance, tracking of all changes made during the downtime period, and a recovery (uptime) phase with medication reconciliation and retrospective entry of the documentation [481]. The SIEC-ATI tender specifications also require the ICU system to be able to create a temporary patient record when the hospital's HIS does not respond [244].

Common Pitfalls

- Confusing the information system (everything related to information in the hospital, including paper) with the computer-based information system (the part that runs on computers).
- Believing that linking every application directly to every other one is simpler than using an integration engine. With 10 applications, there are 45 interfaces instead of 10.
- Interpreting a 90% alert override rate as proof that physicians are negligent. Most often, the alert is poorly configured.
- Reading the ALOS in isolation. A short ALOS accompanied by many readmissions does not mean efficiency.
- Assuming that a hospital with a high EMRAM stage automatically has better clinical outcomes. The available studies show associations, not causal relationships.
- Considering paper procedures useless once the hospital has been computerized.

Key Points

- The HIS consists of modules (ADT, FOCG, CPOE, pharmacy, LIS, RIS/PACS, operating theater, ICU, finance) that must exchange data consistently.
- An integration engine reduces the number of links from n(n-1)/2 to n.
- CPOE has been associated with fewer prescribing errors, but it introduces new, technology-generated errors.
- Interaction alerts are overridden in about 90% of cases. Optimizing them is the task of a multidisciplinary committee.
- Good tender specifications contain measurable, vendor-neutral requirements.
- The ALOS, the occupancy rate, the CMI, and readmissions are interpreted together.
- Every computerized hospital needs downtime procedures that are rehearsed periodically.

Self-Assessment Questions

1. (SA) A hospital has 12 clinical applications that must exchange data with one another. How many interfaces are needed if each application is linked directly to every other application?

a) 12
b) 24
c) 66
d) 132
e) 144

2. (SA) Which module of the hospital information system is the reference source for the patient's administrative data (identity, admission, transfer, discharge)?

a) LIS
b) PACS
c) ADT
d) CPOE
e) the data warehouse

3. (SA) In a 30-day month, a 20-bed ward delivered 480 inpatient days. The bed occupancy rate is:

a) 24%
b) 48%
c) 62.5%
d) 80%
e) 96%

4. (MA) Regarding alerts in clinical decision support systems, which statements are correct?

a) a meta-analysis estimated that drug-drug interaction alerts are overridden in approximately 90% of cases
b) alert fatigue has a single operational definition, used by all systematic reviews
c) the hospitals that published their alert optimization process used a multidisciplinary committee
d) an overridden alert is always a physician's error
e) the override rate and the acceptance rate are frequently used indicators for alerts

5. (MA) Which of the following are part of a well-designed downtime procedure?

a) periodic drills of working without the system
b) backup access to the essential data of inpatients
c) medication reconciliation after the system is restarted
d) deleting the paper documentation after recovery, without entering it into the system
e) paper forms prepared in advance

CHAPTER 13

Chapter 13. Artificial Intelligence in Medicine: Methods, Evaluation, Regulation (2026)

Objectives

By the end of this chapter, the student will be able to:

- distinguish a rule-based system from a model obtained through machine learning and place a clinical example in one of the families of methods (supervised learning, unsupervised learning, reinforcement learning, deep learning, large language models)
- describe the stages of the life cycle of a predictive model, from the clinical question to post-deployment monitoring, and recognize data leakage in a study protocol
- calculate, from a confusion matrix, sensitivity, specificity, the predictive values, accuracy, and the F1 score, then recalculate the predictive values for a different prevalence
- interpret a ROC curve, the area under the curve, and a calibration plot, without confusing discrimination with clinical utility
- choose the appropriate reporting guideline (TRIPOD+AI, CONSORT-AI, SPIRIT-AI, DECIDE-AI, CLAIM, STARD-AI) for a given type of study
- list the main risks of AI tools and the obligations that arise, in 2026, from the AI Act, the MDR, and the ICMJE recommendations.

13.1 What Artificial Intelligence in Medicine Is

13.1.1 Definitions and the Two Major Traditions

The term artificial intelligence (AI) has no single definition. An operational definition used in a recent textbook of medical AI describes AI as the science and technology of computing systems that autonomously solve difficult inferential problems, and machine learning (ML) as the science and technology of systems that learn to solve problems by analyzing data related to those problems [482]. Machine learning is therefore a field of AI, not a synonym for it. The same authors warn that reducing the whole of AI to a single subfield, for instance to deep learning, stands in the way of choosing the method suited to the clinical problem [482]. For an introduction in Romanian, with examples from medical practice, see [3], and for the place of AI within medical informatics as a whole, [2].

The first tradition is **symbolic** AI. Knowledge is written explicitly by experts, in the form of rules and logical relations ("IF fever AND positive blood culture THEN..."), and the system makes deductions. This family includes rule-based systems, with historical examples such as MYCIN, and decision support through independent rules written in Arden Syntax [483]. Drug interaction alerts in electronic prescribing often work in this way. The advantage is transparency: each rule can be read and checked. The disadvantage appears when knowledge has to be maintained by hand for thousands of situations.

The second tradition is **data-driven** AI. Here the model receives not rules but examples, from which it builds the relation between inputs and outputs on its own. Neural networks, probabilistic models, and reinforcement learning belong to this family. Aliferis and Simon call them "ontologically shallow" systems: even a deep network, with many hidden layers, contains no explicit representation of the disease, only a learned relation between inputs and outcome [483]. The distinction matters for interpretation, because a data-driven model can perform well without "knowing" anything about pathophysiology.

13.1.2 Types of Learning

Most clinical applications use **supervised learning**. The model receives labeled examples (radiographs marked "pneumothorax" or "normal," admissions marked "developed sepsis" or "did not") and learns to predict the label for new cases [484]. The quality of the label directly limits the quality of the model.

In **unsupervised learning**, the data have no labels. The algorithm looks for structure, for example groups of patients with similar profiles or outliers [484]. The groups obtained are hypotheses, not diagnoses. Aliferis and Simon show that unsupervised clustering should not be used as a substitute for a predictive model, because a "good" clustering depends on the outcome of interest, which the algorithm does not see [485].

Reinforcement learning trains the algorithm through trial and error: behaviors that bring it closer to the desired outcome are "rewarded" and become more likely [484]. In 2020, its impact on clinical medicine was still small [484], and in the scoping review of AI RCTs published up to 2023, only two of the 27 systems that did not work on images used reinforcement learning [486].

13.1.3 Deep Neural Networks

An artificial neural network is a succession of layers of simple computing units. Each unit forms a weighted combination of its inputs and applies a nonlinear function. The network is "deep" when it has many hidden layers, and the weights are adjusted during training so that the error on the known examples decreases [483,484]. Convolutional networks have dominated medical image analysis (see Chapter 7). The price of this performance is the large number of labeled examples required and the opacity of the model, discussed in Section 13.5.

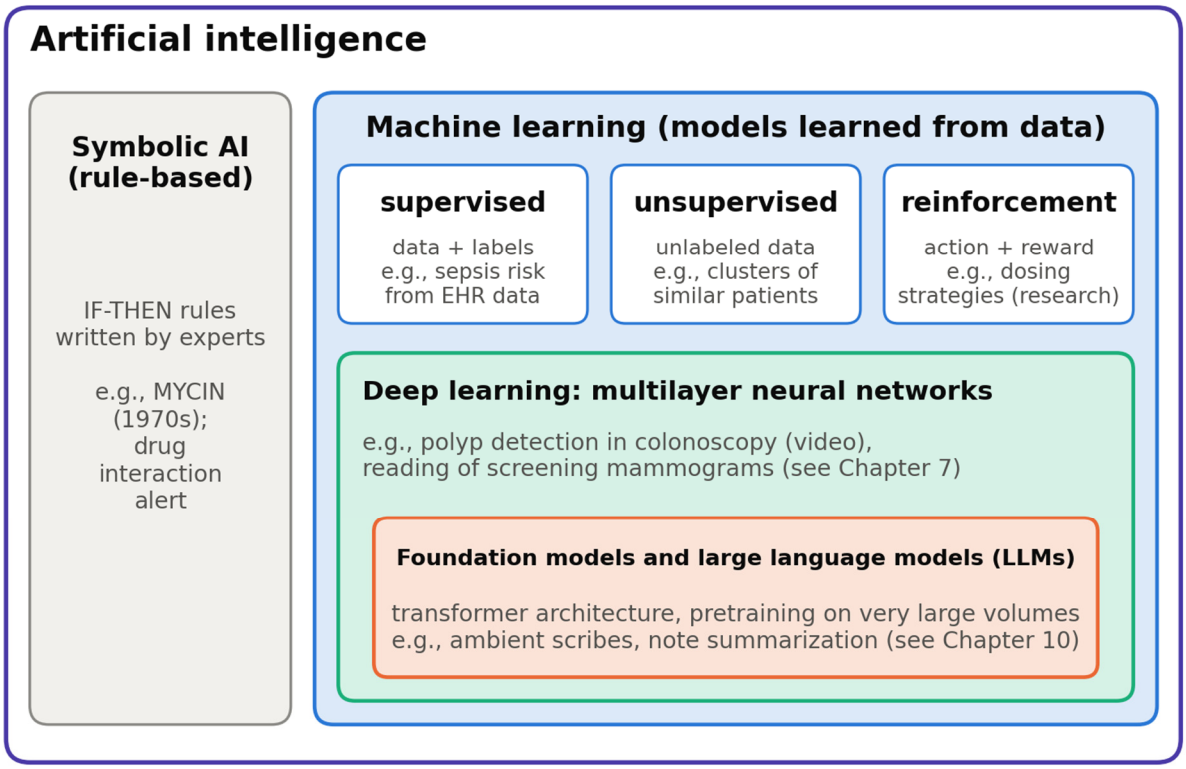


Figure 13.1. The families of artificial intelligence methods and the relations between them, with medical examples. Deep learning is a part of machine learning, and foundation models and large language models are, in turn, deep networks. Teaching diagram based on [483,487].

13.1.4 Foundation Models and Large Language Models

Since 2017, the **transformer** architecture has become the basis of most language processing models. It uses an attention mechanism that weighs, for each word, the importance of the other words in the context [488]. This architecture underlies the models pretrained on very large volumes of text, which can then be adapted to particular tasks through fine-tuning, a process much shorter than the initial training [488].

A **foundation model** is trained through self-supervision on large and diverse data sets and can perform many tasks with few or no task-specific labels. Moor et al. proposed, for medicine, the concept of generalist medical AI, able to combine images, EHR data, laboratory results, genomics, and text [487]. **Large language models** (LLMs) are foundation models for text. They answer freely formulated questions without dedicated training for that task, and chatbots such as ChatGPT are LLMs fine-tuned for dialogue [489].

The practical difference from a classical model is the scope of application. A risk score for a single disease is a focused system whose properties can be verified. An LLM has an unlimited scope, which makes it hard to establish performance and safety for all possible uses [490]. LLM applications in consultation and in mHealth are covered in Chapter 10.

13.2 The Life Cycle of a Model

A model used in patient care passes through distinct stages, from the definition of requirements to integration into the workflow and monitoring [491]. Figure 13.2 summarizes them.

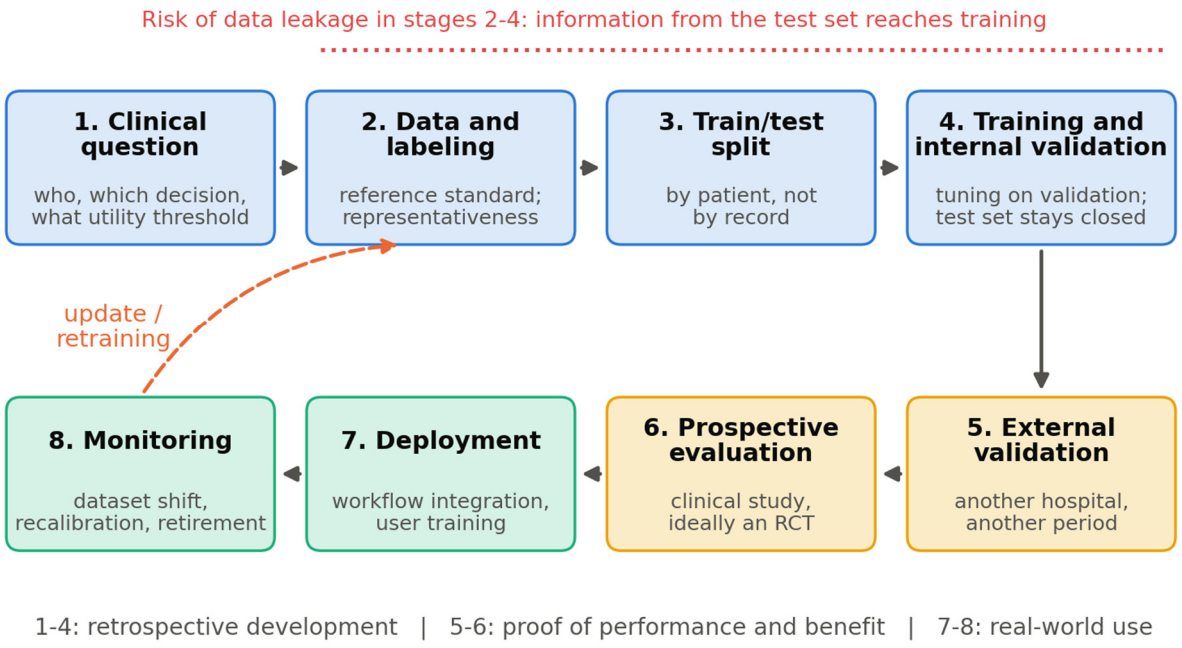


Figure 13.2. The life cycle of a clinical AI model: retrospective development (1-4), proof of performance and benefit (5-6), real-world use (7-8), with the update loop. The dotted line marks the zone in which data leakage usually occurs. Teaching diagram based on [491,492].

The clinical question. The first step is not the choice of algorithm but the definition of the decision that the model supports: who the patient is, at what moment, what action follows a positive result, and what minimum performance would justify use. Aliferis and Simon describe this stage as requirements engineering, with performance targets tied to concrete clinical outcomes [491].

Data and labeling. The data set must resemble the population in which the model will be used. The label is established against a reference standard, the term preferred by the CLAIM 2024 guideline over the expression "ground truth" [493]. Labeling variability and the underrepresentation of some patient groups in the data are sources of bias from this stage onward [494].

Training, validation, testing. The available data are split into three parts: a training set (for the model weights), a validation set (for the choice of variant and hyperparameters), and a test set, opened only once, at the end. The split is made by patient, not by record; otherwise, the same patient appears in both parts.

External validation. The performance obtained on one's own data (internal validation) is almost always optimistic. External validation uses data collected at another site or in another period and explores heterogeneity, fairness, and generalizability [492]. Van Calster et al. stated the conclusion in their title: there is no model "validated" once and for all, because populations, measurement procedures, and practice change [495].

Deployment and monitoring. A model can work well in a study and poorly on the ward. Beyond the edge of the domain it has learned, performance can drop abruptly, a phenomenon that Aliferis and Simon call the "knowledge cliff" [491]. Data set shift occurs when the data encountered in operation no longer resemble the development data, because of technology, the population, or clinical behaviors [496]. In April 2020, the University of Michigan hospital deactivated a sepsis alert model after changes related to COVID-19 altered the relation between

fever and bacterial sepsis [496]. For generative AI, the 2025 NAM report recommends continuous monitoring in the form of "algorithmovigilance" [497].

13.2.1 Data Leakage

Data leakage means that, during development, the model has access to information that it would not have in real use. The result is artificially inflated performance. Kapoor and Narayanan found data leakage in 17 scientific fields that use ML, with 294 papers affected, and proposed a taxonomy of eight types [498]. Typical examples in medicine:

- the same patient (or the same lesion, photographed twice) appears in both training and test data
- normalization or variable selection is performed on the entire data set, before the split
- a variable indirectly contains the outcome: the sepsis diagnosis code, entered after confirmation, used as a predictor of sepsis
- data from the future are used to "predict" the past, for example laboratory tests drawn after the moment at which the model should have raised the alarm.

When errors of this kind were corrected in a reproduction study (outside medicine, on the prediction of civil wars), the complex models no longer outperformed a simple logistic regression [498].

13.3 Performance Evaluation

13.3.1 The Confusion Matrix and Derived Indicators

A binary classifier turns a continuous score (a probability) into a "positive/negative" decision by choosing a threshold. When the decision is compared with the reference standard, each patient falls into one of the four cells of the confusion matrix: true positive (TP), false positive (FP), false negative (FN), true negative (TN). The first six indicators in Table 13.1 are calculated from these four numbers [499,500], and the AUC is derived from the matrices obtained at all possible thresholds.

Table 13.1. Performance indicators calculated from the confusion matrix, with their formulas and interpretation. Source: [499-501].

Indicator	Formula	Question it answers	Depends on prevalence?
Sensitivity (Se)	TP / (TP + FN)	How many diseased people are detected?	No (in principle)
Specificity (Sp)	TN / (TN + FP)	How many healthy people are recognized as such?	No (in principle)
Positive predictive value (PPV)	TP / (TP + FP)	If the test is positive, how likely is the person diseased?	Yes
Negative predictive value (NPV)	TN / (TN + FN)	If the test is negative, how likely is the person healthy?	Yes
Accuracy	(TP + TN) / total	What proportion of decisions are correct?	Yes
F1 score	2·PPV·Se / (PPV + Se)	The harmonic mean of PPV and Se	Yes; ignores TN
Area under the ROC curve (AUC)	The probability that a diseased person has a higher score than a healthy one	How well does the score separate the two groups?	No

Sensitivity and specificity describe the test. The predictive values describe the situation of the patient who already has a result, and they vary with the pretest probability, that is, with the prevalence in the tested population [499]. Accuracy can mislead when the disease is rare: a "test" that declares everyone healthy has an accuracy of 99% in a population with a prevalence of 1%, but it detects no diseased person.

13.3.2 The ROC Curve and the Area Under the Curve

Each threshold produces a different (Se, Sp) pair. The ROC (receiver operating characteristic) curve joins these pairs for all possible thresholds, with sensitivity on the vertical axis and 1 - specificity on the horizontal axis (Figure 13.3). The area under the curve (AUC) ranges from 0.5, for a test that separates no better than chance, to 1, for perfect separation [501]. Values above 0.80 are generally considered clinically useful, but the point estimate is interpreted together with the 95% confidence interval [501].

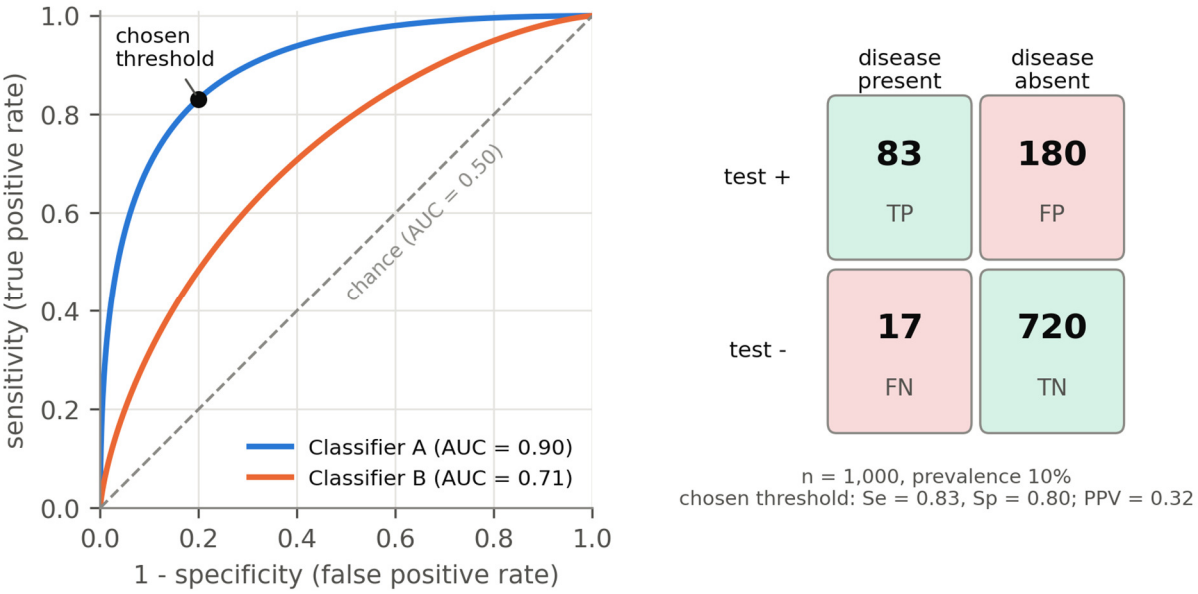


Figure 13.3. Left: the ROC curves of two synthetic classifiers, with normally distributed scores in diseased and healthy people (AUC 0.90 and 0.71). Right: the confusion matrix of classifier A at the marked threshold (Se 0.83, Sp 0.80), for 1,000 patients with a prevalence of 10%. Synthetic, illustrative data, calculated in the figure script.

In the example in Figure 13.3, classifier A, at the chosen threshold, detects 83 of 100 diseased people. However, of the 263 positive results, 180 are false positives, so the PPV is only 83/263 = 0.32. A good AUC therefore does not guarantee that a positive result can be trusted.

A real-world example shows why external validation matters. The Epic Sepsis Model (ESM), deployed in hundreds of US hospitals, was externally evaluated on 38,455 hospitalizations at Michigan Medicine. The AUC was 0.63, the model missed 67% of the patients with sepsis, and it generated alerts for 18% of all hospitalized patients [502]. The second version of the model, evaluated prospectively in four health systems on 227,091 hospitalizations, had an encounter-level AUC between 0.82 and 0.92, but the PPV remained between 0.13 and 0.26, with a heavy alert burden [503].

13.3.3 Dependence on Prevalence: A Numerical Example

Consider an AI tool with Se = Sp = 0.90, applied to 10,000 people in three settings (Figure 13.4):

- **screening in the general population, prevalence 1%**: 100 diseased people, of whom 90 are TP and 10 FN; 9,900 healthy people, of whom 8,910 are TN and 990 FP. PPV = 90/1,080 = 8.3%. Out of every 12 positive results, 11 are false alarms
- **patients referred to a specialist, prevalence 10%**: 900 TP, 100 FN, 900 FP, 8,100 TN. PPV = 900/1,800 = 50%
- **symptomatic patients on the specialty ward, prevalence 30%**: 2,700 TP, 300 FN, 700 FP, 6,300 TN. PPV = 2,700/3,400 = 79.4%, and the NPV drops to 6,300/6,600 = 95.5%.

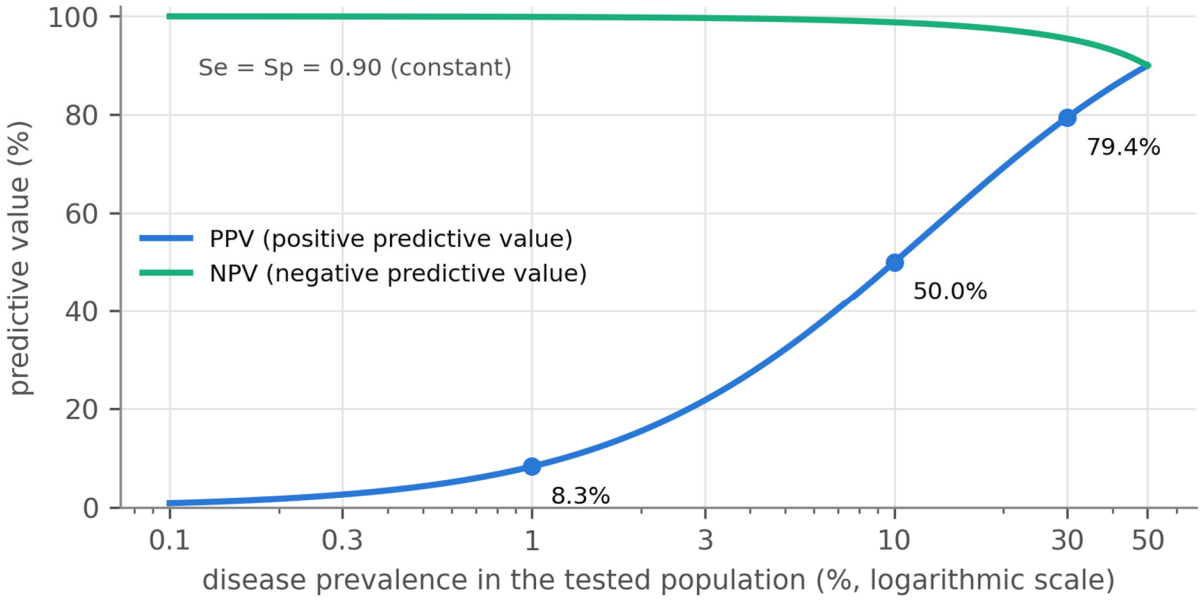


Figure 13.4. Positive and negative predictive value as a function of prevalence, for a test with a constant sensitivity and specificity of 0.90. The dots mark the numerical example in the text (1%, 10%, 30%). Teaching calculation, synthetic data.

The same algorithm, with the same Se and Sp, thus produces positive results with very different meanings. For this reason, performance reported in a tertiary center does not transfer automatically to primary care.

13.3.4 Calibration and Clinical Utility

Discrimination (AUC) tells whether the model ranks patients correctly. **Calibration** tells whether the displayed probabilities are correct: of the patients to whom the model assigns a 20% risk, approximately 20% should have the event. A model can discriminate well and still systematically overestimate risk. In an analysis of 32 performance measures from five domains (discrimination, calibration, overall performance, classification, clinical utility), Van Calster et al. recommend mandatory reporting of the AUC, the calibration plot, a measure of clinical utility such as the net benefit from decision curve analysis, and the distribution of probabilities by outcome category [500]. The F1 score met neither of the two quality criteria analyzed [500].

Decision curve analysis answers a simple question: would using the model do more good than harm? It calculates the net benefit of the model for each threshold probability, that is, for the risk at which a clinician would decide to intervene [504]. It is hard to say how high the AUC must be for a model to be worth using, and the decision curve answers precisely this question [504].

13.3.5 Reporting Guidelines

An AI study can be critically appraised only if it is reported completely. The guidelines in Table 13.2 set out the minimum information, depending on the type of study. TRIPOD+AI, published in the BMJ in 2024, has 27 items, applies to predictive models whether they are built with regression or with ML, and replaces the 2015 TRIPOD checklist [505]. The risk of bias of these studies is assessed with PROBAST+AI (2025) [506].

Table 13.2. Reporting guidelines for studies evaluating AI interventions and models. Source: [493,505-511].

Guideline	Year	Type of study	AI-specific content
SPIRIT-AI	2020	Protocol of a randomized clinical trial	15 new items relative to SPIRIT 2013
CONSORT-AI	2020	Report of a randomized clinical trial	14 new items relative to CONSORT 2010
DECIDE-AI	2022	Early clinical evaluation of a decision support system	17 AI-specific and 10 generic items
TRIPOD+AI	2024	Development or validation of a predictive model	27 items; replaces TRIPOD 2015
CLAIM 2024	2024	AI studies in medical imaging	44 items, "reference standard"
PROBAST+AI	2025	Risk-of-bias assessment (not reporting)	16 + 18 signaling questions
STARD-AI	2025	Diagnostic accuracy of an AI test	18 new or modified items
TRIPOD-LLM	2025	Studies using large language models	19 main items, 50 subitems

CONSORT-AI requires, among other things, a description of the AI intervention, of how the input and output data are handled, and of the human-AI interaction, as well as an analysis of error cases [507]. SPIRIT-AI has the same structure for protocols [508]. DECIDE-AI covers the intermediate stage, in which a system is tried on a small scale in practice, to evaluate safety and human factors before a large study [509]. CLAIM 2024 recommends the terms "internal testing" and "external testing" instead of the ambiguous word "validation" [493].

13.4 Clinical Evidence 2020-2026

Performance on a test set is not the same as benefit for the patient. A 2022 review observed that prospective studies had begun to narrow the gap between research and deployment, especially in image analysis, but it also pointed to persistent obstacles, from lack of data to bias [512]. Benefit is demonstrated in prospective studies, ideally in randomized controlled trials (RCTs). A systematic review published in JAMA Network Open found, up to October 2021, only 41 RCTs of ML-based interventions, with a median of 294 participants. Half were single-center (51%), 37% concerned endoscopy, and none adhered to all CONSORT-AI requirements [513].

A scoping review in The Lancet Digital Health updated the count to November 2023 and included 86 RCTs with AI actually integrated into patient care [486]. Gastroenterology accounted for 43% of the trials, followed by radiology (13%), and 63% were single-center (Figure 13.5). Of the 86 trials, 70 (81%) reported a favorable result for the primary endpoint. The authors nevertheless draw attention to probable publication bias, to poor demographic reporting, and to the fact that only 12 of the 63 trials published after 2021 cited CONSORT-AI [486].

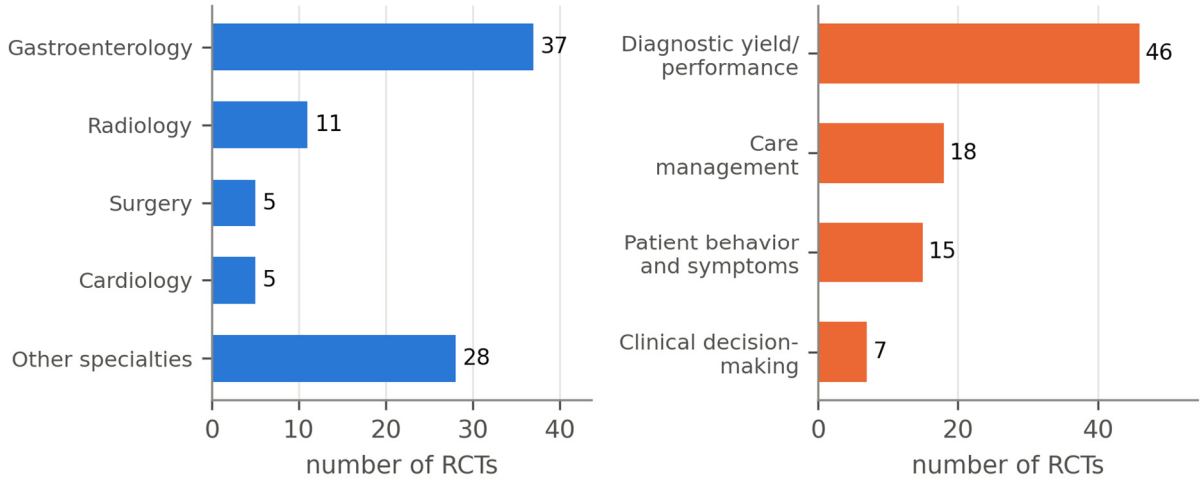


Figure 13.5. The 86 RCTs of AI in clinical practice published between 2018 and November 2023: left, by specialty; right, by type of primary endpoint. Real data taken from [486].

A few representative examples show both what has been achieved and the limits:

- **AI-assisted colonoscopy.** A meta-analysis of 28 RCTs with 23,861 participants found a 20% increase in the adenoma detection rate (RR 1.20) and a 55% decrease in the adenoma miss rate, but no difference for advanced adenomas and 39% more resections of non-neoplastic lesions [514].
- **Mammography screening.** In the Swedish MASAI trial, 105,934 women were randomized to AI-supported reading or standard double reading. The interval cancer rate was noninferior (1.55 vs. 1.76 per 1,000), sensitivity increased from 73.8% to 80.5%, and specificity remained 98.5% in both groups [271]. The technical details are in Chapter 7.
- **LLMs as a diagnostic aid.** In an RCT with 50 physicians, access to an LLM did not improve diagnostic reasoning compared with conventional resources (median score 76% vs. 74%), although the LLM alone scored 16 percentage points higher than the control group [399]. LLM-assisted consultation is discussed in Chapter 10.
- **Ambient scribes.** These systems record the consultation and generate the clinical note. In a pragmatic RCT with 238 physicians from 14 specialties, one of the two applications tested reduced the time spent on notes by 9.5% compared with control, the other produced no statistically significant change (-1.7%), and physicians reported "occasional" clinically relevant inaccuracies [515]. A stepped-wedge randomized study with 66 clinicians found a decrease in burnout (-0.44 points on a 5-point scale) and in time spent on notes (-0.36 hours per day), without a statistically significant increase in professional satisfaction [516]. A third study, a crossover trial with 160 clinicians, compared two products: both reduced burnout scores, but time worked at home, after hours, remained largely unchanged [517]. The 2025 NAM report places ambient scribes among the lower-risk applications, close to current use [497].

The overall picture is therefore mixed. Benefits are best documented for detection tasks in imaging and endoscopy, with intermediate endpoints (detection rate) more often than final clinical endpoints (mortality, complications).

13.5 Explainable AI

Many ML models are "black boxes": they produce a score without a readable justification. Explainable AI (XAI) tries to show what determined a prediction. Two families of methods are encountered most often: attributing importance to each variable, and saliency maps for images.

SHAP values (SHapley Additive exPlanations) borrow from game theory Lloyd Shapley's idea of dividing the payoff fairly among the participants in a coalition. Applied to a model, the method assigns each variable a contribution to the prediction for a patient ("age increased the risk by x, lactate by y"). Aliferis and Simon, however, list nine limitations of Shapley values in explaining ML models, among them the fact that they do not correctly explain even linear models and that they can assign importance to variables unrelated to the outcome; for this reason, they do not recommend them as a routine method [491].

Saliency maps color the pixels that most influenced a network's decision on an image. When evaluated on two public radiograph data sets, all eight saliency methods tested failed at least one trustworthiness criterion (localization, sensitivity to weight randomization, repeatability, reproducibility) and localized lesions less well than networks built specifically for segmentation [518].

Ghassemi et al. called such hopes "false hope": current explanation methods cannot guarantee, at the level of the individual patient, that the explanation reflects the real reason for the prediction. They recommend rigorous internal and external validation as a more direct means of earning trust and warn against imposing explainability as a requirement for clinical models [519]. An example illustrates the problem. Standard deep networks were able to predict the patient's self-reported race from radiographs, CT scans, and mammograms (AUC 0.81-0.99), although no human expert can see this information in the image [520]. A saliency map would not have revealed such a hidden mechanism.

13.6 Risks

Bias and fairness. A model learns from historical data and can reproduce the inequities they contain. Chest radiograph classifiers trained on three large data sets selectively underdiagnosed underserved groups (women, Black patients, patients of low socioeconomic status), with higher rates at the intersection of these categories, for example in Hispanic women [194]. Sources of bias arise in data acquisition, in genetic variation between populations, and in labeling variability [494]. The NAM Code of Conduct makes equity one of its six commitments [521].

Automation bias. The tendency to follow the suggestion of an automated system also affects experts. In a prospective experiment, 27 radiologists read 50 mammograms with purported AI assistance. When the BI-RADS suggestion was wrong, the proportion of correct readings dropped from approximately 80% to 20% (inexperienced readers), 25% (moderately experienced readers), and 46% (very experienced radiologists) [281].

Hallucinations. For generative AI, a "hallucination" (or confabulation) is incorrect or unfounded information presented plausibly, a consequence of the way the model generates text by prediction [497]. When asked to retrieve the literature for 11 systematic reviews, GPT-3.5, GPT-4, and Bard produced hallucinated references (with at least two of the following wrong: title, first author, year) in 39.6%, 28.6%, and 91.4% of cases, respectively [44]. Any reference proposed by an LLM is checked in PubMed or CrossRef before it is used.

Security. Attacks are not aimed only at servers. Replacing just 0.001% of the training tokens with medical misinformation produced models that propagated medical errors, even though these models achieved the same scores on standard benchmarks [522].

Accountability. An AI tool cannot be held legally or ethically accountable. The accountability principle in the NAM code requires identifiable and measurable actions, clear documentation of benefits, and clear controls for possible adverse consequences [521]. The decision and the signature remain the physician's.

Common Pitfalls

- Entering a real patient's data (name, personal numeric code (CNP), images, medical history) into a public chatbot. Data concerning health are a special category, the processing of which is prohibited as a rule and permitted only under the exceptions in Article 9 GDPR [118]. A public chatbot is not a hospital-approved tool for processing them.
- Delegating the decision: "the AI said so" is not a clinical justification. The suggestion is checked, and the physician is accountable for the decision.
- Presenting generated text as one's own, in a term paper, a graduation thesis, or an article, without disclosure.
- Citing a reference proposed by an LLM without checking it in PubMed or CrossRef.
- Judging a model only by its AUC or accuracy, without prevalence, calibration, and external validation.
- Confusing "internally validated" with "shown to benefit patients."

13.7 Regulation in 2026

13.7.1 The European Regulation on AI (AI Act)

Regulation (EU) 2024/1689, known as the AI Act, was adopted on June 13, 2024, and entered into force on August 1, 2024 [22]. As a regulation, it applies directly in all Member States, including Romania, without a transposition law. The approach is risk-based, with four levels (Table 13.3) [523].

Table 13.3. Risk levels in the AI Act, with examples relevant to health. Source: [22,523].

Level	What it means	Examples	Consequence
Unacceptable risk	Prohibited practices (Art. 5)	Harmful manipulation, social scoring	Prohibited since February 2, 2025
High risk	Serious risk to health, safety, or fundamental rights	AI-enabled medical software assessed by a notified body, emergency patient triage	Risk management system, high-quality data, documentation, human oversight
Transparency risk	Users must know that they are interacting with AI	Chatbots, generated content	Informing the user, marking of synthetic content
Minimal risk	Most applications	Spam filters, video games	No new obligations

The link with medical devices is direct. Article 6(1) classifies as high-risk an AI system that is itself a product, or a safety component of a product, covered by the legislation listed in Annex I and required to undergo a third-party conformity assessment, and the MDR and IVDR are part of this annex [22]. In practice, AI-enabled medical software in class IIa or higher, assessed by a notified body, falls under the high-risk regime. Annex III adds, among others, emergency patient triage systems and AI used for risk assessment in health insurance [22].

The timeline was amended in 2026. The Digital Omnibus on AI, Regulation (EU) 2026/1744 of July 8, 2026, published on July 24, 2026, postponed the obligations for high-risk AI systems [197]. As of September 27, 2026, the situation is as follows:

- **February 2, 2025**: the prohibitions (Art. 5) and the AI literacy obligation (Art. 4) apply [22]. The Omnibus reworded Art. 4: providers and deployers take measures to support the AI literacy of their staff, without guaranteeing a specific level for each person [197].
- **August 2, 2025**: the rules for general-purpose AI models, governance, and penalties apply [22].
- **December 2, 2026**: the deadline for the marking of synthetic content by generative systems placed on the market before August 2, 2026 [197].
- **December 2, 2027**: the obligations for stand-alone high-risk AI systems, listed in Annex III [197].
- **August 2, 2028**: the obligations for high-risk AI systems embedded in products listed in Annex I, and therefore also for AI-enabled medical devices [197].

Separately, in December 2025 the Commission proposed a simplification of the MDR and IVDR, including moving them to Section B of Annex I, which would limit the application of the AI Act to medical devices [524]. As of September 27, 2026, this is only a proposal.

13.7.2 The MDR and Medical Software

Regulation (EU) 2017/745 on medical devices (MDR) has applied since May 26, 2020 [284]. For software, Rule 11 of Annex VIII sets the class: software that provides information used to make diagnostic or therapeutic decisions is at least class IIa, class IIb if the decision may cause a serious deterioration of a person's state of health, and class III if it may cause death or an irreversible deterioration [284]. The qualification and classification of software are explained in guidance document MDCG 2019-11, revised in June 2025 [372]. An algorithm intended by its manufacturer to propose diagnoses or treatments is therefore, in general, a medical device, not merely a "program."

Data protection remains a separate framework. Data concerning health are a special category under the GDPR [118], and Romanian Law 190/2018 permits the use of health data for automated decision-making or profiling only with explicit consent or on the basis of express legal provisions, with safeguards [176].

13.7.3 The FDA, Codes of Conduct, and the WHO Guidance

In the United States, the FDA publishes a list of medical devices that use AI. The official file downloaded from the page updated on September 4, 2026, contained 1,614 devices, with decisions up to June 29, 2026, of which 1,230 (76.2%) were in radiology, and in 2025 alone 335 were authorized [267]. The FDA states that the list is not exhaustive. The trend by year is presented in Chapter 7.

Alongside legislation, there are voluntary codes. The Code of Conduct of the US National Academy of Medicine (NAM, 2025) sets out ten principles (engaged, safe, effective, equitable, efficient, accessible, transparent, accountable, secure, adaptive) and six commitments, among them protecting human health as the primary aim, equity, and open monitoring of performance [521]. The NAM report on generative AI from the same year ranks LLM applications by maturity: in the short term, patient education and information synthesis; in the long term, virtual health assistants [497]. On January 18, 2024, the WHO published its guidance on the ethics and

governance of large multimodal models, with more than 40 recommendations and five broad applications: diagnosis and clinical care, patient-guided use, administrative tasks, medical and nursing education, and scientific research and drug development [410,411].

Romania in 2026

- The AI Act and the MDR apply directly, as European regulations. As of September 27, 2026, the obligations for AI-enabled medical devices as high-risk AI systems are scheduled for August 2, 2028 [197].
- Each Member State must have at least one AI regulatory sandbox operational by August 2, 2027 [197].
- The AI literacy obligation (Art. 4), in the form amended in July 2026, already applies to hospitals that deploy AI systems [197].
- Health data used for automated decision-making fall under Article 3 of Law 190/2018 [176].
- For the student and the resident physician, the practical consequence is simple: public AI tools receive no patient data, and every output is checked and endorsed by a human.

13.8 Further Reading

Advanced research methods (digital twins, federated learning, the application of TRIPOD+AI in the actual development of a model) are beyond the level of this course and are covered in the author's doctoral monograph. A synthesis of the state of digital twins in oncology can be found in [525].

Example

What an AI tool can and cannot do in a manuscript. The ICMJE Recommendations, updated in January 2026, set the rules [202]:

- A chatbot cannot be an author, because it cannot be accountable for the accuracy, integrity, and originality of the work.
- The use of AI is disclosed at submission, in the cover letter and in the manuscript: writing assistance in the acknowledgments section; use in data collection, analysis, or figure generation in the methods section.
- AI-generated material cannot be cited as a primary source. Authors check the text, because it can be incorrect, incomplete, or biased, and they ensure that it contains no plagiarism, either in the text or in the images.
- Failure to disclose may require corrections and, in certain situations, may be considered misconduct.

Checklist before submission: (1) I have recorded which tool, which version, and for what purpose I used it; (2) I have checked every statement and every reference against the original source; (3) I have not entered patient data into a public service; (4) the disclosure appears in the appropriate section; (5) all human authors take responsibility for the final text.

Key Points

- Machine learning is a part of AI; deep learning and LLMs are, in turn, particular cases of machine learning.
- A model is judged on external data, prospectively, and is monitored after deployment, because data change.
- Data leakage artificially inflates performance. The split is made by patient, before any processing.
- Sensitivity and specificity describe the test. PPV and NPV depend on prevalence.
- AUC measures discrimination, not calibration and not clinical utility.
- An inventory up to November 2023 found 86 RCTs with AI integrated into practice, most of them single-center and most often with intermediate endpoints.
- Post hoc explanations (SHAP, saliency maps) do not replace validation.
- AI cannot be an author and cannot sign a decision. Its use is disclosed.

Self-Assessment Questions

1. (SA) An AI tool has a sensitivity of 0.90 and a specificity of 0.90. Applied in a population with a disease prevalence of 1%, its positive predictive value is approximately:

a) 90%
b) 50%
c) 8%
d) 1%
e) 99%

2. (SA) Which of the following situations represents data leakage?

a) using a test set from another hospital
b) the presence of the same patients in both the training set and the test set
c) reporting the confidence interval for the AUC
d) recalibrating the model after deployment
e) having the images labeled by two independent radiologists

3. (SA) The reporting guideline intended for studies that develop or validate a clinical predictive model, built with regression or with machine learning, is:

a) CONSORT-AI
b) SPIRIT-AI
c) DECIDE-AI
d) TRIPOD+AI
e) CLAIM 2024

4. (MA) Regarding the ROC curve and the area under the curve (AUC), the following are true:

a) the ROC curve plots sensitivity against 1 - specificity, for all thresholds
b) an AUC of 0.5 corresponds to discrimination no better than chance
c) a high AUC guarantees that the probabilities displayed by the model are well calibrated
d) the AUC does not depend on the chosen decision threshold
e) an AUC of 0.90 guarantees a PPV of at least 0.90

5. (MA) According to the ICMJE recommendations on the use of AI by authors:

a) a chatbot can be listed as an author if it wrote most of the text
b) the use of AI for writing is described in the acknowledgments section
c) the use of AI for data analysis or figure generation is described in the methods section
d) human authors are responsible for AI-generated content
e) an AI-generated text can be cited as a primary source

CHAPTER 14

Chapter 14. Synthesis: The Physician's Competencies in the Digital Environment and the Ethics of Using Artificial Intelligence

Objectives

By the end of this chapter, the student will be able to:

- place the 13 previous chapters on a map with five layers (data, standards, systems, clinical use, governance) and explain, using a clinical example, how the layers are connected
- check their digital competencies at the end of the first year against a concrete list and indicate the practical session in which each competency was practiced
- place correctly in time the deadlines of the AI Act, Regulation (EU) 2026/1744, the EHDS, NIS2, and eIDAS 2 that will change medical practice by 2031
- apply a decision checklist before using an artificial intelligence (AI) tool for a study task or for clinical work
- answer a set of review questions covering Chapters 1-13

14.1 The Map of the Discipline

A student who has worked through 13 chapters may feel that they have learned 13 separate subjects: information theory, databases, codes, the GDPR, signals, images, records, telemedicine, mobile applications, signatures, hospitals, and artificial intelligence. The connection between them becomes visible if we follow the path of a single piece of clinical information. A blood pressure of 150/95 mmHg, measured in a fictitious patient, is first a datum with five components: patient, parameter, value, time, and method (Chapter 2). To be stored, it enters a table with a primary key and validation rules (Chapter 3). To be understood by another system, it receives a LOINC code and becomes a FHIR Observation resource with two components, systolic and diastolic (Chapters 4 and 8, Practical Session 8). It then reaches the hospital's electronic health record (EHR), the health insurance platform, or the patient portal (Chapters 11 and 12). The physician uses it in a teleconsultation or in a risk score (Chapters 9, 10, and 13). All along the way, the rules of confidentiality, security, and liability decide who sees it and what can be done with it (Chapters 5 and 11).

Figure 14.1 summarizes this path in five layers. The order is not accidental. Each layer presupposes that the layers beneath it have been dealt with: an AI algorithm trained on uncoded or incomplete data will perform poorly no matter how sophisticated it is, and a well-drafted legal rule protects nothing if the hospital has no backups and no audit logs. The dotted loop in the figure shows that governance acts back on the data. The GDPR, the EHDS, and the AI Act determine which data may be collected, how long they are kept, and under what conditions they may be reused for research.

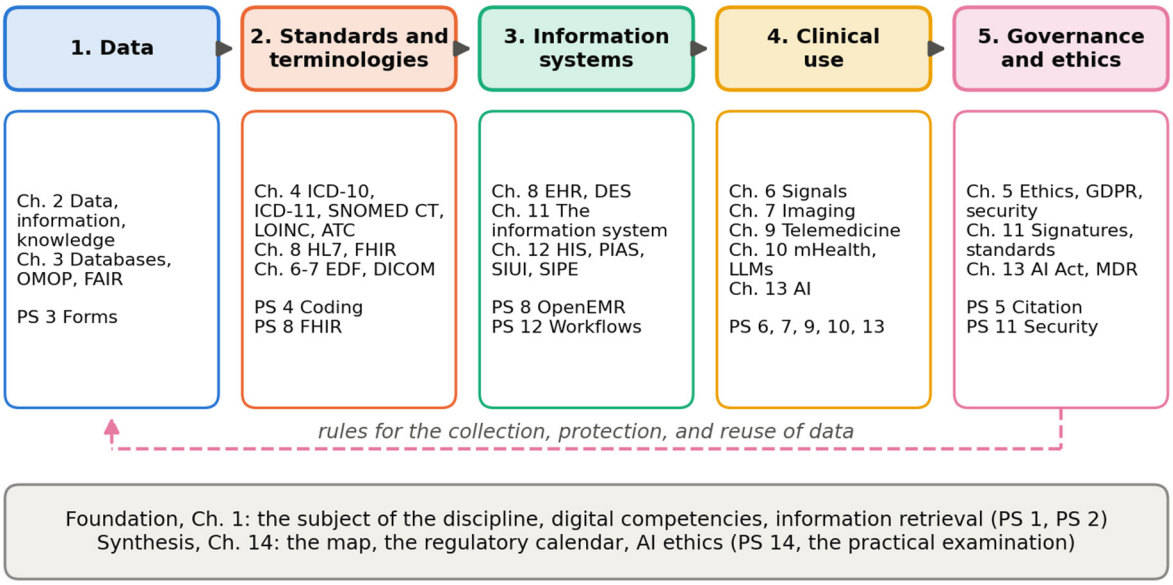


Figure 14.1. Map of the discipline: the five layers (data, standards and terminologies, information systems, clinical use, governance and ethics), with the chapters and practical sessions of the book. The dotted loop shows the effect of the rules on the collection and reuse of data. Teaching diagram.

The map also helps when reading the literature. When an article announces that an AI model "outperformed physicians," the questions are asked layer by layer: which data were used and how representative they were, which reference standard established the label, which system the model was integrated into, which clinical outcome was measured, and who is liable if the model is wrong. Chapter 13 showed that an inventory updated through November 2023 found 86 randomized trials in which AI was actually integrated into care, most of them single-center [486]. The map explains why: each layer adds conditions that a retrospective study does not test.

14.2 Competencies at the End of the First Year

Chapter 1 presented three reference frameworks. The recommendations of the International Medical Informatics Association (IMIA), revised in 2023 on the basis of an international Delphi study, describe the educational needs of three roles, and for the medical student the relevant role is that of the competent user [29]. The DECODE framework groups the digital health competencies of medical education into four domains: digital health professionalism, patient and population digital health, health information systems, and health data science [30]. The European DigComp 3.0 framework keeps five competence areas and integrates AI-related aspects into all of them [34].

For health professionals in the EU, Policy Brief 71 of the European Observatory on Health Systems and Policies found that the emphasis falls on information and data literacy, communication and collaboration, safety and security, problem solving, and digital content creation. The same document considers the training on offer insufficient and notes that participants from Romania reported the English language as a barrier [35]. This textbook, written for the Romanian-, English-, and French-language study lines, addresses this barrier in part through its Romanian edition.

Table 14.1 turns these frameworks into a checklist. Each row describes an action that can be observed and graded, not a general attitude. A student who can check off every row has the foundation required for the practical examination (Practical Session 14).

Table 14.1. Checklist of digital competencies at the end of the first year, aligned with the frameworks in Chapter 1. Source: adapted from [29,30,34,35].

Area (DigComp)	The student can...	Where it was practiced	DECODE domain
Information and data	build a PubMed query with MeSH terms, operators, and filters, and verify a reference by its PMID or DOI	Ch. 1; PS 2	health data science
Information and data	classify a datum by structure and scale and write a data dictionary	Ch. 2, 3; PS 3	health data science
Information and data	code a diagnosis in ICD-10 and ICD-11, stating the version used	Ch. 4; PS 4	information systems
Communication and collaboration	share a document with the minimum necessary role and write a teleconsultation note	PS 1, PS 9	digital professionalism
Communication and collaboration	recognize the appropriate standard (HL7 v2, FHIR, DICOM, IPS) in a scenario	Ch. 7, 8; PS 8	information systems
Content creation	write a text with Vancouver citations and an AI disclosure statement	Ch. 5; PS 5	digital professionalism
Safety	choose a long password, enable two-step verification, and recognize phishing	Ch. 11; PS 1, PS 11	digital professionalism
Safety	apply the GDPR to patient data: special category, pseudonymization, notification	Ch. 5	patient and population health
Problem solving	measure in millimeters on a calibrated DICOM image and calculate SDNN and RMSSD	Ch. 6, 7; PS 6, PS 7	health data science
Problem solving	calculate Se, Sp, PPV, and NPV at different prevalences and verify an answer given by an LLM	Ch. 10, 13; PS 10, PS 13	health data science

The list has an obvious limitation. It checks user skills, not specialist competencies. Students must, however, know when they need a specialist and what questions to ask.

14.3 The 2024-2031 Regulatory Calendar That Will Change Practice

Medical informatics cannot be taught in 2026 without a calendar. Many of the rules discussed in this book have been adopted but apply in stages, and some deadlines have already been amended. Figure 14.2 and Table 14.2 bring together the dates as verified on September 27, 2026.

Artificial intelligence. Regulation (EU) 2024/1689 (the AI Act) entered into force on August 1, 2024. The prohibitions in Article 5 and the AI literacy obligation in Article 4 have applied since February 2, 2025, and the rules on general-purpose AI models, governance, and penalties since August 2, 2025 [22]. The Digital Omnibus on AI, Regulation (EU) 2026/1744 of July 8, 2026, entered into force on July 27, 2026, and shifted the deadlines for high-risk systems. For stand-

alone systems listed in Annex III, which include the triage of patients in emergency care, the obligations apply from December 2, 2027. For systems integrated into products covered by Annex I, and therefore also for AI-enabled medical devices, they apply from August 2, 2028 [197]. The same regulation reworded Article 4: providers and deployers take measures to support the AI literacy of their staff, without having to guarantee a particular level for each person [197]. The MDR and the IVDR remain in Section A of Annex I [526]. The Commission's proposal of December 2025 to move them to Section B had not been adopted at the time of writing [524].

Health data. Regulation (EU) 2025/327 on the European Health Data Space (EHDS) applies, in general, from March 26, 2027. Patient summaries, e-prescriptions, and e-dispensations enter the mandatory regime on March 26, 2029, and medical imaging, laboratory results, and discharge reports on March 26, 2031. Secondary use, through health data access bodies, also applies from 2029 [23,319]. The Data Act has applied since September 12, 2025, and the obligation for connected products to allow access to the data they generate concerns products placed on the market after September 12, 2026 [527]. By contrast, the "Digital Omnibus" proposal that would amend the GDPR had not been adopted as of the update of August 1, 2026 [528]. The GDPR therefore applies in its familiar form.

Security and digital identity. The NIS2 Directive has applied since October 18, 2024, and health is a sector of high criticality. A significant incident is reported in three stages: an early warning within 24 hours, an incident notification within 72 hours, and a final report within one month at most [188]. Romania transposed the directive through Government Emergency Ordinance (OUG) No. 155/2024, published on December 31, 2024, with the National Cyber Security Directorate (DNSC) as the competent authority and national incident response team [187]. The eIDAS 2 Regulation has been in force since May 20, 2024 [434], and the European Commission states that European Digital Identity Wallets will be available by the end of 2026 [437]. Law No. 214/2024 replaced Law No. 455/2001 on the electronic signature [436].

Classifications. ICD-11 has been in force since January 1, 2022, and the WHO fact sheet of May 2026 reports 132 countries at various stages of implementation [140,142]. The current version of the browser is ICD-11 MMS release 2026-01 [141]. In Romania, the diagnoses of hospitalized patients are still coded in ICD-10-AM, 3rd edition, for grouping under RO DRG v.1.1 [139]. We did not identify any legal act or official timeline for the transition to ICD-11.

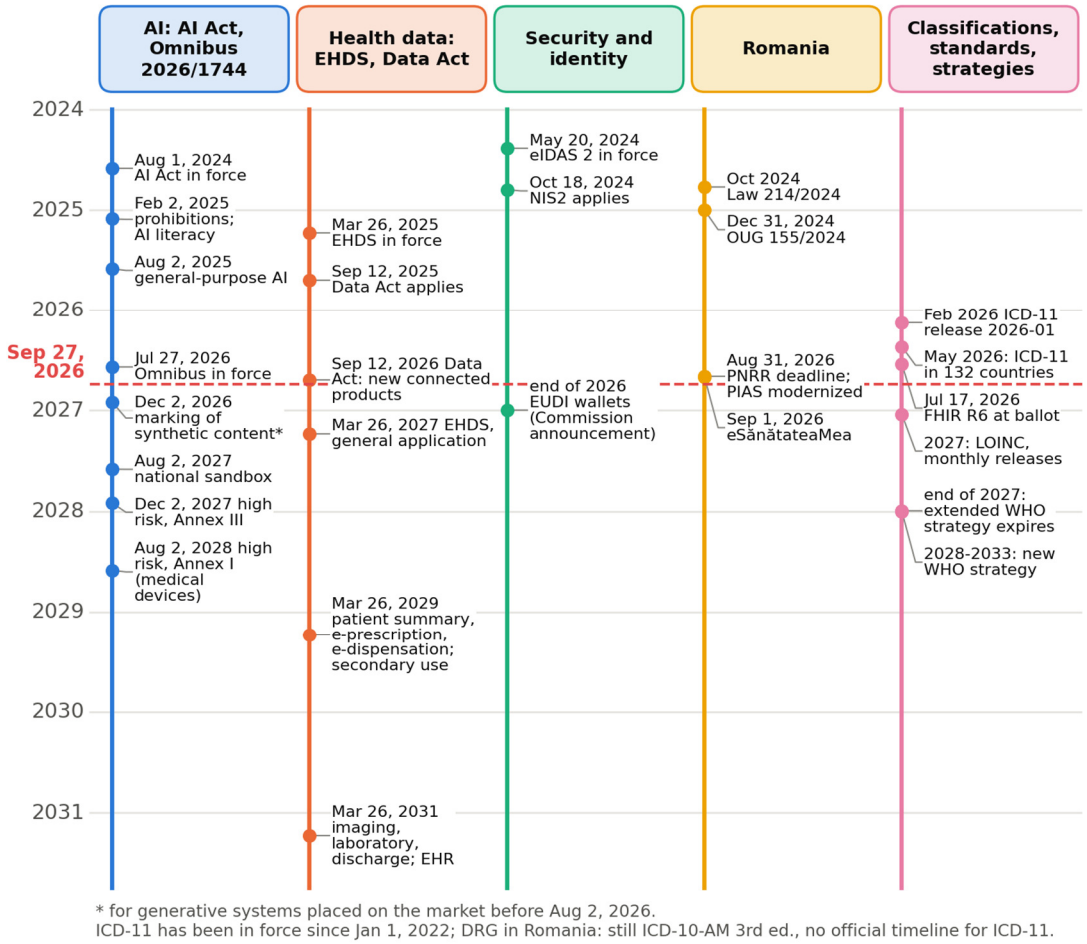


Figure 14.2. Timeline of the regulations and standards that are changing medical practice, 2024-2031, with a line marking the reference date of the book (September 27, 2026). Data from [22,23,38,142,187,188,197,307,323,324,434,436,437,527,529].

Table 14.2. The 2024-2031 regulatory calendar and its consequences for the physician. Source: [22,23,187,188,197,434,437].

Date	Act	What applies	What it means for the physician
May 20, 2024	eIDAS 2 (2024/1183)	framework for the European identity wallet	digital identification recognized throughout the EU
October 18, 2024	NIS2 (2022/2555)	cybersecurity in critical sectors, including health	incidents are reported within 24 h, 72 h, and one month
December 31, 2024	OUG 155/2024	transposition of NIS2 in Romania	DNSC is the competent authority
February 2, 2025	AI Act, Art. 4 and 5	prohibited practices, AI literacy	the institution supports staff training
July 27, 2026	Regulation 2026/1744	new timeline for high-risk systems	the 2026 deadlines no longer apply
March 26, 2027	EHDS (2025/327)	general application	institutions, contact points, access bodies
December 2, 2027	AI Act, Annex III	high risk, stand-alone systems	for example, emergency patient triage
August 2, 2028	AI Act, Annex I	high risk in regulated products	AI-enabled medical devices (MDR)
March 26, 2029	EHDS	summary, e-prescription, e-dispensation; secondary use	mandatory exchange via MyHealth@EU
March 26, 2031	EHDS	imaging, laboratory, discharge; EHR systems	the entire priority record becomes accessible to the patient

Romania in 2026

- On September 1, 2026, the National Health Insurance House (CNAS) launched the eSănătateaMea portal, with authentication through ROeID, the prescriptions of the past 5 years, and the health history of the past 10 years [324].
- The project funded through the National Recovery and Resilience Plan (PNRR) to modernize PIAS, the IT platform of the health insurance system, was declared completed on August 31, 2026, the European deadline for milestones and targets [323,468].
- Romania was not a member of SNOMED International as of September 27, 2026 [151].
- The status of the European Digital Identity Wallet in Romania could not be confirmed from official sources at the time of writing.

14.4 The Ethics of AI Use by Students and Physicians

Chapters 5, 10, and 13 discussed AI from different angles: data protection, consultation with a language model, and model evaluation. The principles that emerge are the same for a student writing a paper and for a physician using a tool on the ward.

Confidentiality. In Romania, information about a patient's condition remains confidential even after death [171]. Data concerning health are a special category whose processing is prohibited as a rule and permitted only under the exceptions in Article 9 of the GDPR [118]. Text entered into a public chatbot ends up on the provider's servers. For this reason, only fictitious cases or genuinely anonymized data are entered into public tools, and keeping "just the initials" anonymizes nothing. Law No. 190/2018 adds a condition for automated decisions based on health data: explicit consent or express legal provisions, with safeguards [176].

Verification. A language model does not consult a database of facts and does not know when it is wrong. Among 636 references, 55% of those generated by GPT-3.5 and 18% of those generated by GPT-4 did not exist [408]. In literature retrieval for 11 systematic reviews, the hallucination rates were 39.6%, 28.6%, and 91.4% for GPT-3.5, GPT-4, and Bard, respectively [44]. Even in a controlled setting for generating clinical notes, 1.47% of sentences contained hallucinations and 3.45% contained omissions, and 44% of the sentences with hallucinations had major errors [406]. The practical rule remains the one from Chapter 1: every statement is checked against a guideline or a primary source, and every reference is looked up in PubMed or CrossRef.

Responsibility. The performance of a model on its own does not guarantee the performance of the person who uses it. In the randomized study by Bean and colleagues, three models working alone identified the relevant condition in 94.9% of the scenarios, but the participants who used them succeeded in fewer than 34.5% of cases [403]. Conversely, a wrong suggestion can drag down even an expert: when the system proposed an incorrect BI-RADS category, the proportion of correct readings by radiologists fell from about 80% to 20-46%, depending on their experience [281]. The Code of Conduct of the US National Academy of Medicine calls for identifiable accountability and for controls against adverse consequences [521]. The decision and the signature remain the physician's.

Disclosure. The ICMJE Recommendations, updated in January 2026, require disclosure of AI use: writing assistance goes in the acknowledgments section, and use for data, analysis, or figures goes in the methods, with the tool and its version. A chatbot cannot be an author, and AI-

generated material is not cited as a primary source [202]. COPE has adopted the same position [203]. Detectors of generated text do not solve the problem. They frequently classify texts written by non-native speakers of English as generated [205], and this includes the texts of Romanian students who write in English.

Equity. An algorithm can be accurate on average and still be inequitable across subgroups. Occult hypoxemia, undetected by pulse oximetry, occurred in 11.7% of the measurements in Black patients, compared with 3.6% in White patients [192]. Chest radiograph classifiers selectively underdiagnosed women, Black patients, and patients of low socioeconomic status [194]. The STANDING Together recommendations call for documenting the composition of data sets and for evaluating performance across population groups [196]. A systematic review of the ethics of language models in medicine found the same recurring concerns (fairness, bias, non-maleficence, transparency, privacy) and the same requirement for human oversight [409].

14.4.1 Decision Checklist: Can I Use an AI Tool for This Task?

The principles above translate into five questions to be asked before any use (Figure 14.3). The order matters, because the first two questions can stop the task before the tool is even opened.

1. Am I entering data that could identify a patient? If so, I stop. I work on a fictitious case, or I use only the tool approved by the institution for clinical data.
2. Do the rules of the context allow AI? The discipline, the examination, the journal, or the hospital may prohibit its use. In the practical examination, AI is used only if the task explicitly requires it.
3. Can I verify every statement and every reference against an independent source? If not, I do not use the content.
4. Does the decision or the final text remain mine, and can I justify it without the tool? If not, I do not use the tool for this task.
5. Could the result disadvantage a group of patients? If so, I check the performance across subgroups and report the limitation.

If every answer allows the task to continue, the tool is used, the result is verified, and the use is disclosed, stating the tool, the version, and the purpose.

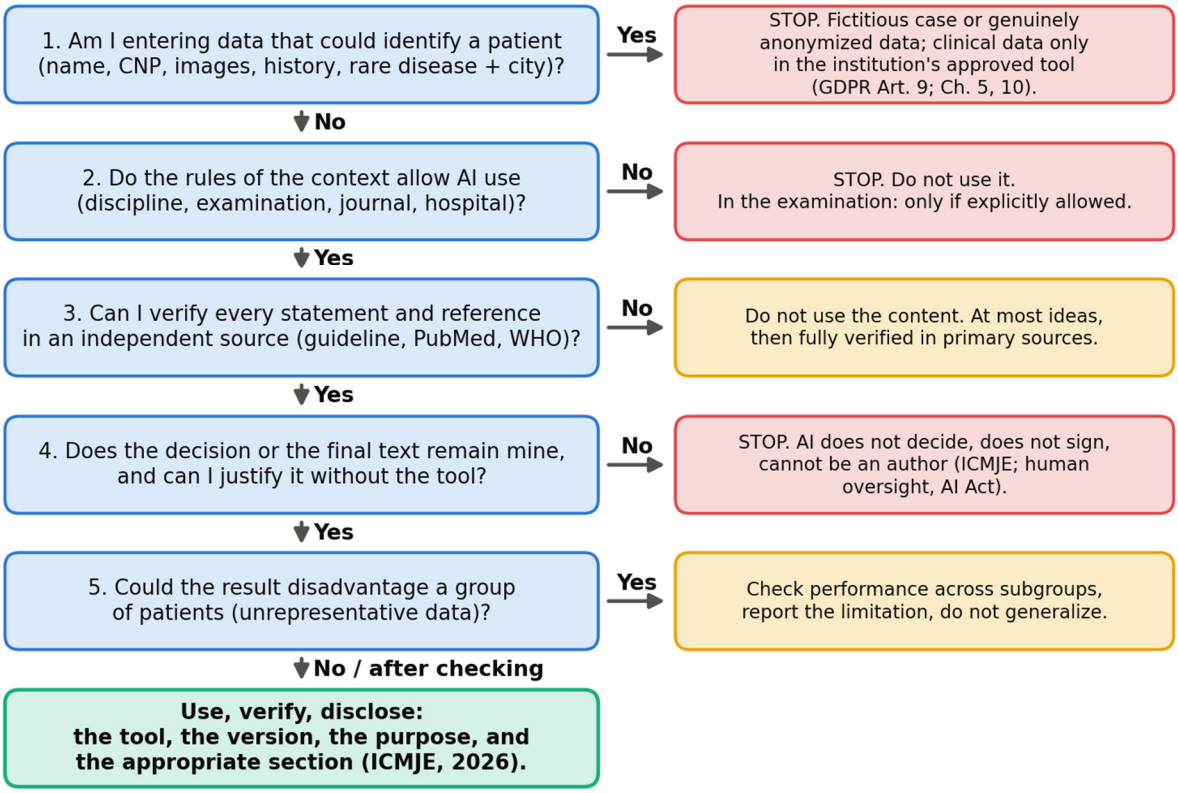


Figure 14.3. The decision checklist "Can I use an AI tool for this task?" as a flowchart. Red exits stop the task; yellow exits require additional measures. Developed from [22,118,196,202].

14.5 Directions Beyond 2026

A few directions can be described with caution, starting from documents that have already been published. The first is the reuse of data. From 2029, researchers will request access to health data through national bodies connected to HealthData@EU, and the data will be processed only in secure environments [23]. The quality of the data collected in hospitals today will determine what can be studied then.

The second concerns the models. Moor and colleagues proposed the concept of generalist medical AI, capable of combining images, EHR data, laboratory tests, genomics, and text, and they anticipate that such models will challenge current approaches to regulation and validation [487]. Validation, however, is never finished. Van Calster and colleagues argue that populations, measurement procedures, and practice all change, so no model is validated once and for all [495]. The example of the sepsis model that was deactivated in 2020, after the pandemic had changed the relationship between fever and bacterial sepsis, shows why post-deployment monitoring becomes a permanent task [496]. The NAM report on generative AI recommends precisely this kind of continuous monitoring [497].

The third concerns the market. The FDA list of AI-enabled medical devices had 1,614 entries in September 2026, 76.2% of them in radiology [267]. The share of radiology nevertheless fell from 85.5% in 2020-2022 to 77.5% in 2023-2025, and 96.2% of the devices analyzed were cleared through the 510(k) pathway, which is based on substantial equivalence [268].

The fourth concerns standards. As of September 27, 2026, FHIR R6 was still at the ballot stage [307], LOINC has announced a move to monthly releases in 2027 [529], and the WHO

global strategy on digital health has been extended until 2027, with a new strategy requested for 2028-2033 [38]. Every number in this book should therefore be checked again a year from now.

Advanced research methods, such as digital twins and federated learning, are covered in the author's doctoral monograph.

14.6 Review Questions

The set below adds 20 questions covering Chapters 1-13 to the five questions of this chapter, in the format of the residency entrance examination (Romanian "rezidențiat"): single answer (SA), with only one correct answer, and multiple answer (MA), with 2-4 correct answers. Each answer can be found in the text of the chapter given in parentheses. The answers are on the last line of the chapter and in Appendix A.

Common Pitfalls

- Using the original AI Act deadlines for high-risk systems. Under Regulation 2026/1744, AI-enabled medical devices fall under the high-risk regime from August 2, 2028.
- Believing that the EHDS has applied in full since 2025. The regulation has been in force since March 26, 2025, but its obligations apply in stages, in 2027, 2029, and 2031.
- Assuming that Romania has adopted ICD-11. As of September 27, 2026, DRG grouping is still based on ICD-10-AM, 3rd edition.
- Considering a case entered into a chatbot "with initials only" to be anonymous.
- Treating a reference proposed by AI as verified because it has a PMID. The PMID may be invented or may belong to a different article.
- Confusing the performance of a model tested on its own with the performance of the physician or patient who uses it.
- Learning the numbers in this book as fixed truths. Standards, platforms, and deadlines change from one year to the next.

Key Points

- The discipline has five layers: data, standards, systems, clinical use, and governance. Each layer depends on those beneath it.
- The competencies of the first year are those of a user: searching, coding, data protection, security, measurement, performance calculation, and verification of AI output.
- AI Act: prohibitions and AI literacy from February 2, 2025; high-risk obligations from December 2, 2027 (Annex III) and August 2, 2028 (medical devices).
- EHDS: general application from March 26, 2027; patient summary and e-prescription from 2029; imaging and laboratory results from 2031.
- NIS2 and OUG 155/2024: a significant incident is reported within 24 hours, 72 hours, and one month.
- Before using AI: no identifiable data, permission, verification, a decision of your own, and attention to equity, followed by disclosure.

Self-Assessment Questions

1. (SA) Under Regulation (EU) 2026/1744, the obligations for high-risk AI systems integrated into products covered by Annex I, including medical devices, apply from:

a) February 2, 2025
b) August 2, 2026
c) December 2, 2027
d) August 2, 2028
e) March 26, 2031

2. (SA) The AI literacy obligation (Article 4 of the AI Act) applies from:

a) August 1, 2024
b) February 2, 2025
c) August 2, 2025
d) July 27, 2026
e) December 2, 2027

3. (SA) As of September 27, 2026, the DRG grouping of hospitalized cases in Romania is based on diagnoses coded in:

a) ICD-11 MMS, release 2026-01
b) ICD-10-AM, 3rd edition
c) ICD-10-CM
d) SNOMED CT
e) ICD-9-CM

4. (MA) Which of the following rules on the use of AI by students and physicians are supported in this chapter?

a) data that could identify a patient are not entered into a public chatbot
b) any reference proposed by a language model is checked in PubMed or CrossRef
c) a chatbot can be a coauthor if it wrote a large part of the text
d) the use of AI is disclosed, stating the tool, the version, and the purpose
e) the performance of a tool is also checked across patient subgroups

5. (MA) Which of the following competencies correspond to the DigComp area "safety," as applied in Table 1.2 and Table 14.1?

a) enabling two-step verification
b) choosing a long password, following the NIST recommendations
c) building a PubMed query with MeSH terms
d) applying the GDPR rules to patient data
e) calculating SDNN from a series of RR intervals

Review Questions

6. (SA) (Ch. 1) The medical informatics course became compulsory for all medical students in Romania in:

a) 1977
b) 1984
c) 1990
d) 1994
e) 2009

7. (SA) (Ch. 2) A 10-second, 12-lead ECG, sampled at 500 Hz and stored at 16 bits, takes up approximately:

a) 12 kB
b) 120 kB
c) 1.2 MB
d) 12 MB
e) 250 MB

8. (MA) (Ch. 3) Regarding keys and normalization, which of the following statements are correct?

a) the primary key uniquely identifies each row of a table
b) a cell containing the value "I10, E11.9" complies with the first normal form
c) a foreign key must point to an existing row in another table
d) the patient's name is a suitable primary key
e) in the third normal form, the year of birth is moved from the Consultation table to the Patient table

9. (SA) (Ch. 4) In ICD-11, the separator that links two stem codes in a cluster is:

a) &
b) /
c) †
d) *
e) the dot

10. (MA) (Ch. 4, 12) Regarding coding and DRG in Romania as of September 27, 2026, which of the following statements are correct?

a) the diagnoses of hospitalized patients are coded in ICD-10-AM, 3rd edition
b) the DRG classification in force is RO DRG v.1.1, dating from March 2024
c) INMSS collects the patient-level minimum data set from hospitals
d) ICD-11 officially replaced ICD-10-AM in 2022
e) procedures are coded in ICD-11

11. (MA) (Ch. 5) Under the NIS2 Directive and OUG No. 155/2024, a significant cybersecurity incident is reported through:

a) an early warning within 24 hours
b) a notification within 72 hours
c) a final report within one month at most
d) reporting only if personal data are affected
e) reporting exclusively to ANSPDCP

12. (SA) (Ch. 6) A signal has useful components up to 150 Hz. According to the sampling theorem, the sampling frequency must be:

a) at least 75 Hz
b) exactly 150 Hz
c) above 300 Hz
d) 256 Hz, regardless of the signal
e) below 50 Hz, to avoid the power-line frequency

13. (SA) (Ch. 7) The lung window, with level L = -600 HU and width W = 1500 HU, displays over the entire gray scale the interval:

a) from -1350 to +150 HU
b) from -600 to +900 HU
c) from -1500 to 0 HU
d) from -160 to +240 HU
e) from -2100 to +900 HU

14. (MA) (Ch. 7) Regarding the DICOM standard, which of the following statements are correct?

a) a study corresponds approximately to a requested imaging procedure
b) each instance has a unique identifier (UID)
c) exporting as JPEG preserves the Pixel Spacing calibration
d) the C-FIND service searches for objects by criteria
e) the display window changes the HU values in the file

15. (SA) (Ch. 8) The HL7 FHIR version on which most current implementations are based, and which is used by the IPS implementation guide and by the public test server in Practical Session 8, is:

a) R2
b) R3
c) R4
d) R6
e) R7

16. (MA) (Ch. 8, 10) Regarding Regulation (EU) 2025/327 on the EHDS, which of the following statements are correct?

a) it applies, in general, from March 26, 2027
b) the obligations for the patient summary, e-prescription, and e-dispensation apply from March 26, 2029
c) medical imaging, laboratory results, and discharge reports are added from March 26, 2031
d) it must be transposed in Romania through a national law
e) EHR systems must include a European logging component

17. (SA) (Ch. 9) The methodological norms for telemedicine in Romania were approved by:

a) Order No. 1133/2021
b) Government Decision (HG) No. 1133/2022
c) Law No. 45/2019
d) OUG No. 155/2024
e) Law No. 214/2024

18. (MA) (Ch. 9) Regarding the evidence on telemedicine, which of the following statements are correct?

a) telemonitoring in heart failure was associated with a reduction in all-cause mortality (OR 0.84)
b) remote consultation used alone brought no benefit in heart failure, in a meta-analysis on cardiovascular disease
c) blood pressure telemonitoring reduced systolic blood pressure by tens of mmHg
d) telestroke had outcomes similar to those of conventional treatment, in a meta-analysis of 12,540 patients
e) telepsychiatry was superior to face-to-face consultation in all diagnoses

19. (SA) (Ch. 10) Under Rule 11 of the MDR, software that provides information used to take diagnostic or therapeutic decisions, without the risks that would place it in a higher class, is in class:

a) I
b) IIa
c) IIb
d) III
e) it is not a medical device

20. (MA) (Ch. 10, 13) Regarding large language models, the data presented in this book show that:

a) participants who used an LLM identified the relevant condition in fewer than 34.5% of the scenarios
b) 18% of the references generated by GPT-4 did not exist, in a study of 636 references
c) in automatically generated clinical notes, the hallucination rate was 1.47%
d) popular chatbots frequently refused to answer questions on topics prone to misinformation
e) scores on licensing examinations guarantee safety in consultation

21. (SA) (Ch. 11) Under eIDAS, the legal effect equivalent to that of a handwritten signature is conferred by:

a) the simple electronic signature
b) the advanced electronic signature
c) the qualified electronic signature
d) the scanned image of the signature
e) the name typed at the end of an e-mail

22. (MA) (Ch. 11) Regarding cryptography and the digital signature, which of the following statements are correct?

a) AES is a symmetric encryption algorithm
b) a hash function uses no key and does not allow the document to be reconstructed
c) when signing, the document's digest is signed with the private key
d) verifying the signature requires the signer's private key
e) TLS protects the data stored on the server

23. (SA) (Ch. 12) A hospital has 10 applications that communicate through an integration engine, in a star architecture. The number of connections required is:

a) 10
b) 20
c) 45
d) 90
e) 100

24. (SA) (Ch. 12) The case-mix index of a hospital is:

a) the mean of the relative weights of the cases discharged in a given period
b) the number of hospital days divided by the number of discharges
c) the number of beds per 100,000 inhabitants
d) the proportion of readmissions within 30 days
e) the number of DRGs used by the hospital

25. (MA) (Ch. 13) An AI tool has a sensitivity of 0.80 and a specificity of 0.95. Which of the following statements are correct?

a) the positive predictive value decreases when the prevalence in the tested population decreases
b) sensitivity and specificity do not, in principle, depend on prevalence
c) an area under the ROC curve of 0.90 guarantees a positive predictive value of at least 0.90
d) accuracy can be misleading when the disease is rare
e) the F1 score takes true negatives into account

PART II

Part II. Practical Sessions

PRACTICAL SESSION 1

Practical Session 1. The Computer, the Operating System, and the Collaborative Workspace (Google Workspace)

Objective

By the end of the session, the student organizes his or her course files in a folder structure with consistent names, shares a document with minimal rights, uses the version history, turns on two-step verification on his or her own account, and explains why patient data have no place in a personal cloud account. All the patient examples in this session are fictitious.

Tools

- The laboratory computer, running Windows 11, or a personal computer (Windows, macOS, ChromeOS, or a Linux distribution).
- An up-to-date browser (Chrome, Edge, or Firefox, the current version as of September 27, 2026).
- A personal Google account or the institutional account, if the faculty provides one. Google Drive, Docs, and Sheets are web applications with no displayed version number. The help pages cited were consulted on September 27, 2026 [530,531].
- Optionally, a trusted password manager [532].

L1.1 The Computer and Files: Basic Notions

A computer has two levels. The physical components (hardware) are the processor, which executes instructions, the working memory (RAM), which temporarily holds the data of open programs, the storage unit (SSD or HDD), where data remain even after shutdown, and the input and output devices. The programs (software) are divided into the operating system and applications. The operating system manages the physical resources, the files, and the users. Applications, such as the word processor or the spreadsheet program, run on top of it. The operating system also separates user accounts, so that each user sees only his or her own files, and it receives updates that fix, among other things, security vulnerabilities. A shared computer, such as the one in the laboratory, therefore calls for two simple habits: signing out of the account when leaving and not saving passwords in the browser of a computer used by several people.

Files are organized in a hierarchical system of folders. Each file has a name and, as a rule, an extension that indicates the format: `.docx`, `.pdf`, `.csv`. The format determines which program can open the file and how easily it can be read ten years from now. In medicine, the choice of format has practical consequences. A radiological image exported as JPEG loses the metadata that the DICOM format preserves, and a study table saved only in the format of a commercial program may become hard to open after a change of license. The DICOM standard was created for images and waveforms, and HL7 FHIR describes clinical data as XML or JSON resources [15]. Multichannel

physiological signals, such as EEG or polysomnography, are frequently exchanged in the EDF or EDF+ format [221]. Table L1.1 summarizes the formats that the student will encounter in this book.

Table L1.1. File formats encountered in medical practice and in the practical sessions. Source: [15,221].

Format	Extension	Typical use	Note
Text document	.docx, .odt	letters, reports, assignments	editable, sent as PDF
Final document	.pdf	results, signed documents	fixed layout on any device
Spreadsheet	.xlsx, .ods	study data, calculations	keeps formulas and formatting
Delimited text	.csv	data exchange between programs	plain text, no formulas
General image	.png, .jpg	figures, photographs	JPEG compression is lossy
Medical imaging	.dcm (DICOM)	radiology, CT, MRI	includes patient metadata
Clinical data	.json, .xml (HL7 FHIR)	exchange between information systems	structured as resources
Signals	.edf (EDF, EDF+)	EEG, polysomnography, ECG	multichannel, open format

File names should be understandable without opening the file. The rule proposed in this session is `YYYY-MM-DD_activity_author_vNN.extension`, for example `2026-10-06_PS01_Popescu_v02.docx`. Writing the date in year-month-day order makes the files line up chronologically when they are sorted by name. Spaces and diacritics in names can cause problems when files are transferred between systems, so they are avoided. The version is written at the end, with two digits. Figure L1.1 shows a folder structure for the first semester.

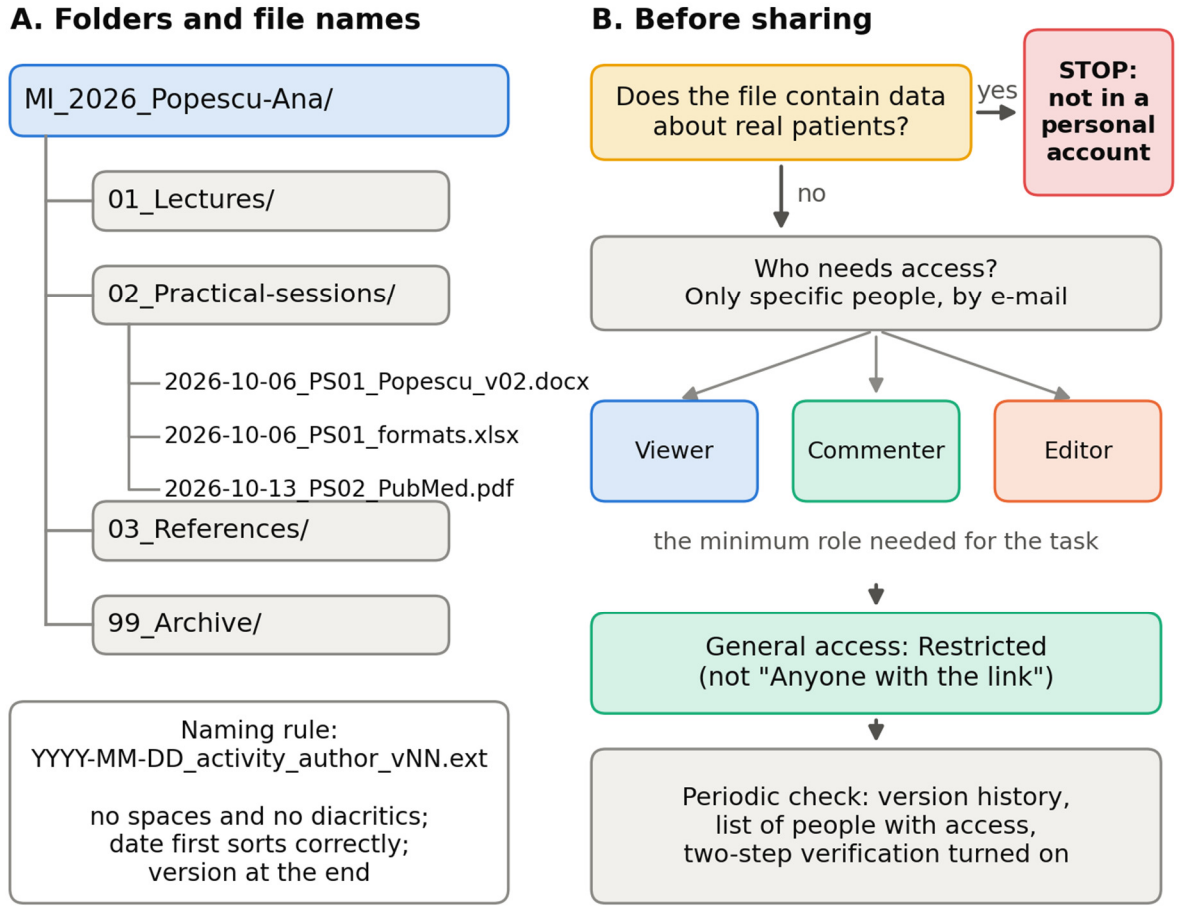


Figure L1.1. Organization of the course files (A) and the sharing decision (B). The roles and access options are those of Google Drive [530].

L1.2 Sharing, Versions, and Account Security

Google Drive offers four roles. The Viewer can view and download the file. The Commenter can, in addition, add comments. The Editor can modify and share the file but cannot change the owner. The Owner has all rights [530]. Separately from the people added by name, there is general access. With the "Restricted" option, only the people added can open the file. With "Anyone with the link", the file can be used by anyone who receives the link, without signing in [530]. The working principle is the minimum necessary role: someone who only reads gets the Viewer role.

The version history shows who modified the document and when. Editing rights are required to see earlier versions. A version can be given a name (at most 40 named versions in a document and 15 in a spreadsheet) and can be restored [531]. The same page warns that some revisions may be merged automatically, so the version history does not replace backups [531].

The account that hosts the files must be protected. Google recommends passwords of at least 12 characters, different for each important account and free of easily guessed personal information, and suggests using a trusted password manager [532]. The US guideline NIST SP 800-63B-4, published in 2025, goes further. Passwords used as the only authentication factor must be at least 15 characters long, and those used together with a second factor at least 8. The guideline prohibits imposed composition rules (of the "must contain a digit and a symbol" type) and forced periodic password changes, but it requires a change when there is evidence of compromise. Verifiers must allow password managers and autofill, and pasting the password from the clipboard is recommended [445]. A long passphrase that is easy to remember and used in only one place is therefore a good choice.

Two-step verification adds a second factor, which is useful if the password has been stolen. Google offers several options: phone prompts; passkeys and security keys, which also protect against phishing scams; authenticator apps with one-time codes; codes by SMS or voice call; and eight-digit backup codes, which are kept in a safe place [533].

Security is not a theoretical problem for hospitals. Between 2016 and 2021, 374 ransomware attacks on healthcare delivery organizations were reported in the United States, exposing the health data of almost 42 million patients. The annual number of attacks more than doubled, from 43 to 91, and 44.4% of the attacks disrupted patient care [183]. Over the same period, the affected systems were less and less often restored from backups [183]. A systematic review analyzed the human behaviors that open security breaches in healthcare institutions, including the vulnerability of staff to social engineering attacks [442].

Backups follow an easy-to-remember practical rule called 3-2-1: three copies of the data, on two different types of media, one of which is kept off-site. The US cybersecurity agency recommends keeping at least one copy of critical data offline and encrypted, and testing restoration regularly, because many ransomware variants search for and delete or encrypt the backups that are accessible from the network [534]. For the student, the minimum option is this: the files in Drive, a copy on the computer, and a copy on an external drive kept disconnected. Automatic synchronization is not a substitute for a backup. If a file is accidentally deleted, or encrypted by malicious software on the computer, the change may propagate to the cloud, and the synchronized copy becomes as unusable as the original.

The last rule is the most important one for a future physician. The General Data Protection Regulation (GDPR) defines data concerning health as personal data related to the physical or

mental health of a natural person, including the provision of healthcare services. They belong to the special categories of data whose processing is prohibited outside the exceptions provided for in Article 9 [118]. The controller, for example the hospital, must implement appropriate technical and organizational measures (Article 32) and may work with a service provider only on the basis of a contract that obliges the provider to protect the data (Article 28) [118]. The personal account of a student or of a physician does not fall under this framework. This is why photographs of medical documents, patient lists, or images of the ward monitor are never uploaded to a personal cloud, even if the folder is "private." Chapter 5 details these rules.

Procedure

1. **Computer components.** On the laboratory computer, open Settings > System > About (Windows 11) or "About This Mac" (macOS). Record in the working document the processor type, the RAM, the storage capacity, and the operating system version.
2. **File extensions.** In File Explorer, turn on the display of extensions (the View > Show menu, the option for file name extensions). Identify three files with different extensions and the program associated with each.
3. **Folder structure.** In Google Drive, create the folder `MI_2026_Surname-Firstname` and, inside it, the subfolders `01_Lectures`, `02_Practical-sessions`, `03_References`, and `99_Archive`, as in Figure L1.1.
4. **The first document.** In `02_Practical-sessions`, create a Google Docs document named according to the rule in Section L1.1, for example `2026-10-06_PS01_Popescu_v01`. Enter the answers to steps 1 and 2 in it.
5. **The spreadsheet.** In the same folder, create a Google Sheets spreadsheet containing Table L1.1, supplemented with two more formats found by the student (for example, the format of a phone app for ECG). Record the source of the information for each format.
6. **Sharing with a minimal role.** Share the document with a classmate as Commenter and with the instructor as Viewer, using the Share button. General access remains Restricted. The student then checks, in a private browser window and without signing in, that the link does not open the document.
7. **Comment and reply.** The classmate adds a comment to a paragraph, and the author replies and marks the comment as resolved.
8. **Version history.** The author makes three successive changes, opens the version history (the button at the top right that shows the last edit), names the current version "after comments", then restores an earlier version and returns to the latest one [531]. Record what information is shown for each version.
9. **Two-step verification.** On his or her own account, in the security section of the Google Account, the student checks whether two-step verification is on and, if it is not, turns it on, choosing a method other than SMS when possible [533]. The backup codes are kept off the computer. No password or code is written in the working document.
10. **The password.** The student compares the password rules of a service that he or she uses with the NIST recommendations (length, no composition rules, no forced periodic change) and notes two differences [445].

11. **The backup plan.** Describe, in five lines, the personal 3-2-1 plan for the course folder: where the three copies are, on what media, and how often they are updated. Download a copy of the document in PDF format and one in DOCX format to the local archive folder.
12. **GDPR scenario (fictitious).** A resident wants to photograph the observation chart of the fictitious patient "I.P., 67 years old" in order to discuss it in the evening with a colleague, via Drive. The student writes down three reasons why this is not allowed and proposes a correct alternative (for example, a discussion in the hospital, on the institution's system, or a fully anonymized case, with no identifiable details).

What to Submit

Submit the link to the folder `MI_2026_Surname-Firstname`, shared with the instructor as Viewer, with general access set to Restricted. The folder contains the working document and the spreadsheet. Screenshots showing personal account data are not submitted. Indicative grading scheme, 10 points:

- folder structure and file names that follow the rule (2 p).
- answers to steps 1 and 2 and the completed table of formats, with sources (2 p).
- correct sharing: minimal roles, Restricted access, check in the private window (2 p).
- version history: a named version, restoration described (1 p).
- comparison of the password rules with NIST and confirmation of two-step verification, without codes (1 p).
- the backup plan and the reasoned GDPR scenario (2 p).

In the practical examination, the student may be asked to redo steps 3, 6, and 8 on a new document within 10 minutes.

Offline Variant

If the internet or the accounts are not available, the session is carried out locally. The folder structure is created in File Explorer, on the computer's disk, and the documents are written in LibreOffice Writer and Calc or in Microsoft Word and Excel. Versions are kept manually, by saving the files successively as `_v01`, `_v02`, `_v03`, and at step 8 the student compares two versions with the document comparison function of the word processor. Sharing is simulated on paper: for three people (a classmate, the instructor, a fictitious patient), a table is filled in with the role granted and the reason. The backup is made on a USB stick, and steps 9 and 10 are discussed on the basis of Section L1.2. The folder is submitted on the stick or uploaded later, when the network is available again.

PRACTICAL SESSION 2

Practical Session 2. Medical Information on the Web: Searching and Evaluating Quality

Objective

The student turns a clinical question into a structured PubMed search (PICO, MeSH, operators, filters), saves and exports the results, and then critically evaluates three types of sources: a Wikipedia article, a web page for patients, and a chatbot's answer, compared with a practice guideline. The clinical question and the patient in the examples are fictitious.

Tools

- PubMed, the free search engine of the US National Library of Medicine, and its user guide, consulted on September 27, 2026 [39,41].
- The MeSH Browser, for checking descriptors and synonyms [535].
- Wikipedia, in Romanian and in English.
- A general-purpose chatbot available free of charge at the time of the session. Record its name and the version displayed by the application.
- Google Sheets or another spreadsheet program, for the checklist in Table L2.1.

L2.1 From the Clinical Question to the Query

A good search starts with a well-formulated question. The PICO format breaks the question down into Population, Intervention, Comparator, and Outcome. It is useful as a starting point, but it does not by itself guarantee an efficient search. A 2021 analysis shows that the evidence on its effectiveness in teaching is limited and recommends that it not be used as an exclusive tool, but together with a critical appraisal of the whole process [536].

The example used in this session is fictitious: "In adults with arterial hypertension, does remote blood pressure monitoring, compared with usual care, lower blood pressure values?" Each PICO element corresponds to a search concept. For each concept, the appropriate descriptor is looked up in the MeSH Browser: Hypertension, Telemedicine, Blood Pressure [535]. MeSH has a hierarchical structure, and a search by descriptor automatically includes the narrower terms, through explosion. The synonyms (entry terms) show what else the concept is called in the literature [13].

Very recent articles may not yet be indexed with MeSH. For this reason, each concept is also expressed through words from the title and abstract, linked to the descriptor with OR. The concepts are then linked to one another with AND. Operators are written in capital letters, parentheses set the order of operations, the asterisk truncates the term (after at least four letters), and quotation marks require an exact phrase [41]. For the example above, the query becomes:

```
("Hypertension"[Mesh] OR hypertens*[tiab]) AND ("Telemedicine"[Mesh] OR telemonitor*[tiab] OR "home blood pressure"[tiab]) AND ("Blood Pressure"[Mesh] OR "blood pressure"[tiab])
```

On September 27, 2026, the first concept alone returned 669,982 records, the second 63,552, the third 546,995, and the combination of all three yielded 3,010 [39]. The filter for systematic reviews and meta-analyses reduced the result to 131 records, and the additional restriction to the years 2021-2026 brought it down to 73. There were 123 randomized controlled trials from 2021-2026 [39]. The numbers will change over time, so each student records his or her own values and the date of the search.

Why is a simple Google-type search not enough? General search engines use partial-match searching: they accept a few words and rank documents by how well they match them. Bibliographic databases also allow exact-match searching, with sets combined through Boolean operators and with attributes such as publication type [13]. Even when it receives an ordinary phrase, PubMed tries to match the words in turn with MeSH descriptors, journal titles, common phrases, and author names, and it searches whatever remains unmatched as free text [13]. The exact translation of the query can be seen on the Advanced page, in the search history and details section. It is worth checking every time, because it shows what the system actually understood. The Clinical Queries function applies validated strategies for questions of therapy, diagnosis, harm, and prognosis on top of the user's terms [13]. In the current version of PubMed, systematic reviews are no longer part of Clinical Queries and are searched for with the article type filter [41].

What happens to the results? The search history on the Advanced page keeps the last 100 queries and expires after 8 hours of inactivity. Saving a search permanently and receiving e-mail alerts require a My NCBI account. Citations are exported with Send to > Citation manager into a `.nbib` file (at most 10,000 citations), which is compatible with many reference management programs [41]. The default sort order, Best Match, uses a relevance algorithm, and sorting by date remains available [41].

The results do not all carry the same weight. The hierarchy of evidence places systematic reviews and meta-analyses above randomized trials, and randomized trials above observational studies, case series, and expert opinion. The hierarchy guides reading but does not exempt the reader from critically appraising each article [42]. Figure L2.1 summarizes the entire workflow.

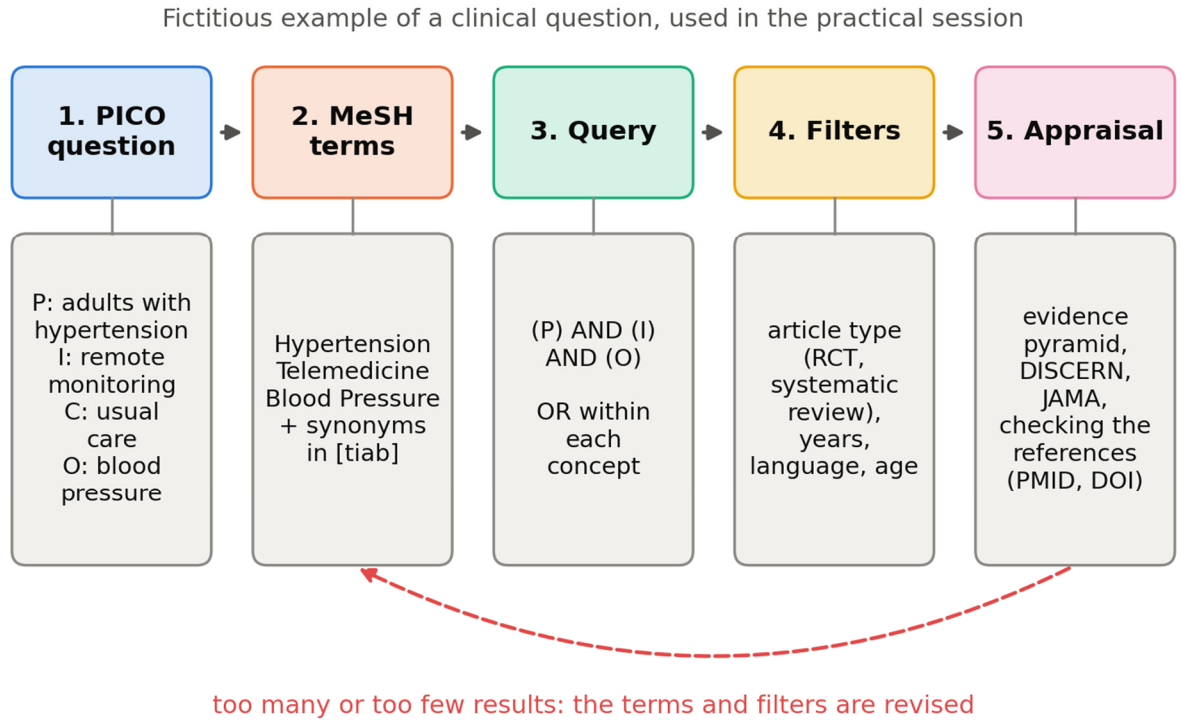


Figure L2.1. The workflow from the clinical question to the evaluation of sources. Fictitious example. Syntax according to [41].

L2.2 Wikipedia: A Starting Point, Not a Source to Cite

Wikipedia is one of the most visited websites in the world. The English-language version contains more than 45,000 medical articles, which are accessed almost 1.6 billion times a year [537]. Their quality varies. A study of the 1,000 most-read health articles on English-language Wikipedia found comparable quality between articles on conditions specific to women and those on conditions specific to men [537]. Differences between languages, however, can be large. The article on stroke existed in 23 of the 30 most widely spoken languages, its content varied greatly from one language to another, and three versions presented corticosteroids as a treatment for stroke. The quality score was positively correlated with the number of references in the article [538].

The practical conclusion: Wikipedia can be used for orientation, for terminology, and for the list of references at the end. The references are opened and checked. In university assignments, the primary source or the guideline is cited, not the encyclopedia.

L2.3 Tools for Evaluating Information for Patients

Two tools appear frequently in studies on the quality of online medical information. DISCERN assesses the quality of health information, with a total score that can reach 80 points. The JAMA criteria check four elements: authorship, attribution of sources, disclosure of conflicts of interest, and the date of the last update [539,540].

The results of applying them are, as a rule, modest. In a study of 65 Arabic-language web pages on dental implants, only 3 pages (4.5%) obtained a high DISCERN score, and only 5 (7.6%) met all the JAMA criteria. The most frequent gaps were the attribution of sources and the author's name [539]. The same type of evaluation can be applied to a chatbot's answers. For 24 questions formulated on the basis of the European guideline on allergen immunotherapy, the answers of ChatGPT-4 were of "acceptable" quality according to DISCERN but had JAMA scores of 0-1, because the author, the sources, the disclosures, and the date were missing. The texts required a university reading level [540].

L2.4 A Chatbot's Answer Compared with the Guideline

Chatbots answer fluently, but their answers must be checked against a reference source. A 2025 study turned 28 recommendations from an international guideline on hereditary angioedema into questions addressed to four models (ChatGPT, Gemini, Perplexity, Copilot). Two clinicians rated the accuracy, appropriateness, clarity, and reliability of the references. ChatGPT and Gemini obtained median scores of 5 out of 5 for accuracy and appropriateness, compared with 3 for the other two models. When the models were explicitly asked to indicate their sources, the consistency of the answers improved [541].

For such a comparison to be reproducible, the way of working must be documented. The CHART reporting guideline, published in 2025 for studies that evaluate health advice provided by chatbots, has 12 items and 39 subitems. They include the identification of the model, the wording of the instructions (prompt), the query strategy, and the method used to evaluate performance [542]. This practical session adopts the same logic in simplified form: the model, the date, the exact text of the question, and the complete answer are recorded.

Table L2.1. Checklist for a source of medical information found on the web or produced by a chatbot. Source: adapted from [42,539,540,542].

Criterion	Verification question	Yes / No / Partly
Author	Is the author named, and does he or she have relevant qualifications?	
Sources	Are verifiable sources cited (PMID, DOI, guideline)?	
Real references	Does each reference exist and say what is attributed to it?	
Timeliness	Is the date displayed, and is it at most 5 years old?	
Conflicts of interest	Are funding and commercial interests disclosed?	
Level of evidence	Does it rely on systematic reviews, RCTs, or guidelines?	
Agreement with the guideline	Does the recommendation match the guideline in force?	
Balance	Are both the risks and the alternatives presented?	
Documentation (chatbot)	Are the model, the date, and the text of the question recorded?	

Procedure

1. **The question.** Each student receives or chooses a fictitious clinical question and writes it in PICO format, in a document named according to the rule in Practical Session 1.
2. **MeSH descriptors.** For each PICO element, look up the descriptor in the MeSH Browser. Record the descriptor, two synonyms (entry terms), and its position in the tree.
3. **Searching by concept.** In PubMed, search for each concept separately, as a MeSH descriptor OR terms in [tiab]. Record the number of results.
4. **Combining.** In Advanced > History, combine the concepts with AND, using the set numbers (for example `#1 AND #2 AND #3`). Record the number of results and the date.
5. **Filters.** Apply, in turn, the filter for systematic reviews and meta-analyses, then the one for randomized trials, then the interval of the last 5 years. Record the effect of each filter.
6. **Saving and exporting.** Students who have a My NCBI account save the search. All students export the first 10 relevant results through Send to > Citation manager and import the `.nbib` file into a reference management program or, if they do not have one, open the file in a text editor and identify the PMID, TI, and AU fields.
7. **A key article.** Choose a recent systematic review from the results. Record the type of study, the year, the number of included studies, and the main conclusion, as read in the abstract.
8. **Wikipedia.** Open the article on the disease in the question in Romanian and in English. Compare their length, the date of the last modification, and the number of references. Check whether two references from the article exist in PubMed.
9. **A page for patients.** Choose the first Romanian-language page returned by a general search engine for the disease in question and apply the checklist in Table L2.1.

10. **The chatbot.** Put the PICO question to the chatbot, without any real patient data, with an explicit request to indicate its sources. Copy the answer in full, and record the model, the date, and the text of the question [542].
11. **Checking against the guideline.** Compare the chatbot's recommendation with the conclusion of the systematic review selected in step 7 and with a practice guideline indicated by the instructor. Look up each reference given by the chatbot in PubMed or by its DOI and classify it as existing and correct, existing but incorrectly described, or nonexistent.
12. **Synthesis.** The student writes a paragraph of at most 150 words about the source he or she would recommend to a patient and the source he or she would use in an assignment, with arguments drawn from the checklist.

What to Submit

Submit a single document (Google Docs or PDF) containing: the PICO question, the table of MeSH descriptors, the final query copied exactly from PubMed, the number of results at each step, with the date, the `.nbib` file or the list of the 10 citations, the checklist completed for the web page and for the chatbot, the result of the reference check, and the synthesis paragraph. Indicative grading scheme, 10 points:

- a clear PICO question and correct MeSH descriptors (2 p).
- a syntactically correct query, with OR within concepts and AND between them (2 p).
- filters applied and documented, export completed (1 p).
- evaluation of Wikipedia and of the web page with the checklist (2 p).
- the chatbot's answer documented, compared with the guideline, and its references checked one by one (2 p).
- a reasoned synthesis (1 p).

In the practical examination, the student may be asked to build a query for a new question and to classify three given references.

Offline Variant

Without internet access, the session is based on materials prepared by the instructor. Students receive, on paper or on a USB stick, a PICO question, an excerpt from MeSH with descriptors and synonyms, and a list of 20 citations previously exported from PubMed, with the publication type. They build the query on paper and determine which citations each filter would keep. For the evaluation, they receive printouts of a Wikipedia article, a page for patients, and a previously saved chatbot answer, together with the reference guideline, and they apply the checklist in Table L2.1. The references are checked later, when the network is available.

PRACTICAL SESSION 3

Practical Session 3. Electronic Data Collection Forms: Google Forms and Epi Info

Objective

The student designs, tests, and uses an electronic data collection form for a fictitious study, then turns the responses into a descriptive table. By the end of the session, the student should be able to:

- write the data dictionary of a small study, with the type, unit, and permitted values of each variable
- configure validation rules in Google Forms and make the questionnaire branch according to the answer given
- run a pilot test and correct the form based on its results
- export the responses to Google Sheets, clean the data, and document every correction
- build a simple descriptive table with a pivot table
- explain why the form contains no names, personal numeric codes (codul numeric personal, CNP), or e-mail addresses

The session applies the concepts of Chapter 3 (the data dictionary, validation at entry, data quality, electronic data capture tools). The spreadsheet part continues the tradition of the discipline's practical sessions: the protocol in [6] is taken up again and updated for Google Sheets 2026, and the exercises from which it started can be found in the 2010 volume [5]. What is new is the direct collection of data through an electronic form, with validation rules set before data collection begins.

The fictitious study. "Blood pressure screening in first-year students" is a teaching scenario. The data are generated through role play among classmates (values are invented within plausible limits, or measurements taken in the practical classroom are used, with no link to the person's identity). No value is interpreted clinically. For context, the WHO defines hypertension as a systolic blood pressure of at least 140 mmHg and/or a diastolic blood pressure of at least 90 mmHg, measured on two different days [543]. A single screening session can therefore flag only a "value elevated at screening," not a diagnosis.

Tools

All tools are free of charge. The web addresses were checked on September 27, 2026.

- **Google Forms** (https://forms.google.com), a web application, current version at the time of access. It requires a Google account. The official documentation describes response validation rules [123], sections and branching to a section based on the answer [124], linking the form to a spreadsheet [125], and downloading the responses in CSV format [128].

- **Google Sheets** (https://sheets.google.com), a web application, for data cleaning and the pivot table [544]. Microsoft Excel or LibreOffice Calc can replace Google Sheets for the CSV file.
- **Epi Info 7 for Windows** (https://www.cdc.gov/epiinfo/), an alternative that needs no online account. The CDC states that development and technical support were discontinued as of September 2025 and that the desktop version was still working on Windows 11 at that date [127]. Here the program is used for teaching purposes only.

Procedure

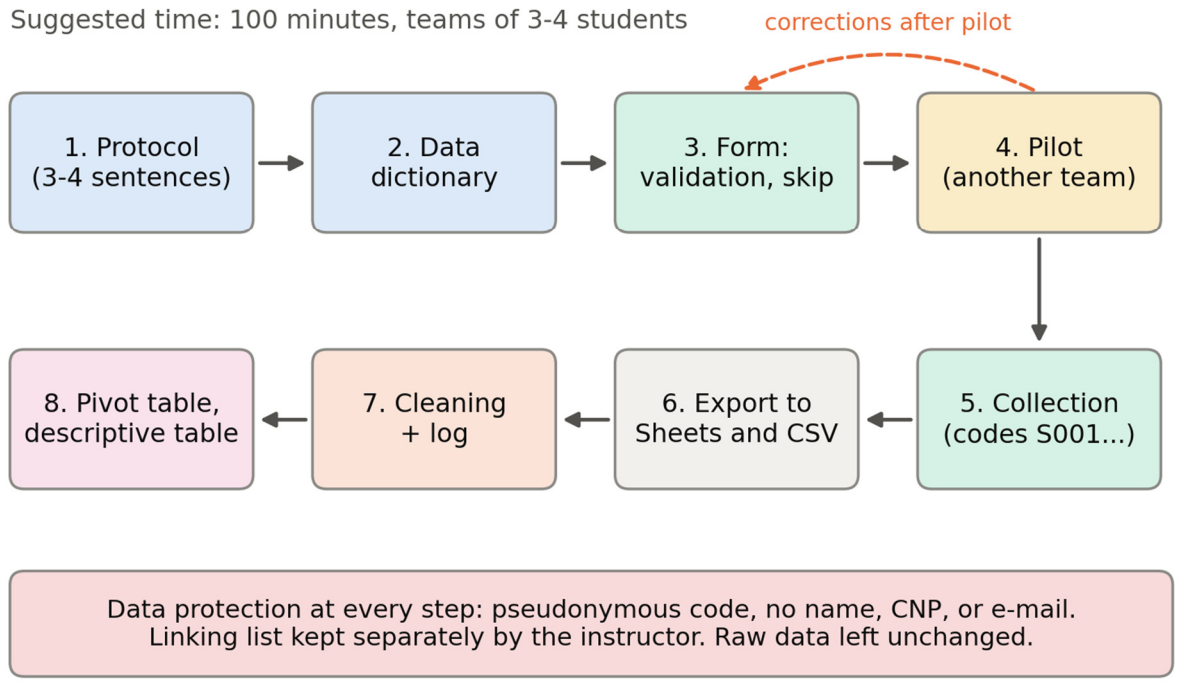


Figure L3.1. Workflow of Practical Session 3: from the protocol and the data dictionary to the form, pilot, collection, export, cleaning, and the descriptive table. Data protection (pseudonymization, no direct identifiers) applies at every step. Original diagram.

Figure L3.1 summarizes the steps. The suggested time is 100 minutes: 25 for the protocol and the data dictionary, 30 for the form and the pilot, 15 for data collection, and 30 for export, cleaning, and analysis. Students work in teams of 3-4.

L3.1 The Protocol and the Data Dictionary

1. In 3-4 sentences, the team writes the protocol of the fictitious study: the population (first-year students aged 18 years or older), the main variable (blood pressure), the secondary variables (age, sex, height, weight, smoking, physical activity, family history of hypertension), and the calculation rule. The fictitious protocol of the session specifies three readings in a seated position, 1-2 minutes apart, and the use of the mean of the last two.
2. Each participant receives a pseudonymous code of the form S001, S002, and so on, from a separate list kept by the instructor. The list linking codes to identities is never entered into the form.
3. The team completes the data dictionary following the model in Table L3.1. For each variable, the team sets the short name (no spaces and no diacritics, for compatibility with statistical

software), the full label, the type, the unit, the permitted values, and the code for a missing value.

Table L3.1. Excerpt from the data dictionary of the fictitious study and the corresponding rule in Google Forms. Fictitious data.

Variable	Label and unit	Type	Permitted values	Rule in Google Forms
code	Participant pseudonymous code	Text	S followed by 3 digits	Regular expression: ^S[0-9]{3}$
age	Age (completed years)	Integer numeric	18-40	Number between 18 and 40
sex	Sex	Nominal	F, M, unspecified	Multiple choice
smoker	Currently smokes	Nominal	yes, no	Multiple choice with skip
sbp2, sbp3	Systolic BP, readings 2 and 3 (mmHg)	Integer numeric	70-250	Number between 70 and 250
dbp2, dbp3	Diastolic BP, readings 2 and 3 (mmHg)	Integer numeric	40-150	Number between 40 and 150
phys_act	Days per week with 30 min of physical activity	Ordinal	0-7	Drop-down list 0-7

The ranges in the table are plausibility rules for the exercise, not clinical thresholds. They must be wide enough not to reject unusual but real values and narrow enough to catch typing errors (for example, 1280 instead of 128).

L3.2 Building the Form in Google Forms

1. Create a new form with the title of the study and a short description: the teaching purpose, the fact that the data are fictitious, and the fact that no identifying data are requested.
2. In the settings, under the Responses section, the collection of e-mail addresses stays turned off. The documentation shows that a form can collect verified addresses or addresses entered by the respondent [128], options that have no place in an anonymous study.
3. The first question is a checkbox for consent to take part in the exercise. It is marked as required.
4. Add the questions in the order of the data dictionary. For each short-answer numeric question, turn on response validation with the rule "Number" and "Between" [123]. The custom error message states exactly what is expected ("Enter systolic BP in mmHg, between 70 and 250").
5. For the pseudonymous code, use validation by regular expression [123]. The expression `^S[0-9]{3}$` accepts only the letter S followed by exactly three digits.
6. **Branching.** The question "Do you currently smoke?" is placed at the end of a section. From the question menu, choose "Go to section based on answer." The answer "yes" leads to a section on the number of cigarettes per day, and the answer "no" leads directly to the measurement section [124]. Branching works only for multiple-choice and drop-down questions [124].
7. Nominal variables are entered as multiple-choice questions, not as free text. This way, "F," "f," "fem.," and "female" can no longer appear in the same file.
8. All questions that correspond to the main variable are marked as required.

L3.3 The Pilot Test

1. Each team sends the form link to another team, which fills it out 3-5 times, including with deliberately wrong values (an age of 17 years, a blood pressure of 1280, a code "s1").
2. The pilot team records the completion time, the ambiguous questions, and the errors that the form did not stop.
3. The authors correct the form. If a validation rule rejected a correct value, the range is widened and the change is recorded in the data dictionary.
4. The pilot responses are deleted before the actual data collection. A short note listing the changes made is kept.

L3.4 Collection and Export

1. Each student fills out the form for 3-5 fictitious "participants," using the codes received.
2. In the Responses tab, the form is linked to a new Google Sheets spreadsheet. The sheet automatically receives the responses in table form [125].
3. At the end, response collection is stopped, and a CSV copy of all responses is downloaded as an archive of the raw data [128]. The raw file is not modified afterward.

L3.5 Data Cleaning in Google Sheets

Cleaning is done on a copy of the response sheet, never on the raw data. Every correction is entered in a separate sheet, "Log," with the participant code, the variable, the old value, the new value, and the reason.

1. Rename the columns with the short names from the data dictionary (code, age, sex, and so on).
2. **Duplicates.** In a helper column, count how many times each code appears: `=COUNTIF(B:B, B2)`. Any value greater than 1 indicates a double entry, which is examined (as a rule, the first complete record is kept).
3. **Missing values.** For each column, count the empty cells with `COUNTBLANK`. The proportion of missing values is included in the report.
4. **Out-of-range values.** Although the form has validated the ranges, check the minimum and maximum of each numeric variable with `MIN` and `MAX`.
5. **Consistency between fields.** Check whether the diastolic pressure is lower than the systolic pressure for each reading: `=IF(G2<F2, "ok", "check")`.
6. **Derived variables.** Calculate the mean of readings 2 and 3 (`=AVERAGE(F2, H2)` for systolic pressure, and likewise for diastolic pressure), the body mass index (`=weight/(height/100)^2`), and the screening category: `=IF(OR(sbp_m>=140, dbp_m>=90), "elevated at screening", "below threshold")`.

The formulas above follow the English-locale convention, in which the decimal separator is the point and function arguments are separated by commas. In locales that use the decimal comma, such as the Romanian regional setting, Google Sheets separates arguments with semicolons, and the same formula is written, for example, as `=COUNTIF(B:B; B2)`.

L3.6 The Descriptive Table

1. Select the cleaned data and insert a pivot table [544]. Put the variable `sex` in the rows, `screening category` in the columns, and the number of participants (COUNTA on code) in the values.
2. In a second pivot table, display the mean and standard deviation of the mean systolic pressure by sex.
3. Copy the result into a final table, following the model in Table L3.2, with the total number of participants (N) in the title and the note "fictitious data."

Table L3.2. Model descriptive table for the report. The cells are filled in with the team's results. Fictitious data.

Characteristic	Female	Male	Total
Participants, n			
Age, mean (SD)			
Mean systolic BP, mmHg, mean (SD)			
Mean diastolic BP, mmHg, mean (SD)			
Value elevated at screening, n (%)			
Smokers, n (%)			

The team discussion starts from three questions. Which errors would have entered the data without the validation rules? How many corrections did the log record? Which variable is missing from the dictionary and would have been needed to interpret the results (for example, the size of the blood pressure cuff)?

L3.7 The Epi Info 7 Variant

Teams working on a Windows computer can rebuild the form in Epi Info 7. The CDC page lists the components of the program: Form Designer for the form, Enter Data for data entry, Classic Analysis and Visual Dashboard for analysis, plus StatCalc [127].

1. In Form Designer, create a new project and a form with the same variables as in the data dictionary.
2. In the Check Code editor, define the range rules, the required fields, and the skip over the question on the number of cigarettes for nonsmokers.
3. Enter 5 fictitious records in Enter Data, including two with deliberate errors, to check the rules.
4. In Classic Analysis or Visual Dashboard, obtain the frequency table by sex.
5. The team compares the two tools using the criteria in Table 3.3 (cost, validation, audit trail, export, long-term support).

Because the program is no longer being developed, the results obtained in Epi Info are not used outside the exercise.

L3.8 Data Protection

Even in an exercise, the rules apply just as they would in a real study. Data concerning health are personal data that reveal information about a person's health status, and pseudonymization requires that the information allowing re-identification be kept separately [118]. The processing of the personal numeric code is subject to additional conditions under Romanian law [176]. For this session, this means in practice that:

- the form contains no name, CNP, e-mail address, telephone number, or full date of birth
- the list linking the codes to persons stays with the instructor, separately from the data
- the spreadsheet is shared only with the team members and the instructor
- at the end, the form is closed to responses, and the spreadsheet is kept only until the assessment

For a real study with patient data, Google Forms is not the right tool. An institutional EDC system with an audit trail is used, together with approval from the ethics committee (Chapter 3).

What to Submit

Each team submits a single document (a PDF or a link with access for the instructor), which contains:

1. the 3-4-sentence protocol and the complete data dictionary (2 points)
2. the form link and screenshots of the validation rules and of the branching, with no personal data of users visible (2 points)
3. the note on the pilot test, with at least two changes made (1 point)
4. the "Log" sheet with the cleaning corrections (2 points)
5. the completed descriptive table, with N in the title and the note "fictitious data" (2 points)
6. a paragraph of 5-8 sentences comparing Google Forms with Epi Info 7 or with REDCap for a real clinical study (1 point)

The grading scheme is indicative, out of 10 points. In the practical examination, the student may be asked to add a new validation rule or to explain a correction in the log.

Offline Variant

If there is no internet connection, the form is simulated in a local spreadsheet (Microsoft Excel or LibreOffice Calc). The data dictionary is written on the first sheet. On the second sheet, each column receives a data validation rule (a list of values or a numeric range), the equivalent of the rules in Google Forms. Branching is replaced by a calculated column that flags incompatible answers (for example, a number of cigarettes entered for a nonsmoker). Epi Info 7, installed beforehand, also works offline. The cleaning steps, the log, and the pivot table remain the same.

Practical Session 4. Coding Diagnoses with the WHO Online Tools (ICD-10 and ICD-11)

Objective

By the end of the session, the student can find and verify an ICD-10 code in the WHO browser, build an ICD-11 code or cluster with the Coding Tool and the ICD-11 browser, order the codes of a cluster correctly for morbidity and for mortality, and explain why an ICD-10-ICD-11 mapping table does not replace manual coding. All cases in this session are fictitious.

Tools

All tools are free and run in the web browser. The web addresses were checked on September 27, 2026.

- **WHO ICD-10 browser, 2019 version**: https://icd.who.int/browse10/2019/en [133]. It contains the tabular list and the inclusion and exclusion notes.
- **WHO ICD-11 MMS browser, release 2026-01**: https://icd.who.int/browse/2026-01/mms/en [141]. The page of each entity displays the description, the exclusions, and the "Postcoordination" section, with the axes allowed for that code.
- **ICD-11 Coding Tool, MMS 2026-01**: https://icd.who.int/ct [545]. In ICD-11, the coding tool replaces the printed index [129].
- **ICD-11 Reference Guide**, for the rules of postcoordination and code ordering [134].
- **ICD-10-ICD-11 mapping tables**, available in the "Info" tab of the ICD-11 browser [135].
- A spreadsheet editor (Google Sheets or Excel) for the answer sheet.

No account is needed. The ICD API requires registration, but it is not used in this session [546].

Procedure

Duration: 100 minutes. Each student works individually and fills in the answer sheet (Table L4.1). Figure L4.1 shows the steps repeated for each case.

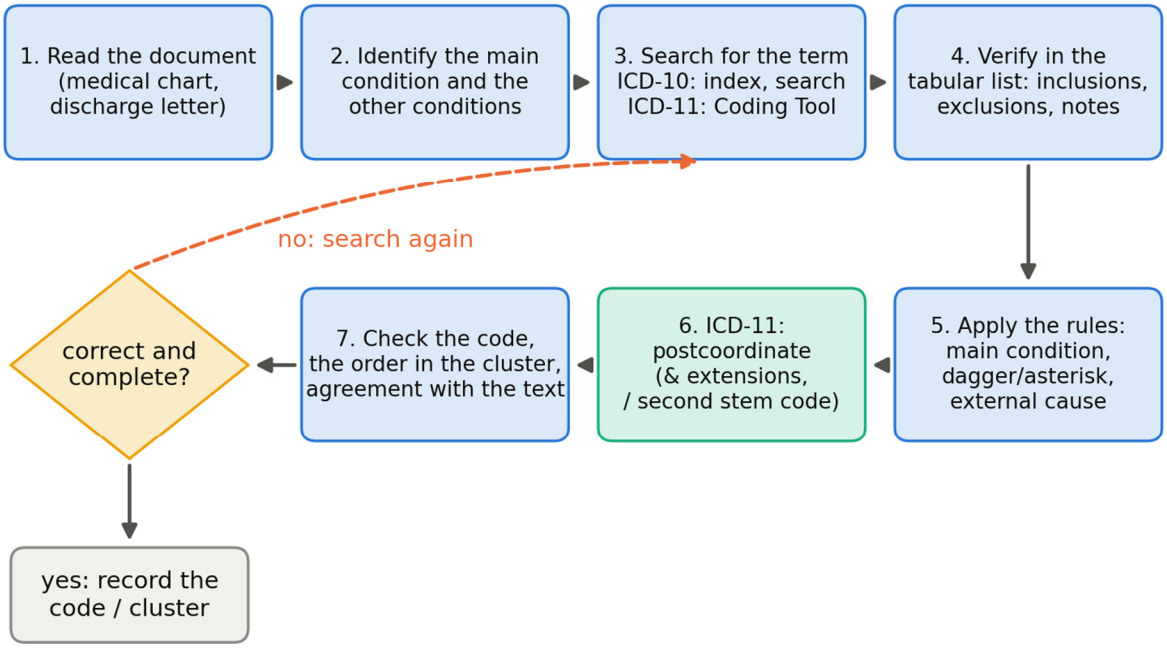


Figure L4.1. Workflow for coding a diagnosis in ICD-10 and ICD-11. Diagram based on the WHO rules [133,134].

L4.1 Preparation (5 minutes)

1. Open three browser tabs: the ICD-10 2019 browser, the ICD-11 2026-01 browser, and the Coding Tool.
2. Check that version 2026-01 appears in the header of the ICD-11 browser and of the Coding Tool. The interface language remains English, because the official code titles are verified in this language version. Codes change between versions, so results are always reported together with the version.
3. Create a spreadsheet with the columns of Table L4.1.

Table L4.1. Answer sheet (to be completed for each case).

No.	Main condition (text)	ICD-10 code	ICD-11 code or cluster	Where you verified it
1	...	...	...	tabular list / entity page

L4.2 The ICD-10 Tabular List (15 minutes)

1. In the ICD-10 browser, open Chapter IX, "Diseases of the circulatory system" (I00-I99), then block I20-I25 and category I21.
2. Read the note at the beginning of block I20-I25: for morbidity, the duration in categories I21, I22, I24, and I25 refers to the interval from the onset of the ischemic episode to admission [133]. Note on the sheet what "acute" means for I21 (the inclusion terms specify a duration of no more than 4 weeks, that is, 28 days, from onset).
3. Open category J13 and read the exclusions. Note where congenital pneumonia due to *S. pneumoniae* is coded.

4. Open G63.2 and observe the asterisk. The manifestation code (*) is used together with the code of the underlying disease marked with a dagger (†), in this case E11.4 for type 2 diabetes [133].
5. Recall the WHO rule: a code found by searching (or in the alphabetical index, Volume 3) is always verified in the tabular list, together with the inclusion and exclusion notes [133].

L4.3 The ICD-11 Browser and the Coding Tool (15 minutes)

1. In the Coding Tool, type "fracture of neck of femur". The list on the left (Word list) suggests words, and the list on the right (Destination Entities) displays the codes, sorted by default by match score. The "[Details]" link opens the description of the entity.
2. Search for the same term in the ICD-11 browser and open NC72.20, "Fracture of neck of femur, subcapital". In the "Postcoordination" section you will find the allowed axes: laterality (XK8G left, XK9K right), fracture subtype, open or closed fracture (XJ7YM, XJ44E), extension into the joint, and the associated external cause [141].
3. Search for "acute cystitis". Observe that the result is not a single code but an already assembled cluster. Write it down and identify which part is the stem code and which is the extension.
4. Search for "COVID-19 with pneumonia". Observe the separator between the two stem codes.

Remember the syntax from the WHO guide: "&" links a stem code to its extension codes, "/" links two stem codes, and an extension code can never occupy the first position [134].

L4.4 Coding Ten Cases (35 minutes)

The main condition is the condition diagnosed at the end of the episode of healthcare that is primarily responsible for the patient's need for treatment or investigation [133]. ICD-11 keeps the idea: the reason for admission, established at the end of the episode [134]. For each case, also note on the sheet the lead term used in the search, usually the noun that names the disease ("infarction", "fracture", "cystitis"). If the first term gives no results, try a synonym.

For each case, identify the main condition, code it in ICD-10 (2019 version) and in ICD-11 (MMS 2026-01), and fill in the sheet. Where the text calls for more detail, use postcoordination. The solutions are given at the end of the session, for the instructor.

1. Man, 58 years old, admitted for chest pain of 2 hours' duration. ECG: ST elevation in the anterior leads. Discharge diagnosis: acute ST elevation myocardial infarction, anterior wall.
2. Woman, 64 years old, follow-up visit for type 2 diabetes mellitus, without documented complications.
3. Man, 71 years old, pneumonia, with *Streptococcus pneumoniae* identified in blood culture.
4. Woman, 55 years old, essential hypertension, without documented organ damage.
5. Woman, 82 years old, tripped on a carpet at home and fell. Diagnosis: closed subcapital fracture of the left femoral neck.
6. Man, 19 years old, operated on for acute appendicitis with localized peritonitis.
7. Woman, 27 years old, allergic asthma, currently without exacerbation.
8. Woman, 30 years old, acute cystitis.
9. Man, 66 years old, COVID-19 confirmed by laboratory testing, with pneumonia.

10. Woman, 59 years old, with type 2 diabetes mellitus, admitted for investigation and treatment of diabetic polyneuropathy, which is the main condition of the episode.

For case 10, compare the two solutions. ICD-10 uses the dagger-asterisk system, whereas ICD-11 uses a cluster with two stem codes. For morbidity, the WHO guide places first the manifestation that meets the definition of the main condition, followed by the diabetes code, as in the example of diabetic retinopathy 9B71.0Z/5A11 [134].

L4.5 Postcoordination Exercise (10 minutes)

Starting from case 1, add the information "lesion of the left anterior descending artery". Search for the anatomy extension in the ICD-11 browser and build the complete cluster. Then answer: what information in the cluster cannot be expressed in ICD-10?

Starting from case 5, build a cluster that contains, in this order, the fracture, the laterality, the closed type of the fracture, and the external cause. The WHO guide gives an example with the same structure for a fracture of the left ulna after a fall: NC32.2&XK8G&XJ7YM/PA60&XE1DA&XE53A [134]. Explain why PA60 is preceded by "/" and not by "&".

L4.6 ICD-10-ICD-11 Mapping (10 minutes)

1. From the "Info" tab of the ICD-11 browser, open the ICD-10-ICD-11 mapping tables.
2. Look up in the table the ICD-10 codes you assigned to cases 1, 5, and 8. Note on the sheet what the table proposes.
3. Compare the table's proposal with the code or cluster you built. Where do they differ? Which details of the clinical text cannot be recovered from the ICD-10 code alone?

The WHO states that the mapping tables serve to compare data across revisions, not to convert data automatically [135]. On real data, the tables worked well for causes of death: applied to 24,645 deaths in Alberta, they successfully mapped 96.6% of the ICD-10 codes in adults [144]. For morbidity, however, information missing from the ICD-10 code remains missing after mapping. The difference in granularity is also visible in the opposite direction: in the Canadian map from ICD-11 to ICD-10-CA, 65% of the target ICD-10-CA codes were broader than the source ICD-11 code [145].

L4.7 Underlying Cause of Death and Order within the Cluster (10 minutes)

1. A fictitious death certificate reads: (a) pneumonia, due to (b) hip fracture, due to (c) tripping on a carpet. Identify the starting point of the sequence, then its code in ICD-10 and in ICD-11. The WHO guide discusses this example and states that the starting point is the external event [134].
2. A patient with type 2 diabetes mellitus is admitted in a diabetic coma. Write the ICD-11 cluster for morbidity. The same patient dies of the coma. Write the cluster for mortality. Explain why the order is reversed.

Common Pitfalls

- Copying the first search result without opening the entity page. Exclusions often point to another code.
- Stopping at the three-character category although subcategories exist (for example, J45 instead of J45.0).
- Entering the external cause code as the main condition in morbidity. In morbidity, the injury is the main condition, and the external cause is added.
- Writing an ICD-11 extension code alone or at the beginning of the cluster.
- Using the "&" separator between two stem codes.
- Reporting a code without its version. The same concept may have a different code in another release.
- Accepting a code proposed by a chatbot without verification. In a 2024 study, large language models given the official code description generated the exact code in fewer than half of the cases [165].

What to Submit

The spreadsheet file (or the printed sheet), with:

- Table L4.1 completed for the 10 cases: 6 points (0.3 points for the correct and complete ICD-10 code, 0.3 points for the correct ICD-11 code or cluster, with the correct order and separators)
- the clusters from the postcoordination exercise and the answers to the questions: 1 point
- the comparison with the mapping table for three cases: 1 point
- the mortality exercise and the explanation of the order within the cluster: 1 point
- awarded by default: 1 point.

Indicative grading: a correct code with the wrong version, or without verification in the tabular list, receives half of the points. For comparison, in the Canadian field test the mean coding time with ICD-11 fell from 23.6 to 9.9 minutes per chart as the coders gained experience [137], and in Korea experienced specialists achieved a mean accuracy of 71.6% when coding isolated diagnostic terms in ICD-11 [138]. Errors in the first cases are normal. The aim is for the student to be able to explain them.

Offline Variant

If the room has no internet connection, the instructor prepares in advance, on paper or as PDF files, excerpts from the ICD-10 tabular list (Chapters IV, VI, IX, X, XI, XIV, XIX, XX, and XXII, only the categories used in the cases) and screenshots of the ICD-11 pages for the entities in the solutions, including the "Postcoordination" section. Students work on the same cases, and searching is replaced by browsing the hierarchy in the excerpts. Alternatively, the WHO allows the ICD API service to be installed locally, for example in a Docker container, on a laboratory computer configured beforehand [546]. Exercise L4.6 is postponed to a session with network access or is done on a printed excerpt of the mapping table.

Solutions (for the Instructor)

All codes were verified on September 27, 2026, in the WHO ICD-10 browser, 2019 version [133], and in the ICD-11 MMS browser, release 2026-01 [141]. The order of the codes in the clusters follows the ICD-11 Reference Guide [134].

Table L4.2. Solutions for cases 1-10 (morbidity). Source: [133,134,141].

Case	ICD-10 (2019)	ICD-11 (MMS 2026-01)	Remarks
1	I21.0	BA41.0&XA7RE3	XA7RE3 = anterior wall of heart
2	E11.9	5A11	5A11 has no subcategories
3	J13	CA40.07	J13 is a three-character code, with no subdivisions
4	I10	BA00.Z	BA00.Z = essential hypertension, unspecified
5	S72.0; external cause W01	NC72.20&XK8G&XJ44E/PA60	XK8G left, XJ44E closed, PA60 fall on the same level
6	K35.3	DB10.01	acute appendicitis with localized peritonitis
7	J45.0	CA23.02	allergic asthma, uncomplicated
8	N30.0	GC00.1&XT5R	GC00.1 infectious cystitis, XT5R acute
9	U07.1 (optional, additional code for pneumonia)	RA01.0/CA40.1Z	CA40.1Z = viral pneumonia, unspecified
10	E11.4† G63.2*	8C03.0/5A11	in morbidity, the main manifestation comes first

Answers to exercises L4.5 and L4.7:

- L4.5, case 1: BA41.0&XA7RE3&XA7NQ7. The affected artery (XA7NQ7, left anterior descending coronary artery) has no counterpart in ICD-10, where I21.0 describes only the wall. The cluster reproduces the example in the WHO guide [134].
- L4.5, case 5: NC72.20&XK8G&XJ44E/PA60. PA60 is a stem code from Chapter 23 (external causes), so it is linked with "/". The extensions XK8G and XJ44E belong to the fracture code and are written immediately after it.
- L4.7.1: the starting point is the tripping, that is, the external event. The code is W01 in ICD-10 ("Fall on same level from slipping, tripping and stumbling") and PA60 in ICD-11 ("Unintentional fall on the same level or from less than 1 metre").
- L4.7.2: morbidity 5A23/5A11 (the coma is the reason for admission), mortality 5A11/5A23 (the diabetes initiated the chain of events). In ICD-10, both situations use the combination code E11.0, type 2 diabetes mellitus with coma [133,134].
- L4.3: the cluster for "acute cystitis" is GC00.1&XT5R, and the cluster for "COVID-19 with pneumonia, SARS-CoV-2 identified" is RA01.0/CA40.1Z [141].

PRACTICAL SESSION 5

Practical Session 5. Correct Citation, Plagiarism Prevention, and Writing a Scientific Essay

Objective

Using a reference manager, the student builds a Vancouver-style bibliography for a review essay of 1,200-1,500 words on a topic from Chapter 5. The student practices formatting five types of sources (journal article, book, book chapter, web page, legislative act), distinguishes paraphrasing from copying, reads a similarity report, and writes the statement on the use of generative artificial intelligence (AI) tools. The session applies the concepts of Section 5.6.

Tools

- **PubMed** (https://pubmed.ncbi.nlm.nih.gov/, accessed on September 27, 2026), for searching and exporting. The Cite button provides the NLM format, on which the Vancouver style is based, and downloads the .nbib file that reference managers import [41].
- **Zotero 10**, downloaded free of charge from https://www.zotero.org/download/ (accessed on September 27, 2026, when the Windows version was 10.0.3), together with the Zotero Connector extension for Chrome, Firefox, or Edge [547]. The plugins for Microsoft Word and LibreOffice are bundled with the program and are installed automatically on first startup, while work in Google Docs is done through the Zotero Connector extension, with the Zotero program open [548]. If needed, the Vancouver style is installed from the style repository, which contains more than 10,000 styles in CSL format [549].
- **Alternative: Mendeley Reference Manager**, version 2.149.0 for Windows 10 or later (https://www.mendeley.com/download-reference-manager/windows, accessed on September 27, 2026), with the Mendeley Cite add-in for Word [550].
- **Text editor**: Microsoft Word, LibreOffice Writer, or Google Docs.
- **The formatting guide** of the US National Library of Medicine, with examples for each source type [201], and the ICMJE Recommendations, updated in January 2026 [202].
- **The anti-plagiarism software** provided by the university, through the institutional account. The session does not depend on any particular product.

Procedure

Figure L5.1 shows the six steps of the session. The loop at the bottom is the most important part: the similarity report does not end the work but sends the student back to writing if unjustified matches or unsupported statements appear.

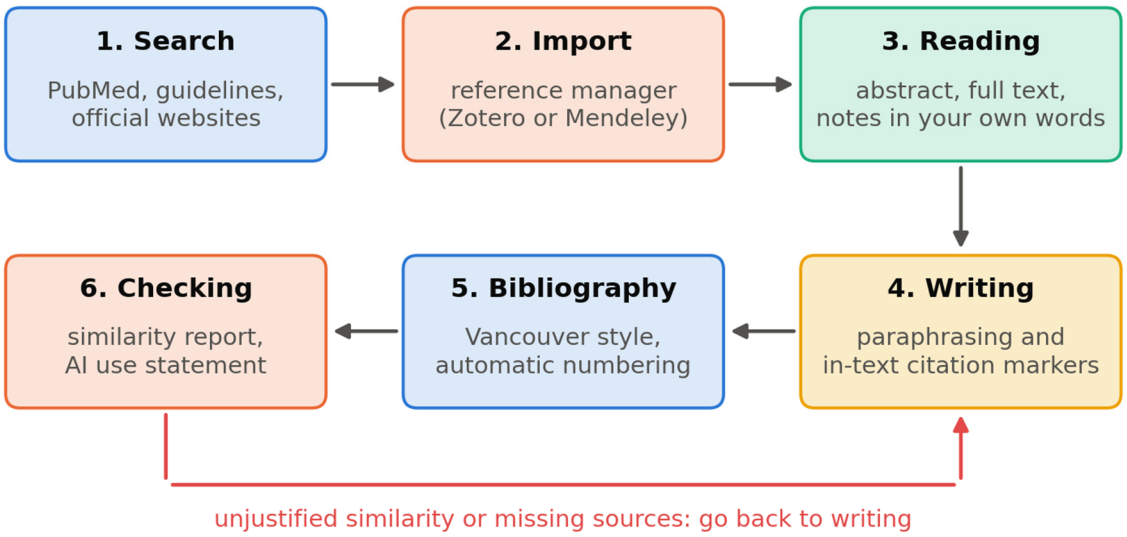


Figure L5.1. Workflow of Practical Session 5: from the search for sources to the final check. The red loop shows the return to writing after the similarity report has been read. Original diagram.

L5.1 Anatomy of a Vancouver Reference

The Vancouver style has two components. In the text, sources are marked with Arabic numerals, in the order in which they are first cited: [1], [2], [3]. A source cited again keeps its original number. At the end, the reference list follows the same numerical order, not alphabetical order [201,202].

The basic rules for the final list, according to the NLM guide [201], are as follows:

1. Authors are written as the surname followed by the initials, with no periods between the initials (Neprash HT). The first six authors are listed, followed by "et al." if there are more.
2. The article title is capitalized only at the beginning and for proper names.
3. The journal name is abbreviated according to the NLM Catalog (JAMA Health Forum, N Engl J Med), without periods.
4. Next come the year, the volume, the issue number in parentheses, and the pages: 2022;3(12):e224873. The page range is shortened (2129-31 instead of 2129-2131).
5. For online sources, [Internet], the date of access [cited ...], and the address after "Available from:" are added.

Table L5.1 gives one model for each source type. All the examples are real references related to Chapter 5, and their bibliographic data were checked in PubMed, CrossRef, or on the issuer's website.

Table L5.1. Vancouver templates and real examples. Source: [62,170,176,183,201,202].

Type	Template	Example
Journal article	Authors. Title. Journal. Year;Vol(No):pages.	Neprash HT, McGlave CC, Cross DA, Virnig BA, Puskarich MA, Huling JD, et al. Trends in ransomware attacks on US hospitals, clinics, and other healthcare delivery organizations, 2016-2021. JAMA Health Forum. 2022;3(12):e224873.
Book	Authors or editors. Title. Edition. Place: Publisher; Year.	Shortliffe EH, Cimino JJ, editors. Biomedical informatics: computer applications in healthcare and biomedicine. 5th ed. Cham: Springer; 2021.
Book chapter	Authors. Chapter title. In: Editors, editors. Book title. Edition. Place: Publisher; Year. p. pages.	Goodman KW, Miller RA. Ethics in biomedical and health informatics: users, standards, and outcomes. In: Shortliffe EH, Cimino JJ, editors. Biomedical informatics: computer applications in healthcare and biomedicine. 5th ed. Cham: Springer; 2021. p. 391-423.
Web page	Author. Title [Internet]. Place: Publisher; date [updated date; cited date]. Available from: URL	International Committee of Medical Journal Editors. Recommendations for the conduct, reporting, editing, and publication of scholarly work in medical journals [Internet]. [place unknown]: ICMJE; [updated 2026 Jan; cited 2026 Sep 27]. Available from: https://www.icmje.org/recommendations/
Legislative act	Title of the act. Official publication. Date;(number).	Legea nr. 190/2018 privind măsuri de punere în aplicare a Regulamentului (UE) 2016/679 al Parlamentului European și al Consiliului din 27 aprilie 2016 privind protecția persoanelor fizice în ceea ce privește prelucrarea datelor cu caracter personal și privind libera circulație a acestor date și de abrogare a Directivei 95/46/CE (Regulamentul general privind protecția datelor). Monitorul Oficial al României, Partea I. 2018 Jul 26;(651).

The formatting errors students make most often are predictable: an alphabetical list instead of numbering in order of citation, all 15 authors of an article instead of the first six followed by "et al.," a journal title written out in full or abbreviated by guesswork, a missing date of access for web pages, and citing the PubMed page instead of the article. Another, more serious error is citing a secondary source (a textbook, an online encyclopedia, an answer from a chatbot) for a figure that comes from a primary study. The figure must be traced to the original article, and it is that article that is cited.

Two further remarks. For a book chapter, the authors of the chapter are cited, not the editors of the book, and the chapter's pages are mandatory. For legislative acts, what matters is the date of publication in the Official Gazette of Romania (Monitorul Oficial), not the date of adoption: Law 190/2018 was adopted on July 18 and published on July 26, 2018 [176].

L5.2 The Reference Manager

The steps below are for Zotero. For Mendeley, the logic is the same, with different menu names.

1. Install Zotero and the Zotero Connector browser extension. Create a new collection named "PS5 Essay."
2. In PubMed, search for the chosen topic, for example `ransomware[ti] AND hospital`, with the year filter set to 2020-2026. For each relevant article, click the Zotero button in the browser toolbar. Alternatively, in Zotero, use the "Add Item by Identifier" button (the magic wand) and paste the PMID or DOI.

3. Check each record manually: the authors' names, the title, the journal, the year, the volume, the pages, and the DOI. Automatic import sometimes makes mistakes with diacritics, corporate authors, and electronic page numbers (of the type e224873).
4. Web pages and legislative acts are added manually ("New Item" > "Web Page" or "Statute," respectively), with the access date filled in.
5. A Zotero tab appears in Word. Choose the style "NLM/Vancouver: Citing Medicine 2nd edition (citation-sequence)," installed if needed from the Zotero style repository [549]. Use "Add/Edit Citation" to insert citations in the text and "Add/Edit Bibliography" to generate the list at the end.
6. Check the numbering: the first source cited in the text must be [1] in the list. If a paragraph is moved, click "Refresh," and the numbering is updated automatically.
7. At the end, export the collection in the chosen format (for example, RIS) by right-clicking the collection and selecting "Export Collection..." from the menu. An export is not a backup: reimported items lose their link to the citations in Word documents already written [551]. A RIS file can also be imported into Mendeley Reference Manager, through "Add New" > "Import Library" [552].

The citations inserted by the manager are active fields in the Word document. If the document is converted to plain text or copied into another program, the link to the library is lost and the numbering no longer updates. For this reason, always keep a copy of the document with the fields intact.

L5.3 Paraphrase, Quotation, Copying

The difference between the correct use of a source and plagiarism does not lie in the number of words changed but in the reference to the source and in the genuine rewording of the idea. Law No. 199/2023 explicitly includes in plagiarism the appropriation of ideas, data, and methods without reference to the source, not only of text [198]. The exercise uses a sentence from the abstract of the article by Neprash and colleagues [183]:

- **Original text** (excerpt): "the annual number of ransomware attacks more than doubled from 43 to 91".
- **Verbatim copying or word-for-word translation** (unacceptable as a paraphrase): "the annual number of ransomware attacks more than doubled from 43 to 91 [1]". The source is indicated, but the original authors' wording appears, without quotation marks, as the essay author's own wording. The same holds for a word-for-word translation of the sentence into another language: the fact that some anti-plagiarism programs do not detect translations does not make it correct.
- **Correct quotation**: Neprash and colleagues show that "the annual number of ransomware attacks more than doubled from 43 to 91" [1]. Direct quotation is used sparingly, for wording that we cannot render better.
- **Correct paraphrase**: "In the United States, the frequency of ransomware attacks on healthcare facilities rose more than twofold in six years, between 2016 and 2021 [1]." The idea is the same, the structure and the words belong to the author of the essay, and the source is indicated.

The student then writes, for two sources of his or her own choice, one paraphrase and one two-sentence summary for each, with citation. A useful method: read the paragraph, close the source,

write the idea from memory, and only then compare it with the original to check the accuracy of the figures.

L5.4 The Similarity Report

The anti-plagiarism program compares the text with databases of published documents and with previously submitted papers. The result is a similarity percentage and a list of sources, with the matching passages highlighted in color. The percentage alone does not show whether plagiarism has occurred: an identical passage may be a correct quotation, and the assessor must analyze each match (Section 5.6.1). The student proceeds as follows:

1. Upload the final version of the essay, with the bibliography included.
2. Read the list of sources in decreasing order of similarity. If a single source exceeds a few percent, check that passage first.
3. Classify each match: a quotation marked with quotation marks and a source (acceptable), a common technical expression such as "data protection impact assessment" (acceptable), an entry in the bibliography (acceptable), or a sentence taken over without quotation marks (to be rewritten).
4. Rewrite the passages in the last category and resubmit the text, if the system allows it.
5. Record on the submission sheet the initial percentage, the final percentage, and a brief note of what was changed.

No attempt should be made to trick the program by replacing letters, by mechanically chosen synonyms, or by automatic back-and-forth translation. The resulting text is hard to read, and the idea is still taken without a source.

L5.5 Declaring the Use of Generative AI

Generative AI tools may be used only within the limits set by the course coordinator. The ICMJE Recommendations require a description of how such tools were used, and AI tools cannot be listed as authors, because they cannot be held accountable for the accuracy and originality of the work. AI-generated material is not cited as a primary source [202]. In addition, COPE requires that the tool used be named [203]. A short section is added at the end of the essay, for example:

Example

- Statement on the use of AI: "I used [tool name, version] for grammar correction of the Introduction section and to propose a structure for the Discussion section. I searched for and read all sources myself in PubMed and on the issuers' websites. I take responsibility for the entire content."
- If AI was not used: "I did not use generative AI tools in writing this essay."

References suggested by a chatbot are not copied into the bibliography until they have actually been found in PubMed, CrossRef, or on the issuer's official page.

L5.6 Structure of the Essay

The review essay follows a structure close to IMRaD (Introduction, Methods, Results, and Discussion), adapted for a short review:

1. An informative **title** of no more than 15 words.

2. **Introduction** (150-200 words): the context, why the topic matters to the physician, and the precise question the essay answers.
3. **Methods** (100-150 words): the databases searched, the search terms, the period (2020-2026), the inclusion criteria, and the number of sources retained.
4. **Results** (500-600 words): a synthesis by theme, not a successive summary of each article. It may include a table of the main studies (author, year, country, N, result).
5. **Discussion** (300-400 words): which conclusion the sources support, where they disagree, and what the limitations of the evidence and of one's own search are.
6. **Conclusion** (2-3 sentences) and a Vancouver **Bibliography** with 8-12 sources, of which at least 6 are articles from journals indexed in PubMed.
7. **The AI statement**.

Suggested topics: occult hypoxemia and pulse oximetry, dynamic consent for research, the clinical impact of ransomware attacks, anonymization of medical images, and journal policies on generative AI.

What to Submit

The following are submitted on the course platform by the announced deadline:

1. The essay (1,200-1,500 words, excluding the bibliography), in .docx format, with the citations inserted from the reference manager.
2. The library export (.ris or .bib file) with all the sources cited.
3. The worksheet containing one reference of each of the five types in Table L5.1, formatted for a source chosen by the student; the two paraphrases from Section L5.3; and the notes on the similarity report from Section L5.4.
4. The AI statement.

Indicative grading scheme (10 points): IMRaD structure and clarity of the question, 2 points. Synthesis and accuracy of statements relative to the sources, 3 points. A complete Vancouver bibliography, without formatting errors and with correct numbering, 2 points. The paraphrases and the analysis of the similarity report, 2 points. The AI statement and meeting the deadline, 1 point. A passage taken over without attribution or a nonexistent reference cancels the score for the synthesis and is discussed with the course coordinator.

Offline Variant

If there is no internet access, the session is done on paper or in a local text editor. The instructor hands out printed sheets with the first pages of five sources (an article, a book, a chapter, a printed web page with the date of access noted, and a legislative act from the Official Gazette). The student writes the Vancouver references by hand, following the templates in Table L5.1, and then compares them with the model answers. Zotero also works offline for the records already saved in the library, so citation in Word can be practiced on a library prepared in advance. The paraphrasing exercises (L5.3) and the essay structure (L5.6) do not require a network connection. The similarity check is postponed until the connection is restored.

PRACTICAL SESSION 6

Practical Session 6. Analysis of ECG and EEG Signals with Free Software

Objective

The student opens public physiological recordings, views them with free programs, and extracts simple, reproducible measurements from them. By the end of the session, the student will be able to:

- identify normal and premature beats on an annotated ECG tracing.
- measure 20 consecutive RR intervals and calculate the mean heart rate, SDNN, and RMSSD in a spreadsheet.
- demonstrate aliasing by downsampling a simple signal.
- recognize the occipital alpha rhythm on an EEG recorded with the eyes closed and explain why it attenuates when the eyes open.

The session applies the concepts of Chapter 6 (sampling, quantization, QRS detection, heart rate variability, HRV). All data are public and anonymized. No clinical patient data are used.

Tools

All tools are free of charge. The addresses and versions were checked on September 27, 2026.

- **PhysioNet** (https://physionet.org/), the open resource of physiological signals maintained by the MIT Laboratory for Computational Physiology [226]. The standard citation requested of users is the 2026 paper describing the platform [227].
- **MIT-BIH Arrhythmia Database**, version 1.0.0: 48 ECG excerpts of 30 minutes each, two channels, 360 samples per second, 11-bit resolution over 10 mV, annotations verified by at least two cardiologists, Open Data Commons Attribution License [213].
- **EEG Motor Movement/Imagery Data set**, version 1.0.0: 109 subjects, 64 electrodes, 160 Hz, EDF+ files. The first two recordings of each subject are one-minute baselines, with the eyes open (R01) and with the eyes closed (R02) [553].
- **LightWAVE 0.71**, the PhysioNet viewer that runs in the browser, with no installation (https://physionet.org/lightwave/) [228].
- **EDFbrowser 2.15** (https://www.teuniz.net/edfbrowser/), a free program under the GPL version 3 license, for 64-bit Windows and Linux. It opens EDF, EDF+, and BDF files, converts other formats (including the MIT format), computes the power spectrum by FFT, and includes a Pan-Tompkins QRS detector with export of the RR intervals. The author states that the program must not be used in applications that are critical for patient safety [216].
- **Spreadsheet**: Google Sheets (web application) or Microsoft Excel.
- **Optional: Python 3** with the numpy library and the wfdb 4.3.1 package, which reads PhysioNet records directly [554,555].

Procedure

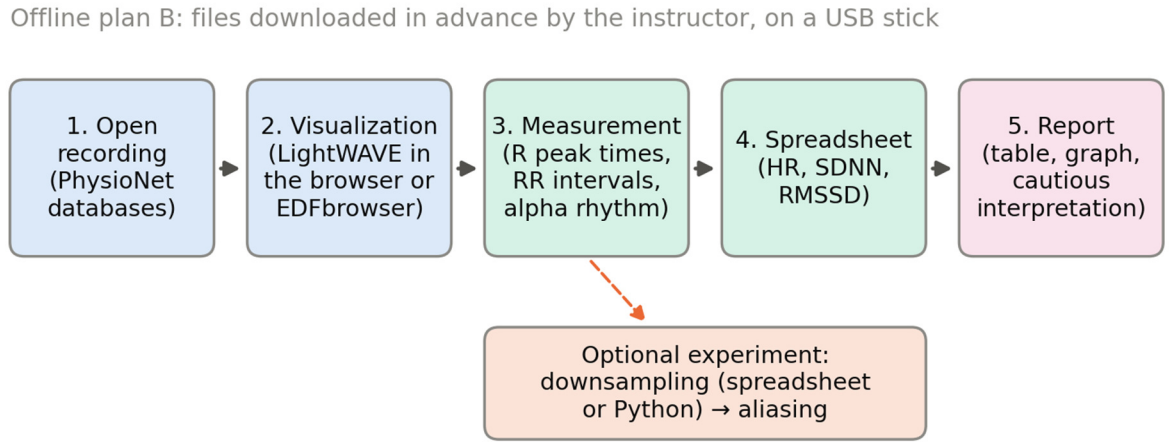


Figure L6.1. Workflow of Practical Session 6: from the open recording to the final report. Original diagram.

Figure L6.1 summarizes the steps. The suggested time is 100 minutes: 45 for the ECG, 20 for the sampling experiment, and 35 for the EEG.

L6.1 Part A. ECG: Viewing and Annotations

1. Open the page of the MIT-BIH Arrhythmia Database (https://physionet.org/content/mitdb/1.0.0/). Write down in your notebook the sampling frequency (360 Hz) and the resolution (11 bits over 10 mV). Calculate the quantization step: 10 mV / 2^11 ≈ 4.9 µV.
2. On the same page, find the section that lists the publications to be cited. Copy them into your report. Using a public database does not exempt you from citing the source.
3. Open LightWAVE with the MIT-BIH database: https://physionet.org/lightwave/?db=mitdb/1.0.0. In the **Choose input** tab, select record **100** from the **Record** list.
4. Switch to the **View/edit** tab. The screen shows 10 seconds of two leads, MLII and V5, with the background grid and the reference annotations (the `atr` row). Normal beats are marked with a dot. In the first 10 seconds, there is also an **A** annotation, that is, an atrial premature contraction [555].
5. Switch to the **Tables** tab. The summary shows the duration of the recording (00:30:06), the frequency of 360 samples per second, and the 2274 annotations: 2239 normal beats (N), 33 atrial premature beats (A), one ventricular extrasystole (V), and one rhythm annotation, `(N` [228]. Here you can also see that each signal has 200 conversion units per millivolt, so one unit corresponds to 5 µV. Compare this with the quantization step calculated in step 1.
6. Check the **View annotations as text** option to obtain the list of annotations with their times. If the list is not displayed, go back to **View/edit** and read the times of the R peaks on the time axis.

A check calculation, useful later: the 2273 beat annotations (2239 + 33 + 1) are spread over 30 minutes and 6 seconds, that is, over 1806 seconds. The mean heart rate over the whole recording is therefore approximately 2273 × 60 / 1806 ≈ 75.5 beats per minute. The value you obtain on a short segment should be of the same order of magnitude. A large difference usually indicates a unit error (seconds instead of milliseconds) or a missed beat.

L6.2 Part B. RR Intervals, Heart Rate, and HRV in the Spreadsheet

1. Choose a segment of approximately 20 seconds **without** A or V annotations. Enter the times of 21 consecutive R peaks in column B of Table L6.1 (in seconds, with three decimal places if you use the annotation list).
2. In column C, calculate the RR intervals in milliseconds: for row 3, `=(B3-B2)*1000`. This gives 20 intervals.
3. In column D, calculate the instantaneous heart rate: `=60000/C3`.
4. In column E, calculate the successive differences: for row 4, `=C4-C3`. This gives 19 differences.
5. Below the table, calculate:

- mean RR: `=AVERAGE(C3:C22)`
- mean HR: `=60000/AVERAGE(C3:C22)`
- SDNN: `=STDEV.S(C3:C22)`
- RMSSD: `=SQRT(SUMSQ(E4:E22)/COUNT(E4:E22))`

In localized versions of the programs, the function names may differ.

6. Build the tachogram: the graph of the RR intervals (vertical axis) against the time of the beat (horizontal axis). Figure L6.2 shows what a tachogram looks like and how a premature beat stands out.
7. Repeat the calculation on a segment that **contains** an A beat. Compare SDNN and RMSSD with the values from the clean segment and explain the difference. Reporting guidelines require ectopic beats to be identified and handled before HRV is calculated, preferably with visual inspection [212].

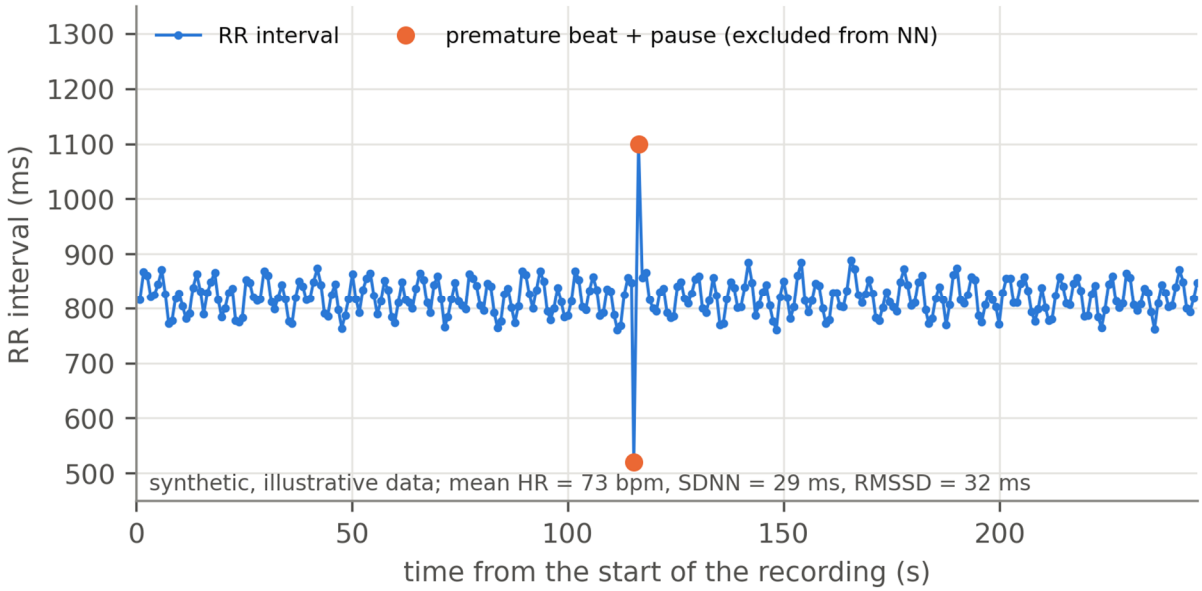


Figure L6.2. RR tachogram calculated from a synthetic series of 300 beats, with one premature beat followed by a pause. The two marked intervals are excluded from the NN series before SDNN and RMSSD are calculated. Synthetic, illustrative data generated by the script fig_L06.py.

Interpret with caution. With 20 intervals, the indices have teaching value only. The guideline report of the Society for Psychophysiological Research notes that RMSSD over windows of only 10 seconds and SDNN over windows of 30 seconds can approximate the values from long recordings only when several windows from a resting period are averaged [212]. RMSSD reflects short-term variability, and SDNN reflects overall variability, which depends on the length of the recording.

Table L6.1. Measurement sheet for RR intervals (spreadsheet template). The formulas are given for the row indicated and are copied down.

A: Beat no.	B: R time (s)	C: RR (ms)	D: HR (bpm)	E: RR difference (ms)
1 (row 2)	value read	not calculated	not calculated	not calculated
2 (row 3)	value read	=(B3-B2)*1000	=60000/C3	not calculated
3 (row 4)	value read	=(B4-B3)*1000	=60000/C4	=C4-C3
...	...	...	...	...
21 (row 22)	value read	=(B22-B21)*1000	=60000/C22	=C22-C21

L6.3 Part C. Sampling Experiment: Aliasing

The aim is for you to see with your own eyes what happens when the sampling frequency falls below twice the frequency of the signal.

1. In a new sheet, enter the numbers 0, 1, 2, ..., 199 in column A.
2. In column B, enter the time for sampling at 100 Hz: `=A2/100`. In column C, enter a 10 Hz sine wave: `=SIN(2*PI()*10*B2)`. Plot C against B for the first 50 rows. Count the cycles in one second.
3. In columns E and F, repeat the calculation for sampling at 12.5 Hz: the time `=A2*0.08` (written `=A2*0,08` in locales that use the decimal comma), and the value `=SIN(2*PI()*10*E2)`. Plot the graph for the first 25 rows (two seconds).
4. Count the cycles per second again. The graph shows a 2.5 Hz wave, that is, |10 - 12.5| Hz, which does not exist in the real signal. Compare it with Figure 6.2 in the chapter.
5. Link the result to the EEG: at 160 Hz, the sampling frequency of the EEG Motor Movement/Imagery data set, the highest frequency that is correctly reproduced is 80 Hz. What would happen to 50 Hz interference if the same EEG were downsampled to 60 Hz without an anti-aliasing filter?

The Python variant (optional) runs the same experiment and automatically calculates the dominant frequency:

```
import numpy as np

fs = 1000                       # initial sampling frequency (Hz)
t = np.arange(0, 2, 1 / fs)     # 2 seconds
x = np.sin(2 * np.pi * 10 * t) # 10 Hz sine wave

for k in (10, 40, 80):          # keep one sample out of k
    y = x[::k]
    fs_new = fs / k
    spectrum = np.abs(np.fft.rfft(y))
    freq = np.fft.rfftfreq(len(y), d=1 / fs_new)
    print(f"fs = {fs_new:6.1f} Hz -> peak at {freq[spectrum.argmax()]:.1f} Hz")
```

When run, the program displays a peak at 10 Hz for sampling at 100 Hz and at 25 Hz, but at 2.5 Hz for sampling at 12.5 Hz. Those who wish to work on the real ECG can read the first 10 seconds of record 100 directly from PhysioNet with the functions `wfdb.rdrecord('100', pn_dir='mitdb', sampto=3600)` and `wfdb.rdann('100', 'atr', pn_dir='mitdb', sampto=3600)`, documented for version 4.3.1 of the package [555]. Crude downsampling of the ECG (`x[::k]`) visibly distorts the QRS complexes.

L6.4 Part D. EEG: The Alpha Rhythm

1. From the S001 folder of the EEG Motor Movement/Imagery data set, open the files `S001R01.edf` (eyes open) and `S001R02.edf` (eyes closed), each approximately 1.2 MB in size [553].
2. Open `S001R02.edf` in EDFbrowser. Select the occipital channels (O1, Oz, O2) for display and, for comparison, a few frontal channels.
3. Set the display window to 10 seconds. Look on the occipital channels for rhythmic, sinusoidal activity of higher amplitude than on the frontal channels.
4. Measure the frequency: count the waves within a one-second interval at three different places in the tracing and take the mean. The posterior alpha rhythm is 8-12 Hz, is most evident with the eyes closed, and attenuates when the eyes open [209].
5. Use the power spectrum (FFT) function in EDFbrowser on channel O1, over 30-60 seconds of signal [216]. Write down the frequency of the main peak between 8 and 12 Hz.
6. Repeat steps 2-5 for `S001R01.edf`, with the eyes open. Compare the visual amplitude and the height of the spectral peak.
7. Observe the artifacts as well: blinks and eye movements appear mainly on the frontal channels, and muscle contraction adds irregular high-frequency activity. Some physiological artifacts, such as eye movements and muscle activity, can even help with interpretation [219].

The IFCN-ILAE minimum standard for routine clinical EEG recommends sampling at no less than 256 Hz and display with 0.5-70 Hz filters [208]. The data set used here, sampled at 160 Hz, is a research data set, suitable for the exercise but not for diagnosis. For the alpha rhythm, 160 Hz is sufficient: the highest frequency that is correctly reproduced is 80 Hz, more than six times the upper limit of the alpha band.

Common Pitfalls

- Measuring RR intervals in seconds and applying the formula HR = 60000/RR, which is meant for milliseconds.
- Including in the calculation the intervals around an A or V beat. The result is greatly increased SDNN and RMSSD values that no longer describe sinus variability.
- Confusing the mean RR with the mean of the instantaneous heart rates. The two do not give exactly the same mean heart rate.
- Looking for the alpha rhythm on the frontal channels or in the recording with the eyes open.
- Reading a narrow spectral peak at 50 or 60 Hz (depending on the country where the signal was recorded) as brain activity. It is power-line interference.

Table L6.2. Observation sheet for the alpha rhythm.

Recording	Channel	Counted frequency (Hz)	Spectral peak (Hz)	Remarks (amplitude, artifacts)
S001R02, eyes closed	O1			
S001R02, eyes closed	O2			
S001R01, eyes open	O1			
S001R01, eyes open	O2			

What to Submit

Each student submits a single electronic report file (a PDF or the spreadsheet document) containing:

1. Table L6.1 completed for the clean segment and for the segment with the A beat, with the mean RR, mean HR, SDNN, and RMSSD (3 points).
2. The tachogram of the clean segment, with labeled axes and units (1 point).
3. The graphs from the aliasing experiment and the calculated apparent frequency (2 points).
4. Table L6.2 completed and a screenshot of the power spectrum for channel O1 in both conditions (2 points).
5. Short answers (no more than 5 lines each) to three questions: (a) why the A beat is excluded from the HRV calculation, (b) what minimum sampling frequency a signal with components up to 150 Hz would require, and (c) why the alpha rhythm decreases when the eyes open (1 point).
6. Correct citation of the databases and programs used (1 point).

Indicative grading scheme: 10 points. In the practical examination, the student may be given another segment of record 100 or another record from the MIT-BIH database and must repeat the steps of L6.2 within 20 minutes.

Offline Variant

If the room has no internet access, the instructor downloads from PhysioNet in advance the files `100.hea`, `100.dat`, and `100.atr` of record 100 (about 1.9 MB in total) [213] and the files `S001R01.edf` and `S001R02.edf` [553]. The files are copied to the computers in the room together with EDFbrowser. The ECG recording is converted to EDF+ with the program's MIT-format conversion utility, and the RR intervals are obtained with the built-in QRS detector [216]. Part C requires no connection. If EDFbrowser cannot be installed either, Part B is done on a synthetic series of RR intervals prepared by the instructor, for example the one in Figure L6.2, generated by the script `fig_L06.py`.

PRACTICAL SESSION 7

Practical Session 7. Medical Image Processing with ImageJ/Fiji

Objective

The student opens a CT image in DICOM format, reads the header, applies a display window, takes calibrated measurements (length, area, mean density), compares two noise reduction filters, segments objects by thresholding, counts them, and exports the results to CSV. All steps are recorded in a macro that can be run again on another image. The session applies the concepts from Chapter 7 (Sections 7.1, 7.4, and 7.6).

Tools

- **Fiji**, the ImageJ distribution that comes bundled with numerous plugins for image analysis [556]. It can be downloaded free of charge from https://imagej.net/software/fiji/downloads (accessed on September 27, 2026). The recommended variant is "Latest," which includes Java 21 and runs on Windows 10 or newer, macOS 11 or newer, and Ubuntu 22.04 or newer. The "Stable" variant, with Java 8, remains available for older computers [557].
- **The ImageJ core**: on the date of access, the latest version in the official release notes was 1.54t, released on May 16, 2026 [558]. The installed version can be checked in Help > About ImageJ.
- **Images**: the ImageJ sample images, opened from the File > Open Samples menu, in particular "CT (16-bit DICOM)" and "Blobs" [559]. Alternatively, a public series from The Cancer Imaging Archive (TCIA) can be used; its collections are de-identified and, for the most part, available under Creative Commons Attribution 3.0 or 4.0 licenses for educational use [560]. Images of real hospital patients are not used.
- **Spreadsheet program** (Excel, LibreOffice Calc, or Google Sheets) for the exported CSV file.

Procedure

Figure L7.1 summarizes the eight steps of the session and the menus used. Table L7.1 can be kept open next to Fiji throughout the session.

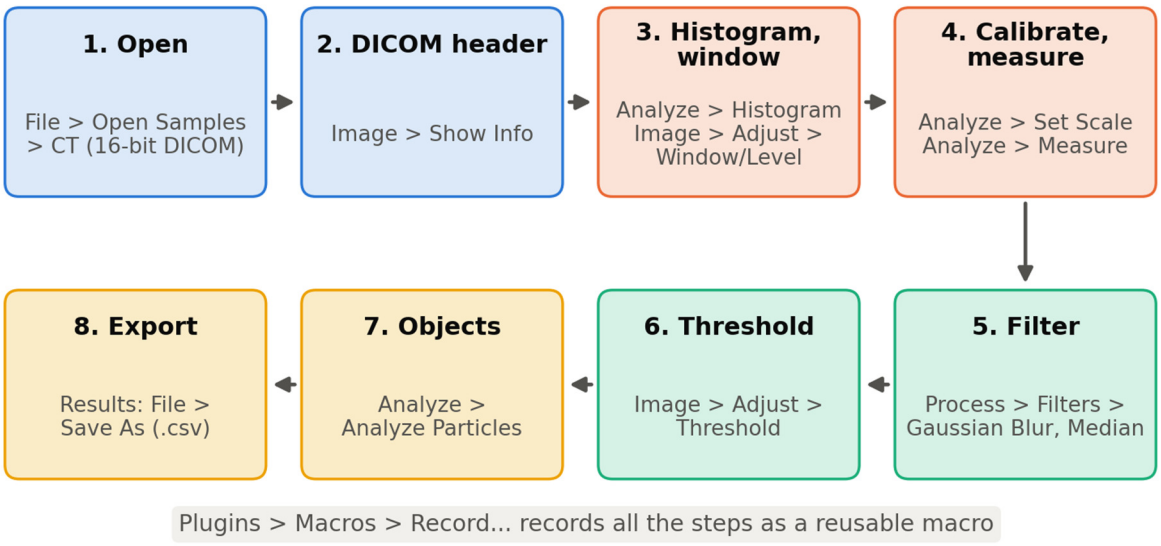


Figure L7.1. Workflow of Practical Session 7 in Fiji, from opening the DICOM image to exporting the results. The macro recorder captures all the steps.

Table L7.1. Operations used in the session and their paths in the Fiji menus. Source: [556,559].

Operation	Menu path	Shortcut	What to record
Opening a sample	File > Open Samples > CT (16-bit DICOM)	-	image size
DICOM header	Image > Show Info	Ctrl+I	tags (0028,0030), (0028,1050), (0028,1052)
Histogram	Analyze > Histogram	Ctrl+H	mean, min, max
Window	Image > Adjust > Window/Level	-	L and W used
Calibration	Analyze > Set Scale	-	mm per pixel
Measured quantities	Analyze > Set Measurements	-	which columns appear
Measuring	Analyze > Measure	Ctrl+M	length, area, mean
Filters	Process > Filters > Gaussian Blur / Median	-	sigma or radius
Threshold	Image > Adjust > Threshold	Ctrl+Shift+T	method, limits
Counting objects	Analyze > Analyze Particles	-	number, mean area
Macro	Plugins > Macros > Record	-	saved .ijm file

L7.1 Installation and Setup

1. Download the Fiji "Latest" archive for your operating system from the official page [557].
2. Unzip it into a folder where you have write permission (for example `C:\LP7\Fiji` or your home folder). Fiji does not require a conventional installation.
3. Start the Fiji executable from the folder. On first launch, the program may offer updates. Accept them only if the teaching assistant tells you to, so that the whole group works on the same version.
4. Write down the version shown in Help > About ImageJ.
5. Create the working folder `C:\LP7\rezultate`, where you will save all the files.

L7.2 Opening the Image and Reading the DICOM Header

1. Open File > Open Samples > CT (16-bit DICOM). The image is a CT slice downloaded from the ImageJ server, so an internet connection is needed.
2. Move the mouse over the image and watch the status bar of the main window: the coordinates and the pixel value appear.
3. Open Image > Show Info (Ctrl+I). The window displays the DICOM header as a list of tags of the form (group,element), with the attribute name and the value.
4. Find and copy onto the worksheet the values of: Modality (0008,0060), Rows (0028,0010), Columns (0028,0011), Pixel Spacing (0028,0030), Bits Allocated (0028,0100), Bits Stored (0028,0101), Window Center (0028,1050), Window Width (0028,1051), Rescale Intercept (0028,1052), and Rescale Slope (0028,1053) [245]. If an attribute is missing, write "absent."
5. Check the calibration of the values: place the cursor in the air outside the patient. If the image is displayed in Hounsfield units, the value should be around -1000. If you see large positive values, apply the relation HU = Rescale Slope · stored value + Rescale Intercept manually for three pixels and write down the calculation.
6. Look in the header for the attributes that would identify a patient, for example (0010,0010) Patient's Name and (0010,0020) Patient ID. Note whether they are empty, de-identified, or filled in. Discuss with the teaching assistant what should be deleted before such an image is sent to a colleague for research.

L7.3 The Histogram and the Display Window

1. Open Analyze > Histogram (Ctrl+H). Observe the two large groups of values: air (around -1000) and soft tissues (around 0-100). Write down the mean, minimum, and maximum shown below the graph.
2. Open Image > Adjust > Window/Level. Press the Set button and enter a level of 40 and a width of 400 (the mediastinal window). Describe what you see in the lungs.
3. Repeat with a level of -600 and a width of 1500 (the lung window). Describe what happens to the soft tissues.
4. Using the formula from Chapter 7, calculate the gray level of a 20 HU pixel in each of the two windows and compare it with what you observe on the screen.
5. **Do not press the Apply button.** It modifies the pixel values, not just the display. If you pressed it by mistake, close the image without saving it and reopen it.

L7.4 Spatial Calibration and Measurements

1. Check the pixel size in Image > Properties. For a DICOM file, it is taken from Pixel Spacing, and the unit should be mm.
2. Open Analyze > Set Scale only to see the value. Do not change it. Explain on the worksheet what would happen to the measurements if you entered 1 pixel = 1 mm.
3. In Analyze > Set Measurements, check: Area, Mean gray value, Standard deviation, Min & max gray value, and Perimeter.
4. With the Straight line tool, draw the transverse diameter of the vertebral body or of the aorta, as appropriate. Press Ctrl+M. The length in mm appears in the Results window.

5. With the Oval tool, draw a region of interest (ROI) inside the aorta or the liver, away from the edges. Press Ctrl+M. Write down the area (mm²), the mean, and the standard deviation, in HU.
6. Add the ROI to Analyze > Tools > ROI Manager (the Add button), so that you can reuse it identically in the next step.
7. Repeat the measurement of the mean in an ROI placed in air. The expected value is approximately -1000 HU [24].

L7.5 Noise Filtering

1. Duplicate the image (Image > Duplicate) twice and name the copies "gauss" and "median."
2. On the "gauss" copy, apply Process > Filters > Gaussian Blur with sigma = 2 pixels.
3. On the "median" copy, apply Process > Filters > Median with radius = 2 pixels.
4. Apply the same ROI from the ROI Manager to each image and measure the mean and the standard deviation.
5. Fill in a table with three rows (original, gauss, median). The standard deviation in a homogeneous region estimates the noise, and filtering is one of the global operations used to reduce it [24]. The Gaussian filter smooths everything, including the edges, whereas the median filter mainly removes isolated points and preserves contours better. Look at the edge of the aorta on the two copies and write down the difference.

L7.6 Threshold Segmentation and Object Counting

Thresholding turns a grayscale image into a binary mask, object versus background. It is a simple form of segmentation, often used, for example, to remove the structures that obscure the region of interest in volume rendering [24]. Figure L7.2 shows the principle on a synthetic image with 14 objects on a noisy background.

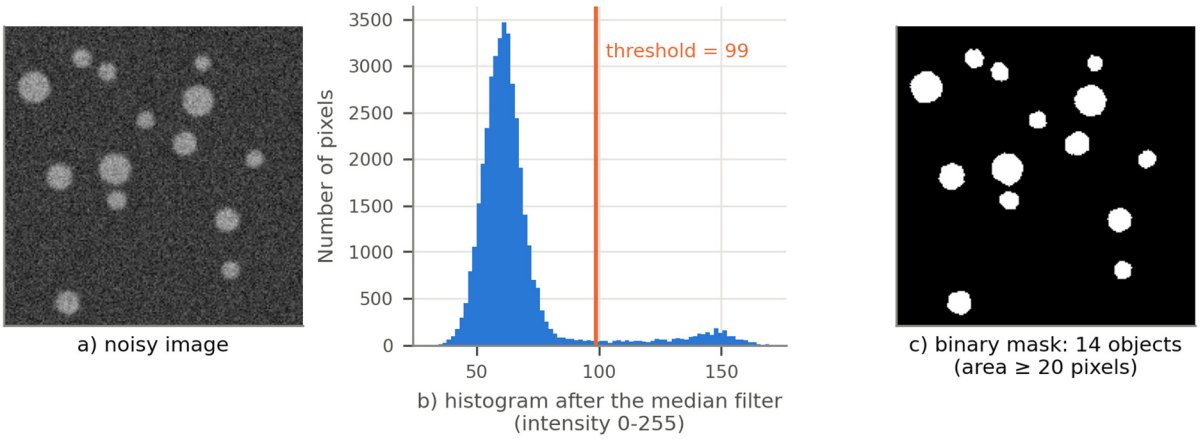


Figure L7.2. Threshold segmentation on a synthetic image: a) the noisy image, b) the histogram after the median filter, with the threshold chosen automatically by Otsu's method, c) the binary mask, in which 14 objects with an area of at least 20 pixels are counted. Synthetic, illustrative data.

1. Open File > Open Samples > Blobs.
2. Open Image > Adjust > Threshold. Choose the Default method, then try Otsu. Observe the area colored in red, which shows the selected pixels. Check or uncheck the Dark background option until the objects, not the background, are selected.
3. Press Apply. The image becomes binary (black and white).

4. Open Analyze > Analyze Particles. Set Size (pixel^2) to 50-Infinity, check Display results, Summarize, and Exclude on edges, then press OK.
5. From the Summary window, write down the number of objects (Count) and the mean area (Average Size). Repeat with Exclude on edges unchecked and explain the difference.
6. Optionally, on the CT image: with Image > Adjust > Threshold and the Set button, select the values from -1000 to -400 HU. The marked pixels correspond to air and lung. Measure the lung area on this slice with Analyze Particles. Compare the result with that of a colleague who chose different limits.

L7.7 Recording a Macro and Exporting the Results

Fiji can be automated through macros and scripts and, for those who program, from Python [561,562]. The macro recorder turns each menu command into a line of code.

1. Close all images and results windows. Start Plugins > Macros > Record.
2. Repeat the steps from L7.6 on Blobs: median filter with radius 2, threshold, Apply, Analyze Particles.
3. Lines similar to those below will appear in the Recorder window (the generated code may differ slightly, depending on the options chosen):

```
run("Median...", "radius=2"); setAutoThreshold("Default"); run("Convert to Mask");
run("Analyze Particles...", "size=50-Infinity display exclude summarize");
```

4. Press Create, save the macro as `C:\LP7\rezultate\lp7_numarare.ijm` and run it on a newly opened image (Macros > Run or Ctrl+R in the editor). The result should be identical to the one obtained manually.
5. In the Results window, choose File > Save As and save `C:\LP7\rezultate\lp7_blobs.csv`. Save the CT measurements in the same way, as `lp7_ct.csv`.
6. Open the CSV files in the spreadsheet program. Calculate the mean and the standard deviation of the object areas and compare them with the values in the Summary window.

L7.8 Discussion: Limitations and Reproducibility

At the end of the session, the group compares the results. The lung area from the optional step in L7.6 usually differs from one student to another, even though the image is the same. The difference does not come from the software but from the choices made: the threshold limits, the shape and position of the ROI, and the filter applied beforehand. The same problem appears, on another scale, in radiomics, where the features extracted from images vary with the scanner, the protocol, and the software used [265]. For this reason, the IBSI initiative defined and verified reference values for 169 features, so that different software packages can be calibrated against one another [266].

The macro saved in L7.7 is the practical answer to this variability: it documents exactly what was done and allows anyone to repeat the analysis. The `.ijm` file is submitted together with the results, just as a scientific article describes its method.

In this session, Fiji is used for teaching purposes only. Software used to make a diagnosis must meet the requirements for medical devices [284], and in the hospital, images are interpreted on calibrated workstations and monitors connected to the PACS. Exporting an image to JPEG "to make measuring easier" is a common mistake: the spatial calibration and the HU values are lost, and with them the possibility of a correct measurement.

What to Submit

The files are uploaded to the discipline's platform in an archive named `LP7_Nume_Prenume.zip` (surname and first name). Indicative grading scheme, 10 points:

- the worksheet with the DICOM attributes from L7.2 and the check of the air value (2 points);
- the description of the two windows and the gray level calculation for 20 HU (1.5 points);
- `lp7_ct.csv` with the length, area, and mean in HU, with the correct units (2 points);
- the table with the original, gauss, and median rows and the comment on noise and edges (1.5 points);
- `lp7_blobs.csv`, the number of objects and the explanation of the Exclude on edges option (1.5 points);
- the macro `lp7_numarare.ijm`, which runs without errors (1.5 points).

In the practical examination, the student may receive an unknown image and must read Pixel Spacing from the header, measure a structure in mm, and count objects by thresholding, explaining each choice.

Offline Variant

If the laboratory has no internet access, before the session, the teaching assistant copies onto each computer (or onto a USB stick) the Fiji archive downloaded from the official page and the sample files `ct.dcm` and `blobs.gif`, which are then opened with File > Open. If the sample images are not available either, a synthetic image equivalent to the one in Figure L7.2 is generated by running the following code in Plugins > New > Macro:

```
newImage("sintetic", "8-bit black", 256, 256, 1); setColor(150); fillOval(30, 40, 22, 22); fillOval(120, 60, 16, 16); fillOval(190, 150, 26, 26); fillOval(70, 180, 18, 18); run("Add Specified Noise...", "standard=18");
```

The steps from L7.5, L7.6, and L7.7 are carried out on this image. Since the background is black, Dark background is checked in the Threshold dialog, and in the macro the threshold line becomes `setAutoThreshold("Default dark");`. The number of objects should be 4. In this case, the DICOM part (L7.2-L7.4) is done on any de-identified DICOM file provided by the teaching assistant, and the questions about the header remain the same.

PRACTICAL SESSION 8

Practical Session 8. Working with an Electronic Health Record: The OpenEMR Demo and FHIR Resources

Objective

This session connects the two faces of the electronic health record (EHR) described in Chapter 8. The first is the face the clinician sees: screens, forms, lists. The second is the face that information systems see: FHIR resources exchanged through web requests. By the end of the session, the student will be able to:

- create a fictitious patient in OpenEMR and record a consultation with vital signs, a problem, an allergy, and a prescription;
- formulate search and read requests to a public FHIR test server in the browser;
- recognize, in a JSON resource, the main elements of the Patient and Observation resources (identifier, name, date, LOINC code, value, unit);
- map OpenEMR screens to FHIR resources;
- compare, at a conceptual level, the flow of a prescription in OpenEMR and FHIR with the flow of the e-prescription in Romania.

The basic rule of this session: only fictitious data are used. The servers used are public, the data entered are visible to other users, and entering a real patient's data would breach professional confidentiality and data protection legislation.

Tools

- An up-to-date **web browser** (Firefox, Chrome, or Edge). By default, Firefox displays JSON files in a structured viewer, which makes Part B easier.
- **OpenEMR, the official online demo**, version 8.4.1, released on September 20, 2026 [563]. The page https://www.open-emr.org/demo/ (accessed on September 27, 2026) offers three identical instances and public accounts for several roles (administrator, physician, clinician, accountant, front office). The passwords are displayed on the page and are not copied into the worksheet. Each instance is reset to its initial state daily at 8:00 UTC, so the data are not kept [564]. OpenEMR is open-source software and provides a FHIR R4 API conforming to the US Core profiles, with SMART on FHIR authorization [565].
- **The public HAPI FHIR test server**, R4 endpoint: https://hapi.fhir.org/baseR4 (accessed on September 27, 2026, when the server software reported the development version HAPI FHIR 8.13.10-SNAPSHOT) [566]. The server page warns that it is not a production server, that personal health information must not be stored on it, and that the database is periodically purged and reloaded with test data [566].
- **The FHIR specification**, for the vital signs profile and the corresponding LOINC codes [567].

- **A text editor** (Notepad, TextEdit, or equivalent) for the worksheet and for copying JSON fragments.

Procedure

The session has three parts (Figure L8.1). The suggested time is 30 minutes for Part A, 35 minutes for Part B, and 25 minutes for Part C.

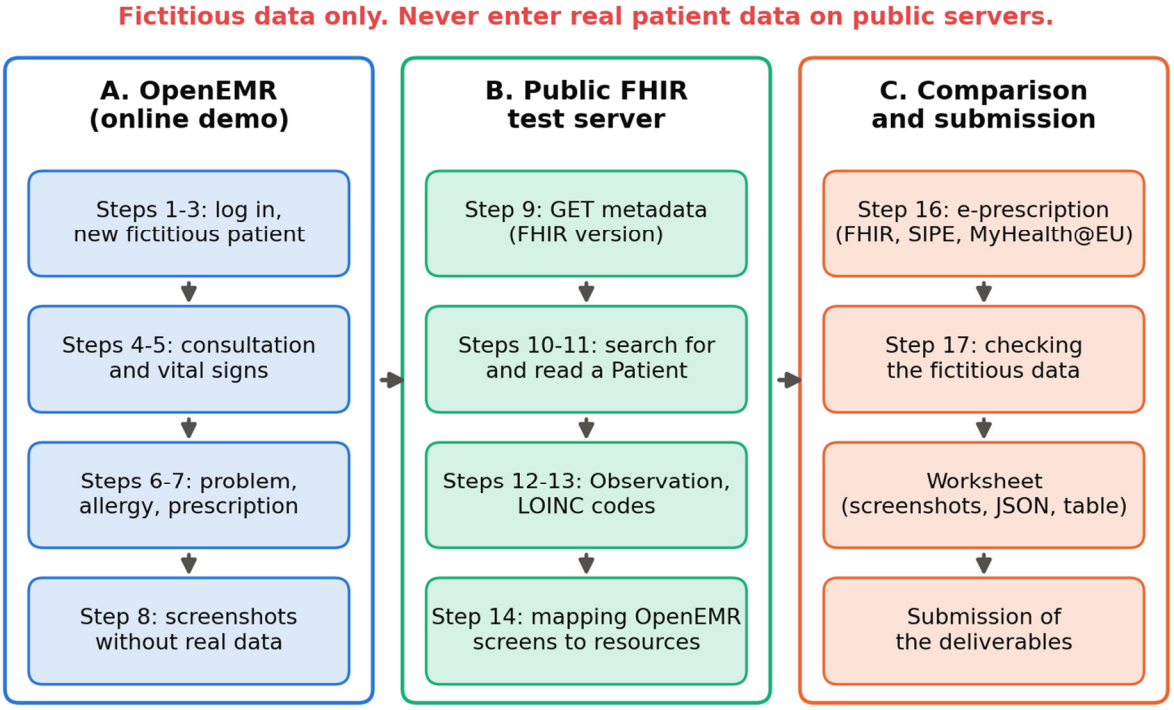


Figure L8.1. Workflow of the practical session: registering a fictitious patient in OpenEMR (A), reading the same types of data as resources on a public FHIR test server (B), comparison with the Romanian e-prescription and submission (C).

Each student chooses a personal working code, made up of the group initials and a number (for example, G3-17). The code becomes part of the fictitious patient's name, so that the record can be found among the records of fellow students.

L8.1 Part A: The Fictitious Patient in OpenEMR

1. Open https://www.open-emr.org/demo/, choose one of the three instances, and log in with the public physician or administrator account shown on the page. If an instance is slow, switch to another one.
2. Create the patient. From the main menu, choose *Patient* and then *New/Search*, fill in the required fields, and confirm with *Create New Patient* [568]. Use the fictitious data: last name Testescu, first name Ana-[your code], date of birth May 12, 1990, sex female. Do not fill in a personal numeric code (CNP), a real telephone number, or a real address.
3. Look at the patient record (*Dashboard*). Write down in the worksheet which sections you see. The OpenEMR documentation describes, among others, the *Medical Problems*, *Allergies*, *Medications*, and *Prescriptions* sections, each with an edit icon [568]. The exact labels may differ slightly from one version to another. If you cannot find a menu, look for it by function, not by name.

4. Open a new consultation (an *encounter*, called a visit in the interface). Fill in today's date, the visit category, and the reason: "Headache for 3 days." Save.
5. In the open consultation, add the vital signs form (*Vitals*) with the following values: blood pressure 150/95 mmHg, pulse 88/min, temperature 36.8 °C, weight 72 kg, height 165 cm. Save and check whether the program has calculated the body mass index. Write down the value displayed.
6. From the patient record, add a medical problem: "Essential hypertension," with today's date as the onset date. If the interface offers a search for a diagnosis code, write down the proposed code. Then add an allergy: penicillin, reaction of the urticaria type, moderate severity.
7. Add a prescription from the *Prescriptions* section (usually through *Edit* and then *Add*): amlodipine 5 mg, one tablet per day, 30 tablets. Save.
8. Take four screenshots: the patient record with the problem, the allergy, and the prescription visible; the vital signs form; the list of consultations; and the prescription screen. Make sure that no real personal data, yours or anyone else's, appear in the screenshots.

What to observe. Some data were entered in structured fields (the vital sign values, the date of birth, possibly the diagnosis code), others as free text (the reason for the consultation). Write down two examples from each category in the worksheet. This is the distinction between structured documentation and free text discussed in Section 8.1.2.

L8.2 Part B: The Same Concepts as FHIR Resources

Part B does not use the data entered in OpenEMR. The OpenEMR FHIR API requires authorization according to the SMART standard, which we do not configure in the laboratory [565]. Instead, we work on the public HAPI FHIR server, which already contains test patients and observations. The requests are typed directly into the browser's address bar. Each GET request is a read: it changes nothing on the server.

9. Check the server. Open https://hapi.fhir.org/baseR4/metadata?_format=json. The response is the capability statement (*CapabilityStatement*). Look for the `fhirVersion` field and write down its value (for this endpoint, 4.0.1, that is, R4).
10. Search for patients by last name: https://hapi.fhir.org/baseR4/Patient?family=smith&_count=5&_format=json. The response is a *Bundle* resource of type `searchset`. Each item in `entry` contains a Patient resource. The server returns the results in pages: the `_count` parameter sets how many results a page contains, and the `link` element with the relation `next`, if present, contains the address of the next page. Open that address and note that it is also a GET request. Then choose a patient and write down the `id` value.
11. Read the chosen patient: https://hapi.fhir.org/baseR4/Patient/[id]?_format=json, replacing [id] with the value you wrote down. Identify in the JSON the elements `name` (with `family` and `given`), `gender`, `birthDate`, and, if present, `identifier`. Copy the fragment into the worksheet.
12. Search for the observations of the same patient: https://hapi.fhir.org/baseR4/Observation?patient=[id]&_count=10&_format=json. If the result is empty, search directly for pulse observations by LOINC code: https://hapi.fhir.org/baseR4/Observation?code=http://loinc.org|8867-4&_count=5&_format=json. For one observation, identify:

- `code.coding` (the system, for example http://loinc.org, the code, and the display name);
- `valueQuantity` (the value, the unit, and the unit code);
- `subject` (the reference to the patient, in the form Patient/[id]);
- `effectiveDateTime` (the time of measurement), if present;
- `category` (for example `vital-signs` or `laboratory`).

13. Link the observation to the vital signs entered in Part A. In the FHIR vital signs profile, pulse has the LOINC code 8867-4, body weight 29463-7, body temperature 8310-5, body mass index 39156-5, and blood pressure is a panel (85354-9) with two components, systolic (8480-6) and diastolic (8462-4) [567]. Explain in two sentences why blood pressure is not represented by a single value.
14. Complete Table L8.1, which maps the screens from Part A to FHIR resources. The element columns are already filled in. For each row, write down in the worksheet the actual value entered in OpenEMR that would end up in the respective element.

Table L8.1. Mapping between OpenEMR screens and FHIR R4 resources. Source: [565,567].

OpenEMR screen	FHIR resource	Main elements	Example from Part A
Demographics	Patient	name, gender, birthDate, identifier	Testescu, Ana-[code], 1990-05-12
Consultation (visit)	Encounter	status, period, reasonCode, subject	the reason "Headache for 3 days"
Vital signs	Observation (vital-signs)	code (LOINC), valueQuantity, effectiveDateTime	pulse 88/min, code 8867-4
Medical problems	Condition	code, clinicalStatus, onsetDateTime, subject	essential hypertension
Allergies	AllergyIntolerance	code, reaction, patient	penicillin, urticaria
Prescriptions	MedicationRequest	medication[x], dosageInstruction, authoredOn, requester	amlodipine 5 mg, 1 tab/day

All of these resources are among those exposed by the OpenEMR FHIR API [565]. In the R4 version used here, the linking elements are called `subject` (for Observation, Condition, Encounter, and MedicationRequest) and `patient` (for AllergyIntolerance). This difference in naming is one of the usual pitfalls when writing searches.

15. Optional, for students who finish early: add the parameter `&_summary=count` to the request from step 12. The response contains only the total number of results. Explain what such a request would be useful for in an application that displays the number of laboratory tests of a patient.

A practical warning. Any user can write to the HAPI server, and the data can be modified or deleted at any time. Do not assume that a given `id` will still exist next week, and do not draw clinical conclusions from these data. Do not use POST, PUT, or DELETE requests in the laboratory: they would write to a public server.

L8.3 Part C: Comparison with the Romanian E-Prescription

In Part A, the prescription was recorded in a single program. In FHIR, the same prescription would become a MedicationRequest resource, and its dispensing in the pharmacy a separate resource, MedicationDispense, which the OpenEMR API also exposes [565]. The ISO 17523

standard specifies what data an electronic prescription must contain, but leaves transmission, validation, and reimbursement to each country [312].

In Romania, e-prescriptions circulate through the Electronic Prescription IT System (Sistemul informatic de prescripție electronică, SIPE), part of the platform of the National Health Insurance House (Casa Națională de Asigurări de Sănătate, CNAS), and since September 1, 2026, patients can view their prescriptions from the last 5 years in the eSănătateaMea portal [322,324]. At the European level, electronic prescriptions and dispensations are two of the priority categories of the European Health Data Space, with mandatory exchange between Member States from March 26, 2029 [23].

16. Answer the following questions in writing, in 150-200 words:

- Who are the actors in the flow (physician, pharmacist, patient, health insurance house), and what information does each of them receive?
- Which FHIR resource would correspond to the issued prescription, and which resource would correspond to dispensing in the pharmacy?
- What would have to happen, technically and organizationally, for a prescription issued in Timișoara to be dispensed in a pharmacy in another Member State? Use the levels of interoperability in Figure 8.2.

17. Check the worksheet and the screenshots once more. No information may allow the identification of a real person.

What to Submit

A single document (PDF or DOCX) is submitted, with the student's name and group on the first page. The grading scheme is indicative, out of 10 points:

- 1 point by default;
- 2 points: the four screenshots from Part A, with the fictitious patient, the consultation, the vital signs, the problem, the allergy, and the prescription visible;
- 1 point: the two examples of structured data and the two of free text (step 8);
- 2 points: the exact addresses of the FHIR requests from steps 9-12 and the copied JSON fragments, with the required elements marked (underlined or highlighted);
- 1 point: the explanation of how blood pressure is represented (step 13);
- 1 point: Table L8.1 completed with your own values (step 14);
- 2 points: the answer to step 16, correct and well argued.

In the practical examination, the student may receive an unfamiliar JSON fragment (a Patient or Observation resource) and must identify, within a few minutes, the patient, the date, the code, the value, and the unit.

Offline Variant

If the laboratory has no internet access, Part A is replaced by a paper-based exercise and Part B by local files.

1. The instructor saves in advance, on the laboratory computers, two official examples from the FHIR R4 specification: https://hl7.org/fhir/R4/patient-example.json and https://hl7.org/fhir/R4/observation-example-heart-rate.json (addresses verified on September

27, 2026). The files are opened with a text editor or with the browser, without a connection. As a fallback, the Patient and HL7 v2 examples from Chapter 8 can be used.

2. The students fill in, on paper, a consultation form for the fictitious patient from Part A, with the same fields (demographics, reason, vital signs, problem, allergy, prescription), and mark the structured fields and the free-text fields.
3. Steps 11-13 are carried out on the local files: identifying the elements in the Patient resource and in the Observation resource for pulse, then writing a JSON fragment of type Observation for the fictitious patient's temperature of 36.8 °C, with the LOINC code 8310-5 and the unit Cel [567].
4. Table L8.1 and the answer to step 16 are completed in the same way. The grading scheme is kept, with the paper form instead of the screenshots and the handwritten JSON fragment instead of the requests to the server.

If the laboratory has a local OpenEMR installation prepared by the instructor, Part A is carried out on it, with the same fictitious data.

PRACTICAL SESSION 9

Practical Session 9. Simulating a Teleconsultation: Preparation, Consent, Documentation, and Limits

Objective

Students work in pairs and simulate a complete video teleconsultation, in which one plays the role of the physician and the other that of a fictitious patient. By the end of the session, each student should be able to: prepare a teleconsultation using a checklist, obtain and record consent, recognize the warning signs that require the remote consultation to be stopped, write the teleconsultation note following the template in Section 9.6, and critically assess a fellow student's note.

The aim of the session is not to teach a clinical technique. The aim is the informational discipline of the remote medical act: who the patient is, where the patient is, which channel we are communicating on, what we cannot see, and how the written trace is preserved. Studies of students show that these skills do not appear on their own. In a comparison of 241 second-year students, those assessed in a simulated telemedicine consultation obtained lower scores for the physical examination and interpersonal communication than those assessed face to face, but higher scores for history taking [569]. A two-hour module dedicated to video communication increased the self-reported confidence of final-year students [570].

The stakes are clinical, not merely formal. In a cohort of 2,393 patients at Mayo Clinic, the provisional diagnosis made by video was confirmed at the in-person examination in 86.9% of cases [342]. The difference of approximately 13% is the reason why the note must state clearly what could not be examined and when the physician will see the patient again. This session practices exactly these two actions.

No real patient data are used. All cases are fictitious and marked as such, and the consultations are not recorded.

Tools

- **A videoconferencing platform approved by the faculty**: students use the institutional account they already have. As a free, open-source alternative, Jitsi Meet can be used (stable version 2.0.11248, released on September 14, 2026, https://meet.jit.si, accessed on September 27, 2026) [571]. On the public meet.jit.si server, the person who creates the room must have an account, while participants do not need one [340]. For this reason, the rooms are created by the instructor, who sends the links to the pairs.
- **Security settings**: the student checks in the platform's documentation whether the stream is encrypted only in transit or also end to end. In Jitsi Meet, end-to-end encryption works only in certain browsers and does not cover chat messages [340].

- **Documents**: the teleconsultation note template (Section 9.6 of Chapter 9), the pre-visit checklist (Table L9.1), the fictitious case sheet received from the instructor, and the peer assessment grid.
- **Text editor**: Google Docs or Microsoft Word, in the institutional account. The note is submitted in PDF format.
- **Working reference**: the WHO guide to implementing telemedicine, for identification, consent, and standard procedures [334].

Procedure

L9.1 Preparation (15 minutes)

1. Pairs are formed at the beginning of the class. Each student will play both roles in turn, with two different cases.
2. The instructor hands out the case sheets. The "patient" sheet contains the age, the reason for the visit, the symptoms, the history, and, in some cases, a warning sign that is revealed only in response to an appropriate question. The "physician" sheet contains only the reason for the appointment. The model is inspired by the standardized patient scenarios used to train situational awareness in telemedicine [572].

The working cases (all fictitious) are chosen so that each pair encounters at least one situation that cannot be resolved remotely:

- **Case A**: a 58-year-old woman with hypertension, scheduled follow-up. The values measured at home over the last week are above target, with no symptoms. The case can be resolved remotely, with adjustment of treatment and reassessment.
- **Case B**: a 34-year-old man with a pigmented lesion that recently appeared on his back and that he cannot show well to the camera. The correct decision is an in-person examination or tele-expertise with a dermoscopic image obtained in a medical office.
- **Case C**: a 62-year-old man who asks for a "prescription for heartburn." When asked about exertion, he reports retrosternal pain with sweating that began an hour ago. The remote consultation is stopped, and the patient is directed immediately to call 112.
- **Case D**: a 21-year-old female student being treated for depression, follow-up visit. The assessment can be done by video, but the question about suicidal ideation is mandatory and is recorded.

3. The "physician" goes through the checklist in Table L9.1 before entering the room: identity, location, backup telephone number, consent, privacy of the space, and the plan for a dropped connection.
4. The "patient" chooses a quiet place, lit from the front, and keeps at hand an imaginary "blood pressure monitor" whose values are read from the sheet.

Table L9.1. Checklist before and at the start of the teleconsultation. Source: based on [334,370].

Moment	Item checked	How it is checked	Recorded
Before	the platform and camera work	audio-video test 5 minutes before	no
Start	the patient's identity	full name and date of birth, compared with the record	yes
Start	the physician's identity	name, specialty, institution	yes
Start	the patient's location	address at the time of the consultation	yes
Start	backup telephone number	the patient dictates it, the physician repeats it	yes
Start	privacy	who else is in the room, whether the patient can speak freely	yes
Start	consent	explanation of the limits and verbal agreement	yes
Start	recording	it is announced that the session is not recorded	yes
During	warning signs	targeted questions at the start of the history	yes
End	plan and safety	recommendations, signs for returning, time frame	yes

L9.2 Consent and the Framework of the Consultation (5 minutes)

5. The "physician" introduces themselves and confirms the "patient's" identity and location.
6. The "physician" explains, in simple words, what can and cannot be done remotely. An example of wording: "I can listen to you, see you, and read the values you have measured. I cannot listen to your heart and lungs. If I consider it necessary, I will ask you to come to the office or to call 112."
7. The "physician" asks for verbal agreement to continue the consultation remotely. The WHO guide specifies that a consultation initiated by the professional requires documented consent, and that the limits of telemedicine must be communicated explicitly [334]. In Romania, the telemedicine norms require, outside emergencies, the patient's express consent, proven by any means of evidence [364]. The verbal agreement is therefore recorded in the note, with the time at which it was obtained.
8. The "physician" sets out the backup plan: "If the connection drops, I will call you within two minutes at the number you gave me."

L9.3 The Consultation Itself (10-12 minutes)

9. **Initial triage.** The first questions look for warning signs: chest pain, dyspnea at rest, sudden weakness of a limb, speech disturbances, major bleeding, altered level of consciousness, suicidal ideation. If any of them appears, the remote consultation is stopped. The "physician" instructs the patient to call 112 or to go immediately to the emergency department and stays on the line until the case is taken over, if the situation requires it (Figure L9.1).
10. **The history** follows the usual structure: the reason for the visit, the history of the present illness, past medical history, medication, and allergies.

11. **The video-guided examination.** The "physician" asks the "patient" to perform simple actions: to come closer to the camera, to show the ankles, to count aloud for 30 seconds, and to read the blood pressure values. Reviews of the virtual physical examination describe precisely this type of maneuver, with its limits [339].
12. **The decision.** The "physician" judges whether the problem can be resolved remotely. If not, the "physician" schedules an in-person examination with a specified date and place.
13. **Closing.** The "physician" summarizes the plan, checks understanding (asks the "patient" to repeat what needs to be done), and gives the safety-netting advice: which signs require a return visit or a call to 112. This advice is often left out of the documentation [371].
14. At any point, the instructor may "interrupt" the connection of a pair (by asking the "patient" to turn off the camera). The "physician" must then apply the announced backup plan.

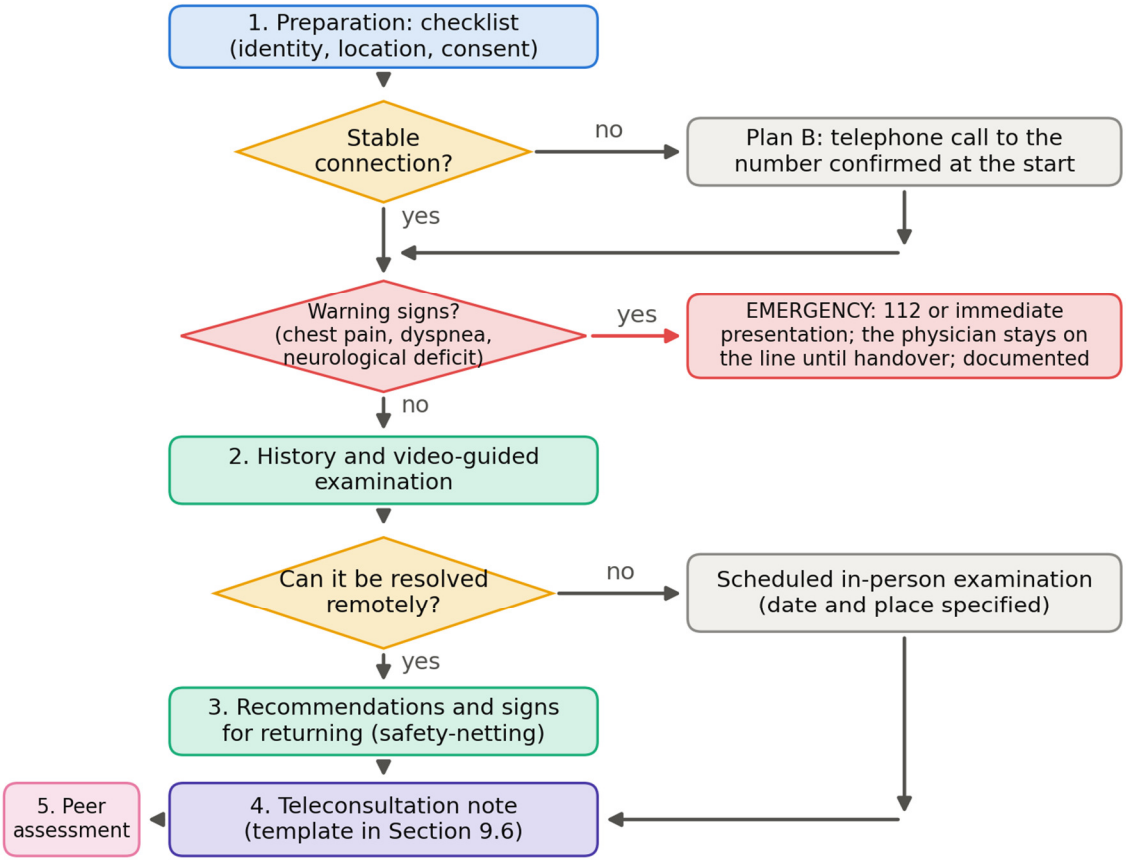


Figure L9.1. Workflow of the simulated teleconsultation, with the decision points: connection stability, warning signs, and whether the problem can be resolved remotely. Teaching diagram based on [334].

L9.4 Writing the Note (15 minutes)

15. Immediately after the consultation, the "physician" writes the teleconsultation note following the template in Section 9.6: identification, location, mode, consent, clinical data, limitations, assessment and plan, safety, and signature. If the simulation included a dropped connection, the technical incident is recorded, as the Romanian norms also require [364]. The note records what was actually observed, not what would have been seen in a complete examination.

16. The note contains no real data. Names are fictitious initials, and the personal numeric code (CNP) is replaced by the label "[fictitious]."
17. The roles are reversed and steps 5-16 are repeated with the second case.

L9.5 Peer Assessment (15 minutes)

18. The pairs exchange their notes. Each student assesses the fellow student's note with the grid below, ticking each item that is present and correct.
19. Peer assessment grid (one point per item):

- the identity of the patient and of the physician, the date, and the start and end times.
- the patient's location and the backup telephone number.
- the mode, the platform, and the persons present.
- consent to the remote consultation and the statement about recording.
- the clinical data, clearly separated into those reported by the patient and those observed by the physician.
- the limitations of the examination, stated concretely.
- the assessment and the plan, including the time frame for reassessment.
- the warning signs explained to the patient.
- the correct triage decision for the case received (remote continuation, in-person examination, or emergency).
- the absence of real identifying data and of unnecessary information.

20. The assessor writes two comments: one item done well and one to be corrected. The instruments developed to assess telemedicine competencies in simulation also use competency-based grids, completed by faculty members or by standardized patients, although their validation is still preliminary [573].
21. Final discussion with the whole group (5 minutes): which warning signs were missed, which limitations were left out of the notes, and what went wrong technically.

L9.6 The Link with Practical Session 10

The anonymized notes produced in this session are kept. In Practical Session 10, the students will ask a large language model (LLM) to summarize or structure a fictitious teleconsultation note and will check, with the same grid, which items the model omitted or invented. The exercise prepares a question they will face in practice: who is responsible for a note written with the help of an automated tool.

The competencies practiced here correspond to the domains of the international digital competency frameworks for medical education, which include digital health professionalism and health information systems [30], as well as to the telemedicine competencies proposed by Delphi consensus for postgraduate training, half of which concern systems-based practice and professionalism [574].

What to Submit

Each student submits, in PDF format, in a single file:

1. the completed checklist for the consultation in which the student played the role of the physician (Table L9.1).
2. the teleconsultation note written according to the template in Section 9.6.

3. the completed assessment grid for the fellow student's note, with the two comments.
4. a paragraph of 80-120 words about a limit of the remote consultation observed in the simulation.

Table L9.2. Indicative grading scheme for the practical examination (10 points).

Criterion	Points
Complete and correct checklist	2
Teleconsultation note: identification, location, mode, consent	2
Note: clinical data, limitations, assessment and plan	2
Note: safety-netting advice and correct triage decision	2
Well-argued peer assessment	1
Reflection paragraph	1

A wrong triage decision in the case with a warning sign (continuing the consultation remotely) leads to the loss of the 2 points for the fourth criterion, regardless of the quality of the rest of the note.

Offline Variant

If the video platform or the network is not working, the simulation is carried out face to face, with the students seated back to back. The "physician" is not allowed to turn around. This simulates a telephone teleconsultation, in which the image is also missing. The guided examination is reduced to what the patient can report, and the student must record this additional limitation in the note. The checklist, the note, and the grid are completed on paper, on the forms printed by the instructor, and are photographed or scanned for submission. The final discussion compares the two situations: what was lost when the image disappeared and which questions compensated for it.

PRACTICAL SESSION 10

Practical Session 10. Simulated Consultation with a Language Model: Verification, Errors, and Disclosure

Objective

The student uses a large language model (LLM) as a conversation partner in a simulated consultation on a fictitious case and treats every answer as a hypothesis to be verified. By the end of the session, each student should be able to write a structured instruction (prompt) for a clinical case, break the answer down into verifiable statements, check each statement against a guideline or a primary source, classify the errors according to a published grid, verify the references proposed by the model in PubMed, and write an AI-use disclosure statement that complies with the ICMJE recommendations.

The rationale for the session is a gap documented in Chapter 10. When tested alone on ten scenarios, three models identified a relevant condition in 94.9% of cases on average, but the people who used them succeeded in fewer than 34.5% of cases, with no advantage over the control group [403]. The difference arose in the conversation: users often gave incomplete information and did not always follow the correct suggestions. Performance also depends on the amount and order of the information received [402]. A student who knows how to present the case in full and how to check the answer is less exposed to these errors.

This session continues Practical Session 9. Students who kept the fictitious teleconsultation note from that session can use it as additional material at step L10.4.

Only fictitious cases, labeled as such, are used. No information about a real patient is entered into the chatbot, not even "just with initials", because data concerning health are a special category protected by Article 9 of the GDPR [118].

Tools

- **An LLM chatbot freely accessible through a browser.** An institutional account is preferred, if the faculty provides one. The session does not depend on any particular product. The student records the name of the tool, the model version shown in the interface, the date and time of the session, and whether the tool had access to web search. These are the details required by the disclosure statement at step L10.7.
- **PubMed** (https://pubmed.ncbi.nlm.nih.gov, accessed September 27, 2026), for checking references and searching for primary sources [39].
- **Free official guidelines**: World Health Organization guidelines approved by its Guidelines Review Committee [575] and the guidelines of European specialty societies, downloaded in PDF format from the official websites. Blogs, commercial summaries, or other AI-generated answers are not used as verification sources.
- **The ICMJE Recommendations** (updated in January 2026), the sections on AI tools [202].

- **The worksheet**: a table with the columns "statement", "verification source", "verdict", "error category", and "severity", received from the instructor in an editable format (Google Sheets or Excel).

Procedure

The total duration is about 110 minutes. Students work individually, and the final discussion involves the whole group.

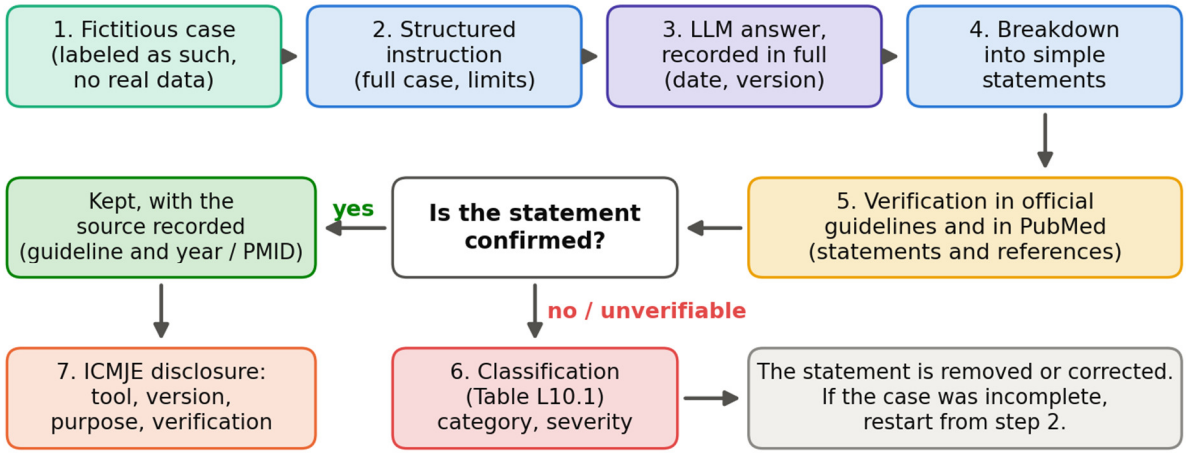


Figure L10.1. The verification loop for a language model's answer in the practical session: from the fictitious case and the structured instruction, through decomposition into statements, verification against sources, and error classification, to the disclosure statement. Any unconfirmed statement goes back into the loop or is removed. Teaching diagram based on [202,406].

L10.1 Choosing the Case (5 Minutes)

The instructor hands out one of the fictitious cases below. Each case contains at least one element that changes management and that can be missed.

- **Case A (fictitious).** A 71-year-old woman with hypertension and diabetes whose smartwatch has shown her an irregular rhythm notification three times in the past week. She has intermittent palpitations, without syncope. Her question: "Do I need to do something, or is it just the watch?"
- **Case B (fictitious).** A 45-year-old man with a headache of sudden onset during exertion, which he describes as "the worst of my life", accompanied by nausea. He asks which painkiller to take.
- **Case C (fictitious).** A 20-year-old female student with a sore throat and fever for two days, without cough. She asks for an antibiotic "so it goes away faster" before her exam session.

L10.2 The Structured Instruction (15 Minutes)

The student writes an instruction in six parts: the role assigned to the model, the context, the complete case, the task, the answer format, and the limits. Vague instructions produce vague answers, and the model fills in missing information with assumptions. For this reason, the case is transcribed in full, with age, sex, duration of symptoms, comorbidities, and medication.

Example

- Role: "Answer as a family physician explaining to a first-year student."
- Context: "The case is fictitious and serves a teaching exercise."
- Case: the full text of the case, without any real personal data.
- Task: "Which diagnoses should be considered, which red flags should be looked for, and what management do you recommend?"
- Format: "Number the statements. For each statement, indicate the source (guideline, year) and, for articles, the PMID or DOI."
- Limits: "State explicitly what you cannot assess without a clinical examination and when the patient must seek urgent medical care."

The instruction is saved exactly as it was sent. In the assessment, what counts is whether it explicitly asked for red flags and limits.

L10.3 Recording the Answer (10 Minutes)

1. A new conversation is opened, with no history, and the instruction is sent only once.
2. The answer is copied in full into the worksheet, with the date, time, tool name, and model version.
3. A second question is then sent, containing a false premise deliberately introduced by the instructor (for example, an incorrect statement about the equivalence of two drugs). The student notes whether the model corrects the premise or accepts it. In a study of five leading models, compliance with illogical requests of this kind reached 100% [407].
4. The answer is not regenerated to obtain a "better" one. Variability between runs is discussed at the end, based on the different answers that classmates received for the same case.

L10.4 Breaking the Answer into Statements and Verifying Them (35 Minutes)

1. The answer is split into simple statements, one per row in the worksheet. For example, "the watch notification confirms atrial fibrillation" is one statement. "An ECG is recommended" is another.
2. For each statement, the student looks for confirmation in an official guideline or in an article found in PubMed. The worksheet records the guideline title, the year, and the page or section, or the PMID of the article.
3. The verdict has three values: confirmed, refuted, or unverifiable. An unverifiable statement is not considered correct.
4. The answer is then compared with the case. Were all the data taken into account (age, comorbidities, duration of symptoms)? Is anything missing that would change management? Omissions are entered in the worksheet as separate rows, for example "did not ask about anticoagulant medication".
5. Students who brought the note from Practical Session 9 can ask the model to summarize it and can check, sentence by sentence, whether the summary contains anything that was not in the note or whether it left out something important. This procedure reproduces, on a small scale, the method used in a study evaluating automatically generated clinical notes, in which clinicians checked every sentence of the note against the transcript of the consultation [406].

L10.5 Classifying the Errors (15 Minutes)

Every refuted statement and every omission receives a category and a severity. The grid in Table L10.1 starts from the taxonomy proposed by Asgari and colleagues for LLM-generated clinical notes: hallucinations of the fabrication, negation, causality, or context type, and omissions, each classified as major or minor depending on whether it could change the diagnosis or management [406]. In that study, the hallucination rate was 1.47% of the 12,999 sentences in the notes, and the omission rate was 3.45% of the sentences in the transcripts [406]. For the consultation, the grid adds three teaching categories: outdated information, unsafe advice, and missed red flag. The last category has empirical support: a consumer triage tool undertriaged 52% of the reference emergencies in a structured test [404].

Table L10.1. Error classification grid, with fictitious examples. Source: adapted from [404,406,408].

Category	Working definition	Fictitious example	Typical severity
Omission	information that changes management is missing	does not ask about anticoagulant medication	major or minor
Fabrication	statement or source with no real counterpart	cites a "2024 study" with a nonexistent PMID	major
Negation	states the opposite of the case data	"the patient has no risk factors", although she has hypertension and diabetes	major
Causality	unsupported causal link	attributes the palpitations to coffee without evidence	minor or major
Context	information that is correct in another context	gives adult management for a pediatric case	major
Outdated information	recommendation from a superseded version	uses a risk score from an old guideline	minor or major
Unsafe advice	recommendation that can cause harm	suggests an antibiotic without an indication	major
Missed red flag	referral to emergency care is missing	sudden-onset headache referred for a checkup in a few days	major

In Case B, a recommendation to take a painkiller and wait, with no referral to the emergency department, is classified as a missed red flag of major severity, however correct the rest of the answer may be.

L10.6 Verifying the References (15 Minutes)

1. If the model did not give any references, it is asked for five, with PMID or DOI, for the main statements.
2. Each reference is searched for in PubMed by PMID and, separately, by title [39].
3. Each reference receives one of four verdicts: correct (it exists and supports the statement), real but with incorrect details (authors, year, journal), real but unrelated to the statement, or nonexistent.
4. The percentage of nonexistent references in the student's own answer is calculated and entered in the worksheet.

The results are then compared with the literature. In a study of 636 references, 55% of those generated by GPT-3.5 and 18% of those generated by GPT-4 did not exist, and some of the real ones contained substantive errors [408]. In another test, the rates of hallucinated references were

39.6% for GPT-3.5, 28.6% for GPT-4, and 91.4% for Bard [44]. A 2026 audit of five popular chatbots found that none of them produced a completely correct reference list [405]. A rate of zero in one's own exercise does not prove that the tool is safe. It only shows that, this time, the references withstood verification.

L10.7 The Disclosure Statement (10 Minutes)

The ICMJE Recommendations require the use of AI tools to be disclosed. Writing assistance is disclosed in the acknowledgments section, whereas use for data collection or analysis or for figures is described in the methods, in enough detail to allow reproduction, including the tool, the version, and the instructions given. Chatbots cannot be authors, humans are responsible for all of the material, and AI-generated material is not cited as a primary source [202]. The student writes a disclosure statement of 60-100 words for the report. A possible model:

Example

- "For this report I used [tool name], version [version shown], on [date], to generate an answer to a fictitious clinical case, based on the instruction in Appendix 1."
- "All statements and references in the answer were verified against official guidelines and in PubMed. The results of the verification are presented in Appendix 2."
- "The text of the report was written by the author, without AI writing assistance. The author takes full responsibility for the content."

L10.8 Reflection and Final Discussion (10 Minutes)

Each student writes 100-150 words on three questions. Which error could have harmed a real patient? Which piece of information in the case changed the model's answer the most? Would the student have noticed the error without verification? The group discussion compares the different answers received for the same case and the number of nonexistent references. The instructor links the conclusions to the WHO guidance on large multimodal models, which warns about false or incomplete statements and about automation bias, that is, the user's tendency to stop noticing the system's mistakes [410].

What to Submit

Each student submits a single PDF file containing:

1. The structured instruction, exactly as it was sent.
2. The model's full answer, with the date, time, tool, and version.
3. The verification worksheet: the statements, the source of each verification, the verdict, and the category and severity of the error.
4. The reference verification table, with the percentage of nonexistent references.
5. The AI-use disclosure statement.
6. The 100-150-word reflection.

Table L10.2. Indicative grading scheme for the practical examination (10 points).

Criterion	Points
Complete structured instruction (full case, red flags, limits)	1.5
Correct breakdown into statements	1.5
Verification with identifiable official sources (guideline and year, or PMID)	2
Classification of errors and severity according to the grid	2
Verification of references in PubMed	1.5
ICMJE-compliant disclosure statement	1
Reflection	0.5

Entering any real patient data into the chatbot results in the cancellation of the submission. A statement marked "confirmed" without an identifiable source is scored as unverified.

Offline Variant

If the network or the chatbot is unavailable, the instructor hands out, on paper, three answers previously generated by an LLM for Cases A, B, and C, together with the instructions used, the date, and the model version. Printed excerpts from the relevant guidelines are also handed out, along with the PubMed records of the references cited by the model, mixed with references that do not exist in PubMed. Students work through steps L10.4-L10.8 on these materials. The references are verified by comparing the generated list with the printed records. Titles and PMIDs that do not appear in the set of records are marked "to be checked" and searched for in PubMed at the first opportunity with network access. The disclosure statement is written for a hypothetical report that would use the answer received.

PRACTICAL SESSION 11

Practical Session 11. Practical Security: Passwords, Multi-Factor Authentication, Verifying an Electronic Signature, and Recognizing Phishing

Objective

By the end of the session, the student will be able to calculate the entropy of a randomly generated password and compare it with current recommendations, enable 2-Step Verification on his or her own Google account, sign and verify a fictitious PDF document with a self-signed certificate, read the TLS certificate of a website, recognize the red flags in a phishing e-mail, and describe the steps for reporting an incident. The session applies the concepts from Chapter 11 and the entropy formula from Chapter 2.

Tools

- A pocket calculator or Google Sheets (the function `LOG(x, 2)` for the base-2 logarithm).
- The student's personal Google account, which the student administers alone. The official help page for 2-Step Verification: https://support.google.com/accounts/answer/185839, accessed September 27, 2026 [533].
- Adobe Acrobat Reader, free, Continuous track version 26.002.21931 (the update of September 17, 2026, according to the release notes at https://www.adobe.com/devnet-docs/acrobatetk/tools/ReleaseNotesDC/index.html, accessed September 27, 2026). Download: https://get.adobe.com/reader/. Help pages used: creating a self-signed digital ID [576] and validating signatures [577].
- Google Chrome 154 (promoted to the stable channel on September 22, 2026, according to https://chromereleases.googleblog.com, accessed September 27, 2026) or Mozilla Firefox, together with the Chrome help page on security symbols [578].
- The European Commission's trusted list browser, https://eidas.ec.europa.eu/efda/trust-services/browse/eidas/tls, accessed September 27, 2026 [579].

No screenshots are taken that contain the student's e-mail address, phone number, or backup codes. All documents and messages used in the session are fictitious.

Procedure

The session has six parts, of 10-35 minutes each, which are completed in order. Part L11.3 is assessed in the greatest detail.

L11.1 How Strong Is a Password? (20 Minutes)

For a password chosen truly at random, each character is an event with N equally probable outcomes. From Chapter 2, we know that such an event carries log2 N bits. A password of L characters therefore has the entropy:

H = L · log2 N (bits)

1. Calculate log2 N for three alphabets: lowercase letters (N = 26, log2 26 ≈ 4.70); uppercase letters, lowercase letters, and digits (N = 62, log2 62 ≈ 5.95); and all 94 printable characters on the keyboard (log2 94 ≈ 6.55).
2. Record in your notebook the entropy for: 8 lowercase letters (8 · 4.70 ≈ 37.6 bits), 8 alphanumeric characters (≈ 47.6 bits), 15 lowercase letters (≈ 70.5 bits), and 12 characters out of 94 (≈ 78.7 bits).
3. A passphrase made of words chosen at random from a list of 7,776 words carries log2 7776 ≈ 12.9 bits per word. Calculate the entropy for 5 and for 6 words (≈ 64.6 and ≈ 77.5 bits).
4. Estimate the average time needed to guess the password by trial, under a teaching assumption of 10^10 guesses per second: t = 2^H / 2 / 10^10 seconds. The approximate results are 10 seconds for 8 lowercase letters, about 3 hours for 8 alphanumeric characters, about 45 years for 5 words, and about 2,700 years for 15 lowercase letters.
5. Compare these results with the NIST SP 800-63B-4 guideline: a password used as a single factor is at least 15 characters long, and one used together with a second factor at least 8. The guideline no longer requires mandatory combinations of character types or periodic changes, but it does require the password to be checked against a list of compromised or commonly used values [445].
6. Write two concluding sentences. Note that length increases entropy more than complexity does and that the formula holds only for random selection. A password invented by a person ("Hospital2026!") has a real entropy much lower than its number of characters suggests, because it follows predictable patterns.

L11.2 2-Step Verification on the Google Account (15 Minutes)

The steps follow the official Google help page [533].

1. Open your Google Account (myaccount.google.com) and select the security and sign-in section.
2. In the section on how you sign in to Google, choose to turn on 2-Step Verification and follow the on-screen instructions.
3. As the second step, prefer Google prompts on your phone or a passkey. The help page states that prompts help protect the account against a takeover of the phone number (SIM swap) [533].
4. Generate the backup codes and keep them offline, in a safe place. Do not photograph them or send them to anyone.
5. If you use a university or work account, the steps may differ, and the Google page recommends contacting the administrator [533].
6. Record in your notebook, without personal data, which verification methods the account offered and which one you chose. Link your choice to the NIST recommendation that

applications offer a phishing-resistant option, while the SMS code remains a restricted authenticator [445].

L11.3 Signing and Verifying a Fictitious PDF in Acrobat Reader (35 Minutes)

The session uses a self-signed certificate created by the student. It is sufficient for observing the mechanism, and the Adobe page recommends it for personal use or for internal processes without a certification authority [576]. However, it does not have the value of a qualified signature, because it is not based on a qualified certificate issued by a trust service provider [433].

1. In Google Docs or Word, create a one-page "medical letter" for a fictitious patient (for example, "Test Patient, 70 years old"), with a diagnosis, a treatment, and the dose of a medication. Export it as a PDF.
2. In Acrobat Reader, create the self-signed digital ID: Menu > Preferences > Signatures, then More under Identities & Trusted Certificates, Digital IDs, the add button, and the option "A new digital ID I want to create now" [576]. Choose the 2048-bit RSA algorithm, not the 1024-bit one (the FIPS 186-5 standard requires a modulus of at least 2048 bits for RSA) [427]. Protect the ID with a new password.
3. Sign the PDF with the Use a certificate tool, using the digital signature option (the exact name may vary between versions), and save a signed copy.
4. Open the Signatures panel and run the validation (Options > Validate Signatures). Open Signature Properties, then Show Signer's Certificate [577]. Record the holder's name, the issuer (identical to the holder, because the certificate is self-signed), the validity period, the algorithm, and the status displayed.
5. Look at the status displayed. The signer's authenticity is confirmed by checking whether the signer's certificate or its parent certificates are trusted [577]. A self-signed certificate does not lead to a certification authority, so on the computer of someone who has not marked it as trusted, the document may appear intact while the signer's identity remains unconfirmed.
6. Add a comment or a note to the signed document, save it, and validate it again. Record what Acrobat reports about the changes made after signing.
7. Check the time stamp. If the signature has no time stamp from a trusted server, the time comes from the clock of the signer's computer [577].
8. Work through the diagram in Figure L11.1 on paper, for your own signature and for a hypothetical scenario: a letter signed by a physician with a qualified certificate issued by a provider on the trusted list of an EU Member State. In the trusted list browser [579], search for Romania and record how many qualified trust service providers are listed on the date of the session. The connection to the legal text: only a qualified electronic signature has a legal effect equivalent to that of a handwritten signature [433].

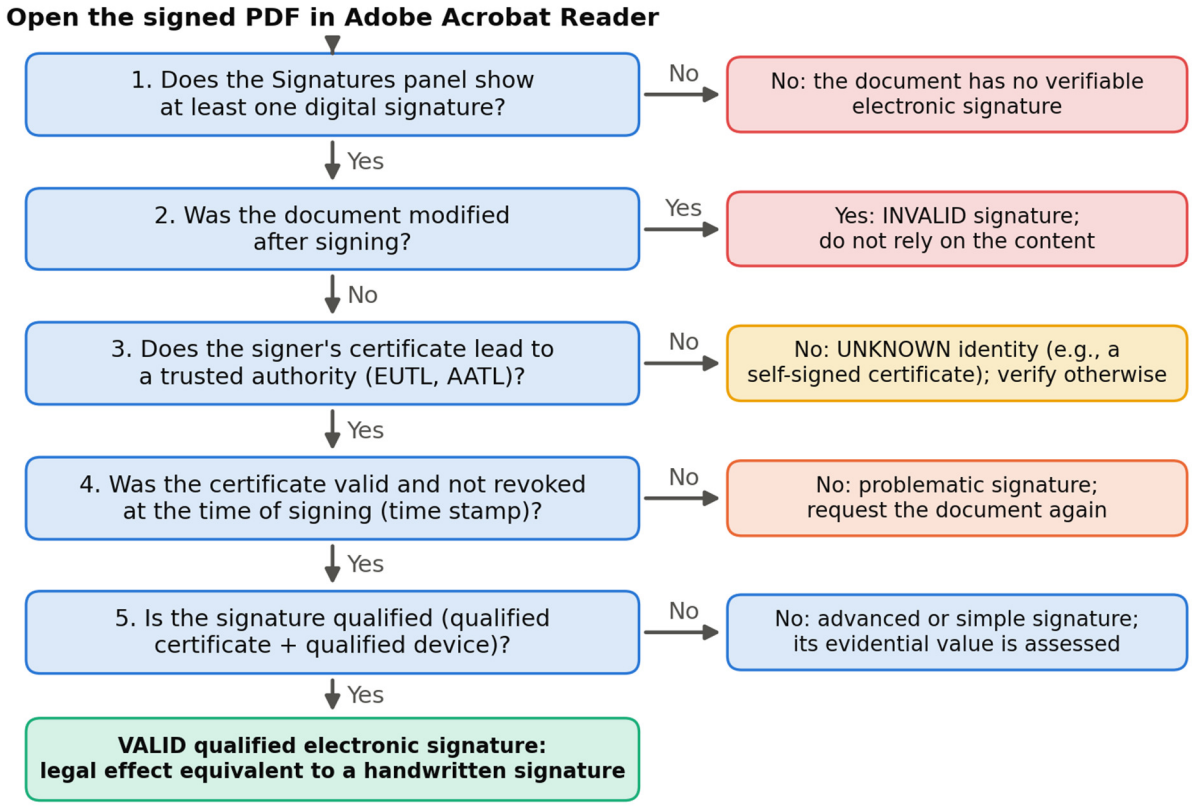


Figure L11.1. Decision diagram for verifying an electronic signature on a PDF in Acrobat Reader. Adapted from [433,577].

L11.4 The TLS Certificate of a Website (10 Minutes)

1. In Chrome, open the public website of an institution, for example https://cnas.ro.
2. Look at the symbol to the left of the address. The Chrome help page explains what the symbols for a secure, not secure, or dangerous connection mean [578].
3. Click the symbol, then the connection information, and then the certificate. Record the name to which the certificate was issued, the issuer, the validity period, and the certification chain up to the root authority.
4. Compare the name in the certificate with the address in the browser bar. Discuss in the group why a padlock on display does not prove that the website belongs to the institution it imitates (a phishing website can also have a valid certificate).

L11.5 Recognizing Phishing: A Test with Fictitious E-mails (25 Minutes)

The instructor projects or hands out the five messages below. All of them are fictitious, and the domains end in .example, a reserved suffix that belongs to no one. Students work individually for 10 minutes and mark each message as "phishing" or "legitimate", using the red flags in Table L11.1. A discussion follows.

1. From "Health Insurance House - Verification" <alerta@cas-verificare.example>. Subject: "Your health insurance card will be suspended in 24 hours". The text asks the recipient to confirm personal data and the PIN code through a shortened link.

2. From "IT Department" <it-support@spital-judetean.example>, the usual internal address. Subject: "Scheduled maintenance on Saturday, 02:00-04:00". It contains no link and asks for no data; it only announces that the system will be unavailable.
3. From "Prof. Dr. X, Head of Clinic" <sef.clinica.x@mail-gratuit.example>. Subject: "Urgent and confidential". It asks the recipient to buy five gift cards and send the codes, without a phone call, because "I am in a meeting".
4. From "Human Resources" <hr@spita1-judetean.example> (the digit 1 instead of the letter l). The attachment "Grila_salarii_2026.xlsm" (salary grid) asks for macros to be enabled.
5. From the faculty's e-learning platform, from the correct address. A message about a new assignment. The displayed link and the real link (seen when hovering the mouse over it) match and lead to the faculty's domain.

Answers for the instructor: 1, 3, and 4 are phishing; 2 and 5 are legitimate. Message 3 illustrates fraud through impersonation of a superior, and message 4 uses a lookalike domain. In an Italian hospital with more than 6,000 employees, messages personalized to the local context were accessed far more often than generic ones (55% versus 7% of the messages sent, in the first campaign) [443]. This is why messages 3 and 4 are the most dangerous.

Table L11.1. Red flags in a phishing message. Source: teaching synthesis; the messages in the test are fictitious.

Red flag	What to check	Example in the test
Urgency or threat	very short deadlines, sanctions, "final notice"	message 1
Request for secret data	passwords, PINs, codes received by SMS	message 1
Suspicious sender	free e-mail domain or one different from the official domain	message 3
Spoofed domain	substituted letters (1 instead of l, rn instead of m)	message 4
Hidden link	the real address differs from the displayed text, shortened link	message 1
Risky attachment	files with macros, archives, executables	message 4
Authority and secrecy	a superior asks for a payment and forbids verification	message 3
Unusual context	a request that does not fit the procedures	messages 1, 3, 4

L11.6 The Incident Reporting Procedure (10 Minutes)

Working in groups of three, students write a procedure of no more than 10 lines for the ward where they do their clinical training, starting from the scenario: "A resident entered their password on the page from message 1". The procedure must include:

1. What not to do: do not click the link again, do not forward the message to colleagues, and do not delete the message before reporting it.
2. What to do immediately: change the password from another device, notify the hospital's IT or security officer according to the internal procedure, and disconnect the suspect device from the network if IT requests it.

3. Who else must be notified: the data protection officer, if patient data may have been exposed (Chapter 5).
4. At the national level, cybersecurity incidents can be reported to the Romanian National Cyber Security Directorate (Directoratul Național de Securitate Cibernetică, DNSC) through the PNRISC platform or by calling 1911, available 24 hours a day [447].
5. What is learned: the ENISA guide for the health sector recommends an incident response process with defined roles and knowledge of the contact person, and, for staff, phishing simulations and role-based training [446].

What to Submit

Each student submits a PDF file named "L11_LastName_FirstName.pdf", which contains:

1. The entropy table from L11.1, with the calculations and the conclusion (2 points).
2. Written confirmation, without screenshots containing personal data, that 2-Step Verification is active, with the method chosen and its justification (1 point).
3. The signed fictitious PDF, plus the observation sheet from L11.3: the certificate details, the validation status before and after the modification, and an explanation of why a signature with a self-signed certificate is not a qualified signature (3 points).
4. The TLS certificate details from L11.4 (1 point).
5. The answers to the phishing test, with at least one red flag from Table L11.1 for each message (2 points).
6. The reporting procedure from L11.6 (1 point).

Total: 10 points. In the practical examination, the student may be asked to calculate the entropy of a given password or to interpret the status of a signature in a PDF.

Offline Variant

If there is no internet access, the calculations in L11.1 and the phishing test in L11.5 are done entirely on paper. For L11.3, a previously installed copy of Acrobat Reader allows the self-signed ID to be created and the document to be signed and validated offline. For L11.2 and L11.4, the instructor presents the steps based on the diagram in Figure L11.1 and on a TLS certificate exported in advance as a file, which students open and analyze locally. The procedure in L11.6 is written on paper.

PRACTICAL SESSION 12

Practical Session 12. Information Flow in the Physician's Office, Outpatient Clinic, and Hospital: Modeling and Indicators

Objective

The student models the path of a fictitious patient from the family physician to the reimbursement of services, in the form of a swimlane diagram. For each step, the student identifies the document issued and the minimum data elements, then calculates, in a spreadsheet, four hospital indicators on a fictitious data set: average length of stay (ALOS), bed occupancy rate, case-mix index (CMI), and 30-day readmission rate. The session applies the concepts from Chapter 12 (Sections 12.4 and 12.6).

Tools

- **draw.io** (also called diagrams.net), a free and open-source program (Apache 2.0 license) that can be used without an account. The web version opens at https://app.diagrams.net, and diagrams can be saved on the user's own computer or in a storage location of the user's choice; the developer states that it has no access to the saved data [580]. The desktop version for Windows, macOS, and Linux is downloaded from the official releases page. As of September 27, 2026, the latest version was 31.5.3, released on September 26, 2026 [581]. The official instructions for swimlane diagrams are in the draw.io documentation [582].
- **Google Sheets** (in the browser, with a Google account) or **Microsoft Excel**. The functions used are SUM, COUNT, COUNTA, AVERAGE, AVERAGEIF, COUNTIF, COUNTIFS, VLOOKUP, and SUMPRODUCT. Their official descriptions are in the Google Sheets function list [583] and in the Microsoft documentation [584], all accessed on September 27, 2026.
- **The fictitious data set** in Table L12.2, copied by hand or imported from a CSV file provided by the instructor. No data from real patients are used.

Working with spreadsheet formulas follows the protocol in [6], revisited and updated for Google Sheets/Excel 2026.

Procedure

The session has two parts. The first (L12.1-L12.3) builds the model of the information flow. The second (L12.4-L12.6) calculates and interprets the indicators. Figure L12.1 shows the expected result of the first part.

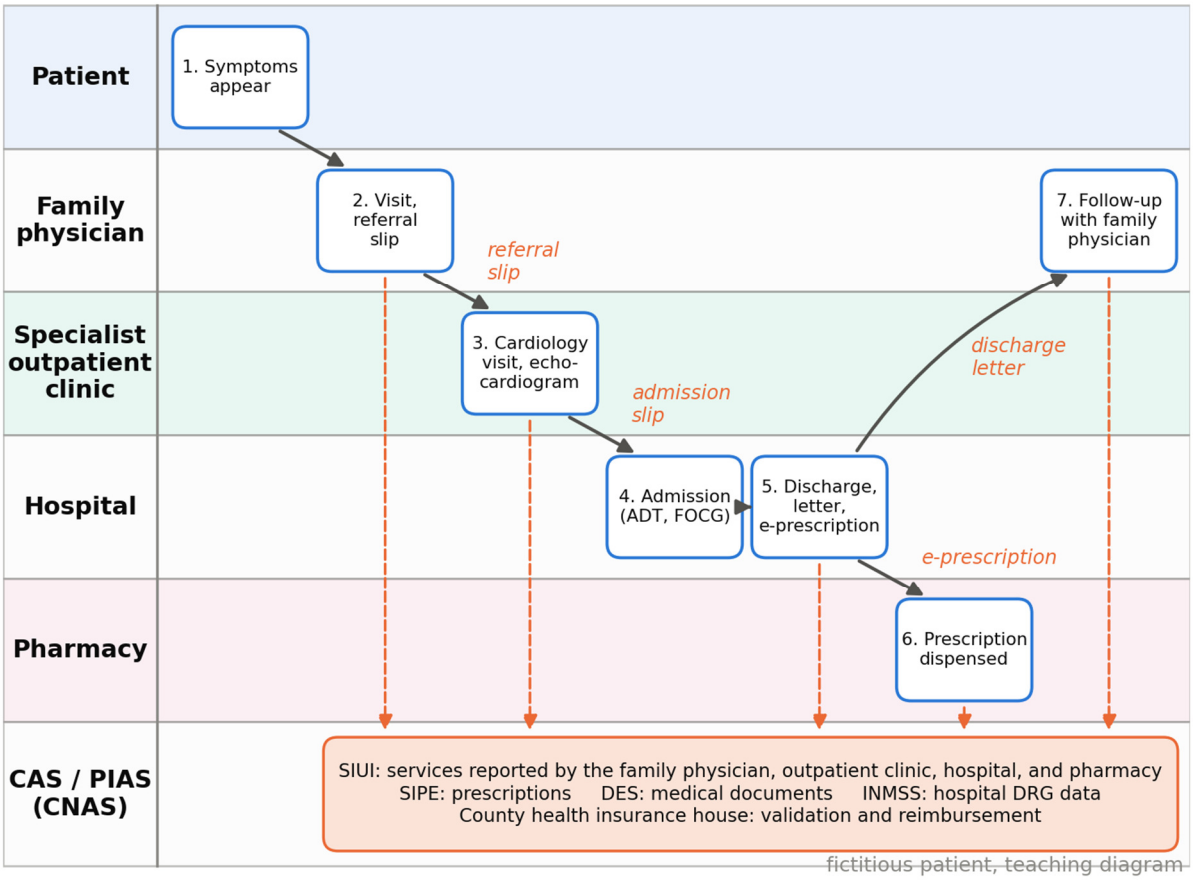


Figure L12.1. Swimlane diagram of the path of a fictitious patient: family physician, specialist outpatient clinic, hospital, pharmacy, and the health insurance platform. Solid arrows show the patient's path and the documents that accompany the patient, and dashed arrows show the reporting to SIUI, SIPE, DES, and INMSS. Teaching diagram.

L12.1 The Scenario

The fictitious patient P.X., aged 64, presents to the family physician with exertional dyspnea that began during the past month. The physician examines him, measures his blood pressure, records an electrocardiogram, and refers him to a cardiologist with a referral slip. In the specialist outpatient clinic, the cardiologist performs an echocardiogram, suspects decompensated heart failure, and issues an admission slip. The patient is admitted for 5 days to the cardiology ward. At discharge, he receives the discharge letter and an e-prescription. The prescription is dispensed at the pharmacy, and two weeks later the patient returns to the family physician for a follow-up visit. Each provider reports its services to the health insurance platform, and the hospital also submits the minimum data set for DRG grouping.

Process mapping helps teams understand a complex system before changing it. A systematic review of 105 studies found, however, that none met all ten quality criteria proposed by the authors and that only 15% of the studies involved patients in producing the map [585]. A widely adopted language is BPMN (Business Process Model and Notation), which has advantages for standardizing workflows but also limitations in highly complex clinical settings [586]. The swimlane diagram used here is a simplified form, sufficient for a first model.

L12.2 The Swimlane Diagram in draw.io

1. Open https://app.diagrams.net. When asked where to save, choose *Device* (your own computer). Create a new diagram and, from the *Basic* template collection, choose *Swimlane Diagram* [582]. Save the file immediately under the name `L12_LastName_FirstName.drawio`.
2. The template contains a container (*pool*) with a few lanes. To add lanes, select a lane and press Alt+Enter or, on the *Arrange* tab, use the *Insert Row Above* or *Insert Row After* commands [582]. You need six lanes.
3. Rename the lanes by double-clicking the title of each one: Patient, Family Physician, Specialist Outpatient Clinic, Hospital, Pharmacy, CAS/PIAS (the county health insurance house and the health insurance IT platform) [582].
4. Drag a rounded rectangle from the *General* library for each step and place it in the lane of the actor who performs it. Number the steps from 1 to 7, as in Figure L12.1. A step belongs in the lane of whoever performs the action, not of whoever requests it.
5. Connect the steps with arrows. Hover the mouse over a shape, drag a connector from one of the direction arrows, and drop it on the target shape [582]. On each arrow, write (by double-clicking) the document that travels along it: referral slip, admission slip, e-prescription, discharge letter.
6. Add dashed arrows from steps 2, 3, 5, 6, and 7 to the CAS/PIAS lane. The dashed style is selected in the *Style* panel, under the line option. On these arrows, write the system that receives the data: SIUI, SIPE, DES, or INMSS (Chapter 12, Table 12.2).
7. Check the diagram against three rules. Each step has a single actor. Each arrow has a direction. No information appears "out of nowhere": any document used at a step was produced at an earlier step.
8. Export the diagram via *File* > *Export as* > *PNG*, at 200% zoom, and also keep the `.drawio` file, which remains editable.

L12.3 Documents and Data Elements

For the scenario in L12.1, fill in a table like Table L12.1. The first rows are given as a model; the student checks them and completes the rest. The "Minimum data elements" column contains only what is needed for the next step or for reporting. For example, the pharmacy does not need the echocardiogram, but it does need the diagnosis code that justifies the reimbursement.

Table L12.1. Documents and minimum data elements at each step of the path (fictitious patient). Source: teaching construct, with step 5 based on [139,465].

Step	Actor and system	Document issued	Minimum data elements
2. Family physician visit	family physician, reporting to SIUI	consultation record, referral slip	patient identification, date, coded presumptive diagnosis, requested specialty, issuing physician
3. Cardiology visit	outpatient clinic, reporting to SIUI	echocardiography result, admission slip	diagnosis, investigations performed, reason for admission
4. Admission	hospital, ADT module, FOCG (general clinical observation chart)	opened FOCG	date and time of admission, ward, bed, type of admission, admission diagnosis
5. Discharge	hospital, reporting to SIUI, INMSS, SIPE, DES	discharge letter, e-prescription, minimum data set	principal and secondary diagnoses (ICD-10-AM), procedures, discharge date, condition at discharge
6. Prescription dispensing	pharmacy, SIPE	dispensed prescription	medicine, quantity, dispensing date
7. Family physician follow-up	family physician	discharge letter attached to the record	discharge recommendations, current treatment

At step 5, diagnoses are coded in ICD-10-AM, 3rd edition, and DRG grouping is performed with the RO DRG v.1.1 classification [139]. The patient-level minimum data set is collected by INMSS from all hospitals [465].

L12.4 The Data Set and the Spreadsheet

The fictitious cardiology ward has 4 beds. In April 2026 (30 days), it discharged 16 patients, all admitted and discharged in the same month. The relative weights of the four DRG groups are **fictitious** and chosen for teaching purposes; they are not the official RO DRG values.

Table L12.2. Fictitious data set of the ward (April 2026). Fictitious DRG weights: A = 0.85 (chest pain without complications), B = 1.20 (heart failure), C = 1.65 (myocardial infarction), D = 2.40 (pacemaker implantation). Fictitious, illustrative data.

ID	Admission	Discharge	DRG group	Readmission within 30 days
P01	01.04.2026	07.04.2026	B	No
P02	01.04.2026	04.04.2026	A	No
P03	02.04.2026	12.04.2026	C	Yes, unplanned
P04	03.04.2026	08.04.2026	B	No
P05	04.04.2026	06.04.2026	A	No
P06	06.04.2026	15.04.2026	D	No
P07	08.04.2026	14.04.2026	B	No
P08	09.04.2026	12.04.2026	A	No
P09	12.04.2026	19.04.2026	C	No
P10	13.04.2026	17.04.2026	B	Yes, planned
P11	15.04.2026	23.04.2026	D	No
P12	16.04.2026	18.04.2026	A	No
P13	18.04.2026	26.04.2026	C	No
P14	20.04.2026	25.04.2026	B	No
P15	23.04.2026	29.04.2026	B	No
P16	24.04.2026	30.04.2026	C	No

Organize the spreadsheet as follows:

1. In row 1, enter the headers: A = ID, B = Admission, C = Discharge, D = Days, E = DRG, F = Weight, G = Readmission (Yes/No), H = Readmission type (planned, unplanned, or blank). Enter the 16 records in rows 2-17. Dates are entered in the regional format of the spreadsheet (dd.mm.yyyy for a Romanian locale, mm/dd/yyyy for a US English locale).
2. In a separate area, build the weights table: M1 = DRG, N1 = Weight, O1 = Cases, P1 = ALOS by group. In M2:M5 enter A, B, C, D, and in N2:N5 the weights 0.85, 1.20, 1.65, and 2.40. In K1 enter "Beds" and in K2 the value 4. In K3 enter 30, the number of days in the month.
3. Days of hospitalization: in D2 enter `=C2-B2` and copy the formula down to D17. The day of discharge is not counted: Eurostat calculates days of hospitalization as the difference between the discharge date and the admission date [474]. If the result is displayed as a date, change the cell format to a number format.
4. Weight of each case: in F2 enter `=VLOOKUP(E2,$M$2:$N$5,2,FALSE)` and copy down to F17. The last argument, FALSE, requests an exact match (0 is equivalent and also works in any interface language). The `$` symbol locks the weights table when the formula is copied.

5. Number of cases per group: in O2 enter `=COUNTIF($E$2:$E$17,M2)` and copy down to O5. The sum O2:O5 must be 16.
6. ALOS by group: in P2 enter `=AVERAGEIF($E$2:$E$17,M2,$D$2:$D$17)` and copy down to P5.

In spreadsheets set to an English (US) locale, as in the formulas above, the argument separator is the comma and the decimal separator is the point. In spreadsheets configured for Romania, the argument separator is the semicolon and the decimal separator is the comma. In Excel with a Romanian-language interface, function names may also appear translated.

L12.5 Calculating the Indicators

Enter the formulas below in a results area (for example, column R) and compare the results with the control values in Table L12.3.

1. **Average length of stay** is the total number of days of hospitalization divided by the number of discharges [474]: `=SUM(D2:D17)/COUNT(D2:D17)`.
2. **Bed occupancy rate** is the ratio of the days of hospitalization actually recorded to the available bed-days, multiplied by 100: `=SUM(D2:D17)/(K2*K3)*100`.
3. **Case-mix index** is the mean of the relative weights of the discharged cases. Calculate it in two ways, which must give the same result: `=AVERAGE(F2:F17)` and `=SUMPRODUCT(N2:N5,O2:O5)/SUM(O2:O5)` [584].
4. **30-day readmission rate**, overall: `=COUNTIF(G2:G17,"Yes")/COUNTA(A2:A17)*100`.
5. **Unplanned readmission rate**: `=COUNTIFS(G2:G17,"Yes",H2:H17,"unplanned")/COUNTA(A2:A17)*100` [583].
6. Build a column chart of ALOS by group (P2:P5), with the DRG groups on the horizontal axis, and add it to the submitted file.

Table L12.3. Calculated indicators and control values (fictitious data).

Indicator	Formula (Google Sheets)	Control value
Days of hospitalization, total	=SUM(D2:D17)	90
Average length of stay	=SUM(D2:D17)/COUNT(D2:D17)	5.625 days (5.6)
Occupancy rate	=SUM(D2:D17)/(K2*K3*)100	75%
Case-mix index	=AVERAGE(F2:F17)	1.375
30-day readmissions, total	=COUNTIF(G2:G17,"Yes")/COUNTA(A2:A17)*100	12.5%
Unplanned readmissions	=COUNTIFS(G2:G17,"Yes",H2:H17,"unplanned")/COUNTA(A2:A17)*100	6.25%
ALOS for groups A, B, C, D	=AVERAGEIF(...)	2.5 / 5.33 / 7.75 / 8.5 days

If a result differs, first check the date format and the absolute references. A date stored as text (for example, with a trailing space) produces an error in the subtraction. Other frequent mistakes: counting the day of discharge, which would give 106 days instead of 90 and an ALOS of 6.6 days; dividing by the number of beds without multiplying by the days in the month; and typing "Yes " with a trailing space, which the COUNTIF function no longer recognizes. For each formula, calculate one case by hand and compare it with the cell.

L12.6 Interpretation

Answer the following questions in writing, in 8-10 lines:

1. Is an occupancy rate of 75% high or low for a cardiology ward? What additional information would be needed to decide whether the number of beds is appropriate?
2. The CMI is 1.375. What does it mean that it is above 1, and what would happen to the index if, in the following month, the ward admitted more cases from group A?
3. ALOS increases from group A to group D. Is this an expected observation? What would a long ALOS signal in a group with a low weight?
4. The overall readmission rate is 12.5%, but the rate of unplanned readmissions is 6.25%. Which of the two says something about the quality of care? A study of 8,878 neurosurgical patients concluded that the 30-day readmission rate, taken alone, is not an adequate measure of quality, because it lacks a uniform definition and requires risk adjustment [477].
5. Why can the ALOS of this ward, calculated for one month and 16 cases, not be compared with the range of 5.2-9.4 days reported for most EU Member States in 2023 [474]?
6. On the diagram from L12.2, indicate the step at which an error in diagnosis coding would change the CMI value calculated by the hospital.

What to Submit

Each student uploads to the course platform a single archive file named `L12_LastName_FirstName.zip`, which contains:

- the swimlane diagram, in `.drawio` and `.png` format (30 points: 6 correctly named lanes, 7 steps placed in the correct lane, the documents written on the arrows, the reporting to CAS/PIAS).
- the table of documents and data elements, completed for all steps (20 points).
- the spreadsheet with the formulas visible (not only the values) and the chart of ALOS by group (30 points: 5 points for each correctly calculated indicator and 5 points for the chart).
- the answers to the interpretation questions (20 points).

In the practical examination, the student may receive a different fictitious data set (a different number of beds or different weights) and must recalculate the indicators within 15 minutes.

Offline Variant

If the laboratory has no internet access, the diagram is drawn in the draw.io desktop application, installed beforehand from the official releases page [581]. The desktop application saves files locally and works without a connection. If this is not available either, the diagram is drawn on an A4 sheet in landscape orientation, with six bands drawn with a ruler, and the student photographs it for submission.

The indicators are calculated in any locally installed spreadsheet program, with the same formulas. Without a computer, the calculation is done by hand on Table L12.2. Add up the 16 lengths of stay (90 days), then divide by 16 for the ALOS and by 4 x 30 = 120 bed-days for the occupancy rate; for the CMI, add up the weights (4 x 0.85 + 6 x 1.20 + 4 x 1.65 + 2 x 2.40 = 22.00) and divide by 16.

Practical Session 13. Critical Appraisal of an AI Tool: From the Confusion Matrix to Debate

Objective

The student calculates the performance indicators of a fictitious AI tool, observes how the predictive values change with prevalence, builds a ROC curve by hand, applies a short checklist derived from TRIPOD+AI to a real article, and defends, in a role-based debate, a reasoned position on responsibility for an imaging report. The session applies the concepts from Chapter 13 (Sections 13.3-13.7). The protocol for calculating diagnostic indicators in [6] is revisited and updated for Google Sheets/Excel 2026.

Tools

- **Google Sheets**, a free web application used with a Google account, accessed on September 27, 2026 (https://sheets.google.com). The functions `SUM`, `COUNTIF`, `ROUND`, `MAX`, and `SUMPRODUCT` are used [583], together with the scatter and line chart types [587].
- **LibreOffice Calc 26.8.0**, a free alternative, installed locally, for the offline variant [588]. Microsoft Excel can be used if available, with the same formulas.
- **The article to be analyzed**: Wong et al., the prospective multicenter validation of the second version of the Epic sepsis prediction model, JAMA Network Open 2026, open access [503].
- **The TRIPOD+AI checklist** (BMJ 2024), in the abridged form in Table L13.3 [505].
- A pocket calculator and graph paper (for the offline variant).

All the data in Part A and Part B are fictitious, constructed for the exercise. No patient data are used.

Procedure

L13.1 Part A. The Confusion Matrix and the Effect of Prevalence (35 min)

The fictitious tool "PneumoScan" reads chest radiographs and flags pneumonia. It was tested on 1,000 radiographs, of which 100 had pneumonia confirmed by the reference standard (prevalence 10%). The results are in Table L13.1.

Table L13.1. Confusion matrix of the fictitious PneumoScan tool (n = 1,000, prevalence 10%) and the formulas for Google Sheets, with cells B2 = TP, C2 = FP, B3 = FN, C3 = TN. Fictitious data. Formulas based on [499,500].

	Pneumonia present	**Pneumonia absent**	**Indicator**	**Formula in Sheets**
PneumoScan positive	TP = 85	FP = 90	Sensitivity	=B2/(B2+B3)
PneumoScan negative	FN = 15	TN = 810	Specificity	=C3/(C3+C2)
			PPV	=B2/(B2+C2)
			NPV	=C3/(C3+B3)
			Accuracy	=(B2+C3)/SUM(B2:C3)
			F1 score	=2*B2/(*2B2+C2+B3)

Steps:

1. Open a new sheet and name it `L13_NameSurname`. In cells A1:C3, build the matrix: row 1 with the headers "Pneumonia present" and "Pneumonia absent"; column A with "Positive" and "Negative"; then the values 85, 90, 15, 810 in B2, C2, B3, C3.
2. In column E, enter the names of the indicators, and in column F, the formulas from Table L13.1. Format column F as a percentage with one decimal place.
3. Check the results: Se = 85.0%, Sp = 90.0%, PPV = 48.6%, NPV = 98.2%, accuracy = 89.5%, F1 = 0.618. If a value differs, check the cell addresses in the formula.
4. In H1:K13, build a table (labels in column H, settings in I, J, K) for three clinical settings with the same sensitivity (0.85) and specificity (0.90): **screening in the general population (2%)**, **outpatients with cough (10%)**, **inpatients with fever and crackles (40%)**. For each column: the number of diseased people `=ROUND(1000*prevalence)`, TP `=ROUND(diseased*0.85)`, FN `=diseased-TP`, TN `=ROUND((1000-diseased)*0.90)`, FP `=(1000-diseased)-TN`, then the same six indicators. The decimal separator in formulas depends on the regional setting of the sheet (a point for a US locale, a comma for a Romanian one).
5. Compare. At a prevalence of 2% you should obtain PPV ≈ 14.8%, at 10% PPV ≈ 48.6%, and at 40% PPV = 85.0%, while accuracy remains between 88% and 90%.
6. Calculate the accuracy of a "test" that declares everyone healthy, at a prevalence of 2% (answer: 98%). In the sheet, in a comment cell, write two sentences: why accuracy is a misleading indicator when the disease is rare, and why the tool's PPV cannot be "carried over" from an infectious diseases hospital to a screening program.

Methodological note: the F1 score ignores true negatives and, in an analysis of 32 measures, it met neither of the two quality criteria analyzed by the authors [500]. It is calculated here only because it appears frequently in AI papers.

L13.2 Part B. The ROC Curve Built Step by Step (35 min)

The same type of tool does not directly produce "positive/negative" but a score between 0 and 1. Table L13.2 contains the scores for 16 fictitious patients, 6 with the disease present (status = 1) and 10 without the disease (status = 0).

Table L13.2. Scores of the fictitious tool for 16 patients, sorted in descending order. Fictitious data, constructed for the exercise.

Patient	Score	Status	Patient	Score	Status
P01	0.95	1	P09	0.41	0
P02	0.88	1	P10	0.34	1
P03	0.81	1	P11	0.30	0
P04	0.76	0	P12	0.26	0
P05	0.67	1	P13	0.19	0
P06	0.62	0	P14	0.12	0
P07	0.58	1	P15	0.08	0
P08	0.49	0	P16	0.05	0

Steps:

1. In a new sheet, copy the data into columns A (patient), B (score), and C (status), starting in row 3. Leave row 2 for the starting point of the curve (0, 0). Check that the table is sorted in descending order by score (menu *Data* > *Sort range*).
2. Each row represents a threshold: patients with a score greater than or equal to the score in that row are declared positive. In D3 enter `=SUM($C$3:C3)` (cumulative true positives), in E3 `=COUNTIF($C$3:C3,0)` (cumulative false positives), and drag the formulas down to row 18. In D2 and E2 enter 0. If the sheet uses a Romanian regional setting, arguments are separated by semicolons: `=COUNTIF($C$3:C3;0)`.
3. In F2, calculate the sensitivity `=D2/6`; in G2, the false positive rate (1 - specificity) `=E2/10`; then copy down to row 18. The last row must give F = 1 and G = 1.
4. Select F2:G18 and insert a scatter chart. In the chart editor, choose column G for the X axis and column F as the series (Y axis). Google Sheets does not connect the points of a scatter chart; for a continuous curve, choose the line chart type, with the same column G on the X axis (Excel and LibreOffice Calc offer a scatter-with-lines variant directly). Add the diagonal (0, 0)-(1, 1) as a second series. The resulting curve should look like Figure L13.1.
5. Calculate the area under the curve with the trapezoidal rule: `=SUMPRODUCT((G3:G18-G2:G17)*(F3:F18+F2:F17)/2)`. The correct result is 0.883.
6. Check it another way: for each of the 6 diseased patients, count how many of the 10 healthy patients have a lower score. The sum of correctly ordered pairs is 53 out of 60 possible, and 53/60 = 0.883. The AUC is therefore the probability that a randomly chosen diseased patient has a higher score than a randomly chosen healthy patient [501].
7. In H2, calculate the Youden index `=F2-G2`, copy it down, and find the maximum with `=MAX(H2:H18)`. The optimal threshold by this criterion is 0.58 (Se 0.83, Sp 0.80, J = 0.63). Discuss in your group why, for a triage tool, the threshold 0.34 (Se 1.00, Sp 0.60) might be preferred, even though its Youden index is slightly lower.

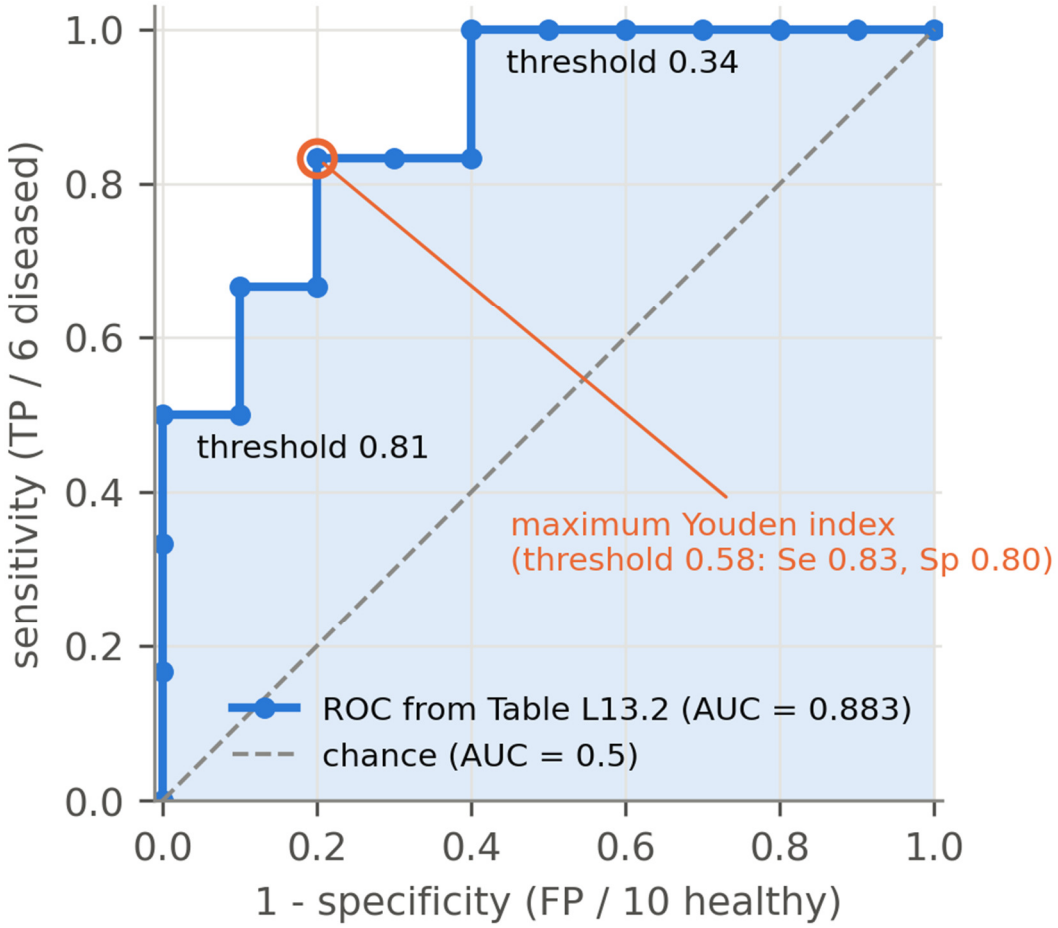


Figure L13.1. The ROC curve built from Table L13.2: each point corresponds to a threshold, each vertical step to a diseased patient, and each horizontal step to a healthy one. AUC calculated with the trapezoidal rule = 0.883, and the circled point marks the maximum Youden index. Fictitious data, calculated in the figure script.

Common mistakes in this part: sorting in ascending instead of descending order (the curve appears below the diagonal), dividing by the total number of patients instead of by 6 and 10, respectively, and omitting the starting point (0, 0), without which the curve no longer starts from the origin of the chart. If two patients had the same score, one diseased and one healthy, the threshold would declare them both positive at once, and the curve would draw a diagonal segment instead of a step. Table L13.2 contains no tied scores, precisely so that the steps are easy to follow.

L13.3 Part C. Critical Appraisal of a Published Study (40 min)

The chosen article prospectively validates, in four US health systems, the second version of a commercial sepsis prediction model. The abstract reports 227,091 hospitalizations, of which 7,401 (3.3%) met the Sepsis-3 criteria, an encounter-level AUC between 0.82 and 0.92 depending on the site, and a PPV between 0.13 and 0.26 [503]. The previous version of the same model had shown an AUC of 0.63 in the 2021 external validation [502].

Work in pairs. Read the full text and complete the checklist in Table L13.3, recording "yes," "partly," or "no" for each item, together with the page or section where you found the information. The checklist is a teaching selection of TRIPOD+AI topics; it does not replace the official 27-item checklist [505].

Table L13.3. Short checklist for a validation study of a predictive model, derived from the TRIPOD+AI topics. Source: adapted from [505].

No.	Question	Where to look
1	Do the title and abstract state that this is a validation study of a predictive model?	Title, abstract
2	Are the data source, the sites, and the period described?	Methods
3	Is the predicted outcome (sepsis) defined explicitly and in the same way at all sites?	Methods
4	Is the model described (version, time of prediction, time horizon)?	Methods
5	Are the number of patients and the number of events reported?	Results
6	Is discrimination reported, with a confidence interval?	Results
7	Is calibration reported (plot or measures)?	Results, appendices
8	Are clinical utility measures reported (PPV, alert burden)?	Results
9	Is performance analyzed in subgroups (sex, age, race, site)?	Results, appendices
10	Are limitations, funding, conflicts of interest, and code availability declared?	Discussion, declarations

Then carry out two calculations:

1. With PPV = 0.13, how many of 100 alerts are false alarms? (87.) With PPV = 0.26? (74.)
2. The number needed to evaluate, that is, the number of patients who must be evaluated to find one true case, is 1/PPV. Calculate it for the two extremes (approximately 8 and 4).

Conclude with a paragraph of 80-120 words that answers the question: "Would you recommend activating this model in a hospital in Romania and, if so, under what monitoring conditions?" Use the concepts of external validation, prevalence, and data set shift from Chapter 13. For a formal risk-of-bias assessment, there is the PROBAST+AI tool [506], optional for interested students.

L13.4 Part D. Structured Debate (40 min)

The motion: "An AI system should be allowed to sign an imaging report on its own."

Roles. The group is divided into five teams of 2-3 students, each with a role, plus a moderator (the instructor or a student) and the audience:

- **the clinician (radiologist)**: workload, human error, experience with assistance systems
- **the patient**: the right to know who read the examination, trust, the path for filing a complaint
- **the developer**: reported performance, cost, model updating
- **the regulator**: the classification of software as a medical device and the obligations for high-risk AI systems
- **the ethicist**: accountability, fairness, automation bias.

Rules. Each team has 3 minutes for its opening statement, followed by a round of 1-minute rebuttals and 5 minutes of questions from the audience. Every factual argument must be supported by a source from Chapter 13 or from the list below. Unsourced claims earn no points. For the final tally, the audience votes twice, before and after the debate, and the moderator records the change.

Minimum recommended sources: the AI Act obligations for high-risk AI systems, including human oversight [523], MDR Rule 11 for the classification of software [284], the MASAI trial on AI-supported mammography reading [271], the experiment on automation bias in radiologists [281], and the ICMJE position that an AI tool cannot be an author because it cannot be accountable for the work [202].

Team scoring rubric (10 points): clarity and coherence of the argument (0-3), correct use of evidence, with numbers and sources (0-3), response to counterarguments (0-2), adherence to the role and the time limit (0-1), a conclusion stated in one sentence (0-1).

At the end of the debate, the moderator summarizes the points of agreement and disagreement in five minutes. The summary must separate three questions that are easily conflated: whether the system reads well enough (performance evidence), who is accountable when it errs (legal and ethical responsibility), and what the law in force on September 27, 2026, allows (regulation). Each team records on its sheet one argument from the other side that it considers the hardest to refute.

What to Submit

A single file is uploaded to the discipline's platform (a Google Sheets file shared with the instructor, or an .xlsx/.ods file), with four sheets:

1. **Part A**: the matrix, the six indicators, and the table for the three prevalences, with the two interpretive sentences (3 points).
2. **Part B**: the cumulative table, the ROC chart, the trapezoidal AUC, the pairwise check, and the Youden index (3 points).
3. **Part C**: the completed checklist from Table L13.3, the two calculations, and the recommendation paragraph (2 points).
4. **Part D**: the team sheet, with the arguments and sources used, plus the converted rubric score (2 points).

Indicative grading scheme at the practical examination: the student receives a different confusion matrix and a different prevalence and must calculate Se, Sp, PPV, NPV, and accuracy within 10 minutes, then explain why the PPV has changed.

Offline Variant

If there is no internet access, Parts A and B are completed in LibreOffice Calc 26.8.0 installed locally, with the same formulas (in the Romanian-language version of the program, the argument separator may be the semicolon). If no computer is available either, the work is done on paper: the matrix and the indicators with a pocket calculator, the ROC curve on graph paper, with 10 squares horizontally (one healthy patient = one step to the right) and 6 vertically (one diseased patient = one step up), and the AUC by counting the squares under the curve (53 out of 60). For Part C, the instructor distributes a printed copy of the article. The debate does not depend on the network.

PRACTICAL SESSION 14

Practical Session 14. The Practical Examination: Sample Paper, Grading Scale, and Review

Objective

This session prepares students for the practical examination. According to the curriculum of the discipline, the practical examination accounts for 30% of the final grade. By the end of the session, the student knows the format of the examination, has solved a complete sample paper, knows how each point is awarded, and has a review plan for Practical Sessions 1-13. The paper below is a teaching model proposed by this book: the number of tasks, the time allowed, and the scoring may be set differently by the course leader, and the official version is announced in the lecture. All cases, values, and files in this session are fictitious.

Tools

The tools are those of the previous practical sessions, all free of charge, with their addresses verified on September 27, 2026:

- Google Forms and Google Sheets (or LibreOffice Calc), for tasks S1 and S3 [123,124].
- The WHO ICD-10 browser, 2019 version, the ICD-11 MMS browser, release 2026-01, and the ICD-11 Reference Guide, for task S2 [133,134,141].
- Fiji, in the version installed in the laboratory, with an anonymized DICOM file and the Blobs sample image, for task S4 [245,557].
- PubMed, for verifying the references in task S6 [41].
- A pocket calculator, paper, and a pen, for the offline variant.

Students do not use a chatbot during the examination. Task S6 works on a language model's answer printed in advance by the instructor.

Procedure

The proposed model has six tasks, 90 minutes, and 10 points (Figure L14.1). The tasks can be solved in any order. Every answer is saved in a single submission document named `L14_Name_Surname`, and the working files (the spreadsheet and the screenshots, which contain no personal data) are attached at the end.

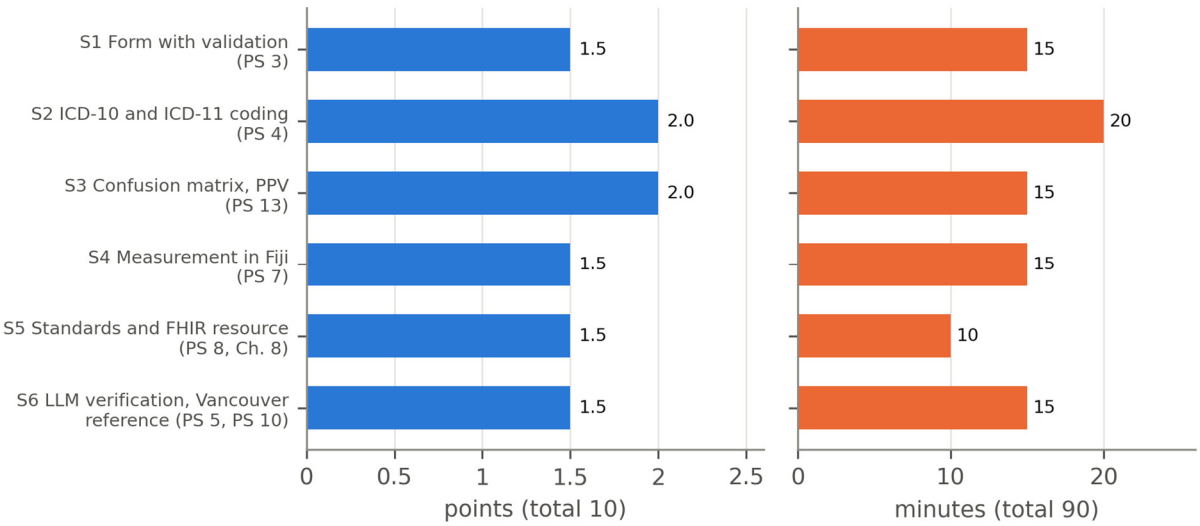


Figure L14.1. The practical examination model proposed in this book: the points (left) and the approximate time (right) for each of the six tasks, with the practical session from which each task is drawn. A teaching model, not a regulation.

L14.1 The Sample Paper

S1. A Form with Validation (1.5 points, 15 minutes, Practical Session 3). The fictitious study "Sleep and coffee consumption in first-year students" collects the pseudonymous code, age, the hours of sleep the previous night, coffee consumption (yes/no), and, only for those who drink coffee, the number of cups per day. Requirements:

1. Write the data dictionary for the five variables: short name, label, type, unit, and permitted values.
2. Build the form in Google Forms, with validation of the code by the regular expression `^S[0-9]{3}$`, of age between 18 and 40 years, and of the hours of sleep between 0 and 14 [123].
3. Set up the branching so that the section with the number of cups is reached only after the answer "yes" [124].
4. Explain in one sentence why the form does not ask for an e-mail address.

S2. Coding in ICD-10 and ICD-11 (2 points, 20 minutes, Practical Session 4). Code the main condition of each fictitious case in ICD-10 (WHO 2019 version) and in ICD-11 (MMS 2026-01). Write the version next to each code.

1. A 64-year-old woman, follow-up visit for type 2 diabetes mellitus, with no documented complications.
2. A 30-year-old woman, acute cystitis.
3. An 82-year-old woman tripped over a rug in her home and fell. Diagnosis: closed subcapital fracture of the left femoral neck.

S3. Performance of an AI Tool (2 points, 15 minutes, Practical Session 13). The fictitious tool "OtoScan" was tested on 1,000 patients, 100 of whom had the disease according to the reference standard. The results were TP = 80, FP = 45, FN = 20, and TN = 855.

1. Calculate the sensitivity, specificity, positive predictive value (PPV), negative predictive value (NPV), and accuracy [499].
2. Recalculate the PPV and NPV for a screening program with a prevalence of 2%, keeping the sensitivity and specificity unchanged.

3. Explain in two sentences why the PPV changed.

S4. Measurement in Fiji (1.5 points, 15 minutes, Practical Session 7). Open the anonymized DICOM file you have received (a fictitious CT slice for teaching purposes).

1. Read Pixel Spacing (0028,0030), Window Center (0028,1050), Rescale Slope (0028,1053), and Rescale Intercept (0028,1052) from the header [245].
2. A pixel has a stored value of 1064, with Slope = 1 and Intercept = -1024. Calculate its value in Hounsfield units.
3. Pixel Spacing is 0.78 mm. A line drawn across the diameter of a structure is 64 pixels long. Calculate its length in millimeters and check it with Analyze > Measure.
4. On the Blobs image, apply the Default threshold, then run Analyze Particles with a size of 50-Infinity and Exclude on edges. Record the number of objects and explain what the Exclude on edges option does.

S5. Interoperability Standards (1.5 points, 10 minutes, Practical Session 8, Chapter 8). The fictitious patient A.T. is admitted to the hospital. For each step, name the main standard or component involved:

1. the analyzer sends the blood glucose result to the laboratory system and to the hospital system, with a code for the test
2. the CT scanner sends the series to the image archive
3. an authorized mobile application reads the blood pressure from the electronic record through an API
4. a physician in another Member State requests the patient's summary
5. the physician signs the medical letter so that the signature has the legal effect of a handwritten one.

Then, in the JSON fragment below, identify the patient, the measured parameter, the value, the unit, and the time of the measurement [567].

```
{"resourceType": "Observation", "status": "final",
 "code": {"coding": [{"system": "http://loinc.org", "code": "8867-4"}]},
 "subject": {"reference": "Patient/demo-17"},
 "effectiveDateTime": "2026-10-12T08:30:00+03:00",
 "valueQuantity": {"value": 88, "unit": "/min"}}
```

S6. Verifying a Generated Answer and Citing (1.5 points, 15 minutes, Practical Sessions 5 and 10). You receive the printed answer of a chatbot to the fictitious case "a 45-year-old man with a headache of sudden onset during exertion, the worst of his life, with nausea, who asks which painkiller to take." The answer contains four numbered statements and two references with PMIDs.

1. Classify each statement as confirmed, refuted, or unverifiable, and indicate the error category according to the grid used in Practical Session 10 [406].
2. Look up the two PMIDs in PubMed and record whether each reference exists and whether it supports the statement.
3. Write, in Vancouver style, the reference for the article with the following metadata: authors Walters WH and Wilder EI, title "Fabrication and errors in the bibliographic citations generated by ChatGPT," journal Scientific Reports, year 2023, volume 13, issue 1, article 14045, DOI 10.1038/s41598-023-41032-5 [201].

4. Write, in two sentences, the AI use disclosure statement for a report that would cite this answer, in line with the ICMJE recommendations [202].

L14.2 Solutions and Grading Scale

The solutions below are those the instructor uses for grading. Table L14.1 summarizes the scoring.

S1. The data dictionary contains, for example: `code` (text, S followed by three digits), `age` (integer, years, 18-40), `sleep_hours` (numeric, hours, 0-14), `coffee` (nominal, yes/no), and `cups_per_day` (integer, completed only if `coffee` = yes). The e-mail address is not requested because it is a direct identifier that serves no purpose in the study, and data concerning health are collected only to the extent necessary.

S2. The solution follows Table L4.2 [133,134,141]:

- case 1: E11.9 in ICD-10 and 5A11 in ICD-11
- case 2: N30.0 in ICD-10 and the cluster GC00.1&XT5R in ICD-11
- case 3: S72.0, with W01 as the external cause, in ICD-10, and the cluster NC72.20&XK8G&XJ44E/PA60 in ICD-11. The laterality and closed fracture extensions are linked to the fracture code with "&", and the external cause code PA60 is linked with "/".

S3. The sensitivity is 80/100 = 80.0%, the specificity 855/900 = 95.0%, the PPV 80/125 = 64.0%, the NPV 855/875 = 97.7%, and the accuracy 935/1,000 = 93.5%. At a prevalence of 2%, there are 20 patients with the disease: TP = 16 and FN = 4. Of the 980 healthy individuals, TN = 931 and FP = 49. This gives PPV = 16/65 = 24.6% and NPV = 931/935 = 99.6%. The PPV falls because, at the same specificity, the number of false positives remains high, while the number of true positives decreases along with the prevalence.

S4. HU = 1 × 1064 - 1024 = 40 HU. The length is 64 × 0.78 = 49.9 mm. The Exclude on edges option removes from the count the objects that touch the edge of the image, because they are incomplete and would distort the mean area. The number of objects is compared with the value obtained by the instructor on the same file.

S5. (1) HL7 v2, with a LOINC code for the test. (2) DICOM, the C-STORE service. (3) HL7 FHIR through the REST interface, with SMART App Launch authorization. (4) The International Patient Summary (IPS), through MyHealth@EU, mandatory in the EU from March 26, 2029. (5) The qualified electronic signature, under eIDAS. In the JSON fragment, the patient is the reference `Patient/demo-17`, the parameter is the pulse rate (LOINC 8867-4), the value is 88, the unit is /min, and the time is October 12, 2026, at 08:30.

S6. Any recommendation to take a painkiller and wait, without immediate referral to the emergency department, is classified as a missed red flag of major severity. The correct reference is: Walters WH, Wilder EI. Fabrication and errors in the bibliographic citations generated by ChatGPT. Sci Rep. 2023;13(1):14045. doi:10.1038/s41598-023-41032-5. The disclosure statement names the tool, the version, the date, and the purpose, and states that the author has verified the statements.

Table L14.1. Grading scale for the sample paper (10 points). A teaching model.

Task	What is scored	Points	What cancels the score
S1	dictionary 0.5; at least three validations 0.5; branching 0.25; justification for omitting the e-mail 0.25	1.5	collecting names, personal numeric codes (CNP), or e-mail addresses
S2	for each case: ICD-10 0.3 and ICD-11 0.3; versions recorded 0.2	2.0	an extension code alone or in first position
S3	five indicators 1.0; recalculation at 2% 0.6; explanation 0.4	2.0	formulas without visible calculation cells
S4	header attributes 0.4; HU 0.3; length 0.4; counting and explanation 0.4	1.5	measurement on an image exported as JPEG
S5	five standards, 0.2 each; JSON elements 0.5	1.5	none
S6	statements classified 0.5; PMIDs verified 0.3; Vancouver reference 0.4; disclosure 0.3	1.5	"confirmed" without an identifiable source

Partial credit is awarded for correct reasoning that contains a calculation error. A correct code without a version receives half the points, as in Practical Session 4.

L14.3 Frequent Mistakes

Common Pitfalls

- In S1: numeric fields left as free text, without validation, or branching placed on a short-answer question, where Google Forms does not allow it.
- In S2: stopping at the three-character category, placing "&" between two stem codes, omitting the version.
- In S3: dividing by the total number of patients instead of the number of patients with or without the disease, confusing the PPV with sensitivity.
- In S4: changing the scale with Set Scale, pressing the Apply button in Window/Level, measuring on a screenshot.
- In S5: "HL7" without a version, confusing DICOM with PACS, describing the IPS as a complete record.
- In S6: accepting a statement as true because it sounds plausible, accepting a reference without searching PubMed, listing references alphabetically instead of in citation order.
- At any point: real data of actual persons in screenshots or files.

All these mistakes were discussed in the chapters and practical sessions from which the tasks are drawn. The competency checklist in Table 14.1 can be used for self-assessment before the examination.

L14.4 Preparation Plan

Review is based on the practical sessions, not on lecture notes. The plan below assumes three weeks.

1. **Three weeks before.** Redo Practical Sessions 3 and 4 on new cases. For Practical Session 4, choose five diagnoses from Chapter 4 and code them without looking at the solutions, then compare.
2. **Two weeks before.** Redo the calculations of Practical Session 13 with other values and other prevalences, until working through a matrix takes less than five minutes. Redo Practical Session 7 on an anonymized DICOM image and save the macro.
3. **One week before.** Go through Practical Session 8 and Chapter 8 for the standards. Draw up on paper a "situation, standard" table with at least ten rows. Repeat the verification of a generated answer, following Practical Session 10, and write three Vancouver references without a reference manager.
4. **In the last few days.** Solve the sample paper under timed conditions, in 90 minutes. Grade your own work with Table L14.1 and note the task on which you lost the most points.
5. **On the day of the examination.** Check your access to the institutional account and your passwords before the examination, without writing them down on the answer sheet.

The practical sessions that use spreadsheets (Practical Sessions 3, 12, and 13) continue the protocols of [6], revised and updated for Google Sheets/Excel 2026. Practicing them on new data is the most effective form of review, because the examination changes the values, not the method.

What to Submit

At the examination, students submit a single document with the answers to tasks S1-S6, accompanied by the spreadsheet (S1, S3), the link to the form shared with the instructor with the Viewer role (S1), and the Fiji screenshots without personal data (S4). At the preparation session, students submit the solved sample paper, self-graded with Table L14.1, with the score obtained for each task. Indicative grading scheme: Table L14.1, 10 points in total.

Offline Variant

Without a network connection, the tasks are solved as follows. S1 is completed in a local spreadsheet, with data validation rules and a column that flags incompatible answers, as in the offline variant of Practical Session 3. For S2, the instructor hands out printed excerpts from the ICD-10 Tabular List and the ICD-11 pages of the entities used, including the postcoordination section. S3 is solved with a pocket calculator. For S4, a locally installed copy of Fiji opens the DICOM file copied to the computer in advance, and the Blobs image is replaced with the synthetic image from Practical Session 7. S5 requires no network. For S6, the PubMed records of the references are printed and mixed with nonexistent references, and the student compares them with the list received.

Appendix A. Answers to the Self-Assessment Questions

SA denotes single-answer questions (one correct option) and MA multiple-answer questions (two to four correct options).

Table A.1. Answer key by chapter.

Chapter	Correct answers
1	1-c; 2-b; 3-c; 4-a,c,d; 5-a,b,c,e
2	1-d; 2-d; 3-b; 4-a,b,d; 5-a,b,c
3	1-b; 2-a; 3-b; 4-a,b,d; 5-a,b,c,d
4	1-d; 2-c; 3-b; 4-a,b,d; 5-a,b,d
5	1-b; 2-b; 3-c; 4-a,b,d; 5-a,d,e
6	1-d; 2-c; 3-b; 4-a,b,d; 5-a,b,d
7	1-d; 2-e; 3-b; 4-a,b,d,e; 5-a,b,c
8	1-c; 2-c; 3-d; 4-a,b,c,e; 5-a,b,d
9	1-b; 2-c; 3-b; 4-a,b,c,e; 5-a,b,d
10	1-c; 2-c; 3-c; 4-a,b,d; 5-a,c,d
11	1-e; 2-b; 3-b; 4-a,b,c; 5-a,c,d
12	1-c; 2-c; 3-d; 4-a,c,e; 5-a,b,c,e
13	1-c; 2-b; 3-d; 4-a,b,d; 5-b,c,d
14	1-d; 2-b; 3-b; 4-a,b,d,e; 5-a,b,d; 6-c; 7-b; 8-a,c,e; 9-b; 10-a,b,c; 11-a,b,c; 12-c; 13-a; 14-a,b,d; 15-c; 16-a,b,c,e; 17-b; 18-a,b,d; 19-b; 20-a,b,c; 21-c; 22-a,b,c; 23-a; 24-a; 25-a,b,d

Appendix B. Glossary

- **Accuracy**: the proportion of correct decisions made by a test, (TP + TN) / total. It can be misleading when the disease is rare. (Ch. 13)
- **AI literacy**: the knowledge and skills of staff who use AI systems, which providers and deployers must support under Article 4 of the AI Act. (Ch. 13, 14)
- **Alert fatigue**: the tendency to ignore alerts because of the large number of irrelevant warnings. (Ch. 8, 12)
- **Algorithmic bias**: a systematic error of an algorithm whose consequences fall disproportionately on a particular group of patients. (Ch. 5, 13)
- **Aliasing**: the appearance of a false, lower frequency when a signal is sampled at less than twice its maximum frequency. It can no longer be corrected after digitization. (Ch. 6)
- **Anonymization**: the irreversible removal of the link between data and a person. Anonymous data fall outside the GDPR, but true anonymization is difficult to achieve. (Ch. 5)
- **Archetype**: a reusable model of a clinical concept, for example blood pressure, defined by clinicians in openEHR. (Ch. 8)
- **Area under the ROC curve** (AUC): the probability that a randomly chosen patient with the disease has a higher score than a randomly chosen person without it. It measures discrimination, not calibration. (Ch. 13)
- **Asymmetric encryption** (*public-key cryptography*): encryption with a pair of mathematically linked keys, one public and one private, the latter kept only by its holder. (Ch. 11)
- **Audit trail** (*audit log*): the record of every access to and modification of the data, with the author, the time, and the action. (Ch. 3, 5, 11)
- **Automation bias**: the tendency to accept the suggestion of an automated system even when it contradicts other evidence. (Ch. 7, 13)
- **Bed occupancy rate**: the inpatient days delivered divided by the available bed-days, multiplied by 100. (Ch. 12)
- **Biomedical informatics**: the interdisciplinary field that studies and pursues the effective use of biomedical data, information, and knowledge for research, problem solving, and decision making, with the aim of improving health. (Ch. 1)
- **Bit** (*binary digit*): the amount of information obtained from a choice between two equally probable alternatives. Eight bits make up a byte. (Ch. 2)
- **Broad consent**: consent that covers a category of future research rather than a single study. (Ch. 5)
- **Calibration**: the agreement between the probabilities displayed by a model and the actual frequency of the event. (Ch. 13)
- **Case-mix index**: the mean of the relative weights of the cases discharged by a hospital over a given period. (Ch. 4, 12)
- **CIA triad**: confidentiality, integrity, and availability, the three properties by which information security is described. (Ch. 5, 11)
- **Classification**: a set of exhaustive, mutually exclusive classes that group concepts which no longer need to be distinguished in the analysis, such as the ICD. (Ch. 4)
- **Clinical datum**: a single observation about a patient, described by the patient, the parameter, the value, the time, and the method. (Ch. 1, 2)

- **Clinical terminology**: a controlled nomenclature with stable identifiers, synonyms, and relations between concepts. (Ch. 4)
- **Cluster**: in ICD-11, the group of codes linked by postcoordination that describes a single clinical concept. (Ch. 4)
- **Confidentiality**: the property of information of being accessible only to authorized persons. It refers to information, unlike privacy, which refers to the person. (Ch. 5)
- **Confusion matrix**: the four-cell table (true positive, false positive, false negative, true negative) from which the performance measures of a test are calculated. (Ch. 13)
- **Data concerning health**: personal data related to the physical or mental health of a person, including the provision of healthcare services, which reveal information about his or her health status. They are a special category under the GDPR. (Ch. 5)
- **Data dictionary** (*codebook*): the description of each variable in a data set: name, label, type, unit, permitted values, and the code for missing values. (Ch. 3)
- **Data lake**: a repository that keeps data in their raw form and structures them only at the time of analysis. (Ch. 3)
- **Data leakage**: access by a model, during development, to information it would not have in real-world use, which artificially inflates performance. (Ch. 13)
- **Data warehouse**: a database into which data from routine systems are periodically copied, cleaned, integrated, and organized for analysis. (Ch. 3)
- **Data set shift**: the situation in which the data encountered in operation no longer resemble those used to develop a model, because of changes in technology, population, or practice. (Ch. 13)
- **Decision curve analysis**: a method that calculates the net benefit of a model for each threshold probability at which the clinician would intervene. (Ch. 13)
- **Downtime procedures**: the procedures by which a hospital operates when the information system is unavailable. (Ch. 12)
- **Dynamic consent**: consent managed on a digital platform, through which participants update their preferences over time and can withdraw their agreement. (Ch. 5)
- **Electronic health record** (EHR): an electronic repository of information about a person's health and care, with functions that support clinical practice and the other legitimate uses of the data. (Ch. 8)
- **Electronic medical record** (EMR): the electronic documentation of a single institution, practice, or hospital. (Ch. 8)
- **End-to-end encryption**: encryption in which the message is decrypted only on the recipient's device, and the intermediate server cannot read it. (Ch. 9, 11)
- **Entropy**: the average information expected from a variable, that is, its uncertainty before observation, $H = -\Sigma\, p \cdot \log_2 p$. (Ch. 2)
- **Explainable AI** (XAI): methods that attempt to show what determined a model's prediction, for example SHAP values or saliency maps. (Ch. 13)
- **Extension code**: an ICD-11 code from Chapter X that adds detail to a stem code and cannot stand alone or in the first position of a cluster. (Ch. 4)
- **External validation**: the evaluation of a model on data collected at a different place or during a different period from those used for development. (Ch. 7, 13)
- **FAIR principles**: the requirement that data be findable, accessible, interoperable, and reusable. It does not mean public data. (Ch. 3)

- **FHIR resource**: the basic unit of exchange in FHIR, with elements, data types, and a human-readable section, represented in JSON or XML. (Ch. 8)
- **Foreign key**: a column that must point to an existing row in another table. (Ch. 3)
- **Foundation model**: a model trained by self-supervision on large and diverse data sets, which can perform many tasks with few or no task-specific labels. (Ch. 13)
- **Hallucination**: a statement presented as fact by a generative model but not supported by the source or by reality. (Ch. 10, 13)
- **Hash function**: a function that computes a fixed-length fingerprint from data of any length. It has no key and serves to verify integrity. (Ch. 11)
- **Heart rate variability** (HRV): the fluctuation of the intervals between normal sinus beats, described by measures such as SDNN and RMSSD. (Ch. 6)
- **Hospital information system** (HIS): the computerized part of a hospital's information system, made up of modules such as ADT, FOCG, CPOE, LIS, and RIS/PACS. (Ch. 12)
- **Hounsfield unit**: the unit of CT values, with water at 0 HU and air at approximately -1,000 HU. (Ch. 7)
- **ICD-11 Foundation** (*Foundation Component*): the semantic knowledge base from which the ICD-11 linearizations are derived. (Ch. 4)
- **IHE integration profile**: a specification that sets out which standards and which options are used in a given clinical scenario, for example Scheduled Workflow. (Ch. 7, 11)
- **Information**: data that have been analyzed and organized so that they acquire meaning. (Ch. 1, 2)
- **Information system**: all the processing of data, information, and knowledge in an organization, together with the human or technical actors that carry it out. The computer-supported part is the computer-based information system. (Ch. 11, 12)
- **Integration engine**: a communication server at the center of a star architecture that receives, queues, transforms, and distributes messages between hospital applications. (Ch. 12)
- **International Patient Summary** (IPS): a minimal set of clinical data for care across organizational boundaries, with three mandatory sections: problems, allergies, and medication. (Ch. 8)
- **Interoperability**: the ability of two systems to exchange data and to use them without manual intervention, at the technical, syntactic, semantic, and organizational levels. (Ch. 8)
- **Knowledge**: the result of the formal or informal analysis of information, which includes study results, heuristics, and models. (Ch. 1, 2)
- **Large language model** (LLM): a neural network trained on very large volumes of text to predict the next word fragment. It does not consult a base of verified facts. (Ch. 10, 13)
- **Least privilege**: granting access by role, limited to the needs of that role. (Ch. 5, 11)
- **Linearization**: the extraction of entities from the ICD-11 Foundation into a classification for a specific purpose, without changing their meaning. The MMS is the statistical linearization. (Ch. 4)
- **Machine learning**: a field of AI in which systems learn to solve problems by analyzing data related to those problems. (Ch. 13)
- **Main condition**: the condition diagnosed at the end of the episode of care that is primarily responsible for the need for treatment or investigation. (Ch. 4)
- **Medical device software**: software with a medical purpose declared by the manufacturer for an individual patient, regulated by the MDR. (Ch. 10)

- **Medical informatics**: the application of biomedical informatics focused on disease and on the physician's work. In this book, the discipline as the medical student encounters it. (Ch. 1)
- **Medical teleassistance**: in Romanian regulations, the real-time integration of data or opinions from remote professionals for the performance of medical procedures. (Ch. 9)
- **Multi-factor authentication**: authentication that combines at least two factors, for example a password and a code generated on a phone. (Ch. 5, 11)
- **Nomenclature**: a set of coded terms that name medical concepts as precisely as possible, without grouping them, such as SNOMED CT. (Ch. 4)
- **Normalization**: the decomposition of a redundant table into smaller tables linked by keys, so that each fact is recorded only once. (Ch. 3)
- **Ontology**: a representation of the concepts of a domain through formal relations ("is a", "has finding site"), from which a program can infer relations that are not explicitly stated. (Ch. 2, 4)
- **Patient-generated health data**: clinically relevant data captured by patients or their devices outside the traditional care setting. (Ch. 2, 10)
- **Personal health record** (PHR): a record controlled by the patient, in which patients view their results, communicate with the care team, and can add data measured at home. (Ch. 8)
- **Phishing**: a message that imitates a trusted source in order to obtain a password, a click, or the opening of an attachment. (Ch. 5, 11)
- **Pixel** (*picture element*): the element of a two-dimensional digital image that represents the intensity of a small area. The corresponding volume element is called a voxel. (Ch. 7)
- **Plagiarism**: presenting as one's own work texts, ideas, data, or methods taken from the works of other authors, without reference to the source. (Ch. 5)
- **Positive predictive value**: the probability that a patient with a positive test has the disease, TP / (TP + FP). It depends on prevalence. (Ch. 13)
- **Postcoordination**: linking two stem codes, or a stem code with extension codes, to describe a clinical concept, in ICD-11 with "&" and "/". (Ch. 4)
- **Primary key**: the attribute or group of attributes that uniquely identifies each row of a table. (Ch. 3)
- **Primary use**: the use of health data for the care of the person concerned, including in another Member State. (Ch. 5, 8)
- **Pseudonymization**: the processing of data in such a way that they can no longer be attributed to a person without additional information, which is kept separately. Pseudonymized data remain personal data. (Ch. 5)
- **Qualified electronic signature**: an advanced signature created with a qualified device and based on a qualified certificate. It has the legal effect of a handwritten signature. (Ch. 11)
- **Redundancy**: the partial repetition of information, which allows errors caused by noise to be detected and sometimes corrected. (Ch. 2)
- **Rule 11**: the rule in Annex VIII of the MDR that determines the risk class of medical device software. (Ch. 10, 13)
- **Sampling**: measuring the value of an analog signal at equal time intervals. (Ch. 6)
- **Secondary use**: the reuse of health data for research, public health, statistics, policy, or education, through health data access bodies. (Ch. 5, 8)
- **Semantic interoperability**: the ability of the receiving system to understand the data in the exact sense in which they were sent, through shared terminologies. (Ch. 4, 8)

- **Sensitivity**: the proportion of patients with the disease detected by the test, TP / (TP + FN). (Ch. 13)
- **Signal-to-noise ratio** (SNR): the ratio between the power of the useful signal and the power of the noise, expressed in decibels. (Ch. 6)
- **Specificity**: the proportion of people without the disease correctly identified as such by the test, TN / (TN + FP). (Ch. 13)
- **Stem code**: an ICD-11 code that can be used on its own. (Ch. 4)
- **Sycophancy**: the tendency of a language model to accept the user's false premise. (Ch. 10)
- **Symmetric encryption**: encryption in which the sender and the recipient use the same secret key, for example AES. (Ch. 11)
- **Teleconsultation**: a remote consultation between patient and physician, usually synchronous, by video or telephone. (Ch. 9)
- **Tele-expertise**: remote consultation between professionals, synchronous or asynchronous, for a second opinion. (Ch. 9)
- **Telemedicine**: the remote delivery of health services through information and communication technology, for diagnosis, treatment, monitoring, or prevention. (Ch. 9)
- **Telemonitoring** (*remote monitoring*): the transmission of parameters measured by medical devices to a team that receives alerts when values fall outside the set ranges. (Ch. 9)
- **Time stamp**: electronic data that bind a set of data to a particular moment and prove that the data existed at that moment. (Ch. 11)
- **Underlying cause of death**: the disease or injury that initiated the train of morbid events leading directly to death, or the circumstances of the accident or violence that produced the fatal injury. (Ch. 4)
- **Wearable**: a sensor worn on the body, usually on the wrist, finger, chest, or arm, that transmits its data to a phone. (Ch. 10)
- **Wellness application**: software intended for use by a natural person to process health data for purposes other than medical care, as defined in the EHDS. (Ch. 10)
- **Windowing**: the choice of the level and width of the range of Hounsfield units displayed across the entire gray scale. It changes only the display, not the data. (Ch. 7)

References

1. Mihalaș GI, Lungeanu D. Informatică medicală și biostatistică. Timișoara: Editura Victor Babeș; 2009.
2. Focșa MA. Informatica medicală în era digitală: de la date la inteligența artificială. Timișoara: Editura Victor Babeș; 2025.
3. Mihalaș GI. Inteligența artificială în medicină: ghid de conversație pentru medici, studenți și pacienți. Timișoara: Editura Victor Babeș; 2025.
4. Focșa M, Tudor A, Câtu C, Muntean C, Mada L. Informatică medicală: îndrumător de lucrări practice. Timișoara: Editura Victor Babeș; 2010.
5. Vernic C, Mada L, Lungeanu D, Muntean C, Apostol AS, Câtu CO, et al. Aplicații practice de informatică medicală și biostatistică. Timișoara: Editura Victor Babeș; 2010.
6. Vernic CV, Frandeș M, Tudor A, Apostol SA, Voicu A, Șerban C, et al. Prelucrarea datelor biomedicale: aplicații practice. Timișoara: Eurostampa; 2018.
7. Shortliffe EH, Chiang MF. Biomedical Informatics: The Science and the Pragmatics. In: Shortliffe EH, Cimino JJ, editors. Biomedical Informatics: Computer Applications in Health Care and Biomedicine. 5th ed. Cham: Springer; 2021. p. 3-44. doi:10.1007/978-3-030-58721-5_1
8. Mooney SD, Tenenbaum JD, Altman RB. Bioinformatics. In: Shortliffe EH, Cimino JJ, editors. Biomedical Informatics: Computer Applications in Health Care and Biomedicine. 5th ed. Cham: Springer; 2021. p. 273-98. doi:10.1007/978-3-030-58721-5_9
9. LaVenture M, Ross DA, Staes C, Yasnoff WA. Population and Public Health Informatics. In: Shortliffe EH, Cimino JJ, editors. Biomedical Informatics: Computer Applications in Health Care and Biomedicine. 5th ed. Cham: Springer; 2021. p. 613-36. doi:10.1007/978-3-030-58721-5_18
10. European Federation for Medical Informatics. The Romanian Society of Medical Informatics [Internet]. EFMI; 2021 Sep 20 [cited 2026 Sep 27]. Available from: https://efmi.org/2021/09/20/the-romanian-society-of-medical-informatics/
11. Shortliffe EH, Chiang MF. Biomedical Data: Their Acquisition, Storage, and Use. In: Shortliffe EH, Cimino JJ, editors. Biomedical Informatics: Computer Applications in Health Care and Biomedicine. 5th ed. Cham: Springer; 2021. p. 45-75. doi:10.1007/978-3-030-58721-5_2
12. Grossi E. The long journey of artificial intelligence in medicine: an overview. Clin Exp Rheumatol. 2025;43(5):815-21. doi:10.55563/clinexprheumatol/oamfed
13. Hersh W. Information Retrieval. In: Shortliffe EH, Cimino JJ, editors. Biomedical Informatics: Computer Applications in Health Care and Biomedicine. 5th ed. Cham: Springer; 2021. p. 755-94. doi:10.1007/978-3-030-58721-5_23
14. Musen MA, Middleton B, Greenes RA. Clinical Decision-Support Systems. In: Shortliffe EH, Cimino JJ, editors. Biomedical Informatics: Computer Applications in Health Care and Biomedicine. 5th ed. Cham: Springer; 2021. p. 795-840. doi:10.1007/978-3-030-58721-5_24
15. Jaffe C, Nguyen V, Kubick WR, Cooper T, Leftwich RB, Hammond WE. Standards in Biomedical Informatics. In: Shortliffe EH, Cimino JJ, editors. Biomedical Informatics: Computer Applications in Health Care and Biomedicine. 5th ed. Cham: Springer; 2021. p. 205-40. doi:10.1007/978-3-030-58721-5_7
16. Kaul V, Enslin S, Gross SA. History of artificial intelligence in medicine. Gastrointest Endosc. 2020;92(4):807-12. doi:10.1016/j.gie.2020.06.040
17. Jumper J, Evans R, Pritzel A, Green T, Figurnov M, Ronneberger O, et al. Highly accurate protein structure prediction with AlphaFold. Nature. 2021;596(7873):583-9. doi:10.1038/s41586-021-03819-2
18. Mugaanyi J, Cai L, Cheng S, Lu C, Huang J. Evaluation of Large Language Model Performance and Reliability for Citations and References in Scholarly Writing: Cross-Disciplinary Study. J Med Internet Res. 2024;26:e52935. doi:10.2196/52935
19. Kung TH, Cheatham M, Medenilla A, Sillos C, De Leon L, Elepaño C, et al. Performance of ChatGPT on USMLE: Potential for AI-assisted medical education using large language models. PLOS Digit Health. 2023;2(2):e0000198. doi:10.1371/journal.pdig.0000198
20. Singhal K, Azizi S, Tu T, Mahdavi SS, Wei J, Chung HW, et al. Large language models encode clinical knowledge. Nature. 2023;620(7972):172-80. doi:10.1038/s41586-023-06291-2
21. Peek N, Sujan M, Scott P. Digital health and care: emerging from pandemic times. BMJ Health Care Inform. 2023;30(1):e100861. doi:10.1136/bmjhci-2023-100861
22. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No

167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act). Official Journal of the European Union. 2024 Jul 12;L 2024/1689.

23. Regulation (EU) 2025/327 of the European Parliament and of the Council of 11 February 2025 on the European Health Data Space and amending Directive 2011/24/EU and Regulation (EU) 2024/2847. Official Journal of the European Union. 2025 Mar 5;L 2025/327.
24. Rubin DL, Greenspan H, Hoogi A. Biomedical Imaging Informatics. In: Shortliffe EH, Cimino JJ, editors. Biomedical Informatics: Computer Applications in Health Care and Biomedicine. 5th ed. Cham: Springer; 2021. p. 299-362. doi:10.1007/978-3-030-58721-5_10
25. Sheikh A, Anderson M, Albala S, Casadei B, Franklin BD, Richards M, et al. Health information technology and digital innovation for national learning health and care systems. Lancet Digit Health. 2021;3(6):e383-96. doi:10.1016/S2589-7500(21)00005-4
26. Cronin RM, Jimison H, Johnson KB. Personal Health Informatics. In: Shortliffe EH, Cimino JJ, editors. Biomedical Informatics: Computer Applications in Health Care and Biomedicine. 5th ed. Cham: Springer; 2021. p. 363-89. doi:10.1007/978-3-030-58721-5_11
27. Tenenbaum JD, Shah NH, Altman RB. Translational Bioinformatics. In: Shortliffe EH, Cimino JJ, editors. Biomedical Informatics: Computer Applications in Health Care and Biomedicine. 5th ed. Cham: Springer; 2021. p. 867-911. doi:10.1007/978-3-030-58721-5_26
28. Gulkesen KH, Sonuvar ET. Data Privacy in Medical Informatics and Electronic Health Records: A Bibliometric Analysis. Health Care Anal. 2025. doi:10.1007/s10728-025-00519-0
29. Bichel-Findlay J, Koch S, Mantas J, Abdul SS, Al-Shorbaji N, Ammenwerth E, et al. Recommendations of the International Medical Informatics Association (IMIA) on Education in Biomedical and Health Informatics: Second Revision. Int J Med Inform. 2023;170:104908. doi:10.1016/j.ijmedinf.2022.104908
30. Car J, Ong QC, Erlikh Fox T, Leightley D, Kemp SJ, Švab I, et al. The Digital Health Competencies in Medical Education Framework: An International Consensus Statement Based on a Delphi Study. JAMA Netw Open. 2025;8(1):e2453131. doi:10.1001/jamanetworkopen.2024.53131
31. Khurana MP, Raaschou-Pedersen DE, Kurtzhals J, Bardram JE, Ostrowski SR, Bundgaard JS. Digital health competencies in medical school education: a scoping review and Delphi method study. BMC Med Educ. 2022;22(1):129. doi:10.1186/s12909-022-03163-7
32. Zhang W, Han J, Han X, Xu H, Wang L, Lin T, et al. Medical Students' Attitudes, Perceptions, and Self-Reported Familiarity With AI in Health Care: Systematic Review and Meta-Analysis. JMIR Med Educ. 2026;12:e89411. doi:10.2196/89411
33. Vuorikari R, Kluzer S, Punie Y. DigComp 2.2: The Digital Competence Framework for Citizens - With new examples of knowledge, skills and attitudes. EUR 31006 EN. Luxembourg: Publications Office of the European Union; 2022. doi:10.2760/115376
34. European Commission, Joint Research Centre. DigComp 3.0 [Internet]. Seville: JRC; 2025 [cited 2026 Sep 27]. Available from: https://joint-research-centre.ec.europa.eu/projects-and-activities/education-and-training/digital-transformation-education/digital-competence-framework-digcomp/digcomp-30_en
35. Williams GA, Falkenbach M, van Schaaijk A, Batenburg R, Wismar M; BeWell Consortium. Closing the digital skills gap in health care: identifying core digital skills and competencies and education and training opportunities for health professionals in the European Union. Policy Brief 71. Copenhagen: European Observatory on Health Systems and Policies, WHO Regional Office for Europe; 2025. Available from: https://www.ncbi.nlm.nih.gov/books/NBK618031/
36. World Health Organization. Global strategy on digital health 2020-2025. Geneva: WHO; 2021. ISBN 978-92-4-002092-4. Available from: https://www.who.int/publications/i/item/9789240020924
37. World Health Organization. World Health Assembly endorses extension of the global digital health strategy to 2027 [Internet]. Geneva: WHO; 2025 May 23 [cited 2026 Sep 27]. Available from: https://www.who.int/news/item/23-05-2025-world-health-assembly-endorses-extension-of-the-global-digital-health-strategy-to-2027
38. World Health Assembly. Decision WHA78(22). Global strategy on digital health 2020-2025: extension [Internet]. Geneva: World Health Organization; 2025 May 27 [cited 2026 Sep 27]. Available from: https://apps.who.int/gb/ebwha/pdf_files/WHA78/A78_(22)-en.pdf
39. National Library of Medicine (US). PubMed [Internet]. Bethesda (MD): National Center for Biotechnology Information; 2026 [cited 2026 Sep 27]. Available from: https://pubmed.ncbi.nlm.nih.gov/
40. Giustini D, Amar-Zifkin A, Chen E, Kung JY. Automated indexing in MEDLINE and the Medical Text Indexer (MTI), 2000-2025: a scoping review. J Med Libr Assoc. 2026;114(3):191-207. doi:10.5195/jmla.2026.2406

41. National Library of Medicine (US). PubMed User Guide [Internet]. Bethesda (MD): NLM; 2026 [cited 2026 Sep 27]. Available from: https://pubmed.ncbi.nlm.nih.gov/help/
42. Wallace SS, Barak G, Truong G, Parker MW. Hierarchy of Evidence Within the Medical Literature. Hosp Pediatr. 2022;12(8):745-50. doi:10.1542/hpeds.2022-006690
43. Bhattacharyya M, Miller VM, Bhattacharyya D, Miller LE. High Rates of Fabricated and Inaccurate References in ChatGPT-Generated Medical Content. Cureus. 2023;15(5):e39238. doi:10.7759/cureus.39238
44. Chelli M, Descamps J, Lavoué V, Trojani C, Azar M, Deckert M, et al. Hallucination Rates and Reference Accuracy of ChatGPT and Bard for Systematic Reviews: Comparative Analysis. J Med Internet Res. 2024;26:e53164. doi:10.2196/53164
45. Asai A, He J, Shao R, Shi W, Singh A, Chang JC, et al. Synthesizing scientific literature with retrieval-augmented language models. Nature. 2026;650(8103):857-63. doi:10.1038/s41586-025-10072-4
46. Carl N, Hetz MJ, Wies C, Haggenmüller S, Winterstein JT, Mangold MH, et al. Enhancing clinicians' trust in large language models via transparent source attribution: A randomized controlled evaluation in uro-oncology. Eur J Cancer. 2026;233:116168. doi:10.1016/j.ejca.2025.116168
47. Bedi S, Liu Y, Orr-Ewing L, Dash D, Koyejo S, Callahan A, et al. Testing and Evaluation of Health Care Applications of Large Language Models: A Systematic Review. JAMA. 2025;333(4):319-28. doi:10.1001/jama.2024.21700
48. International Medical Informatics Association. Romanian Society of Medical Informatics [Internet]. IMIA; [cited 2026 Sep 27]. Available from: https://imia-medinfo.org/wp/romanian-society-of-medical-informatics/
49. Societatea Română de Informatică Medicală. Romanian Society of Medical Informatics (SRIM) [Internet]. România: SRIM; 2026 [cited 2026 Sep 27]. Available from: https://www.srimed.ro/
50. Stoicu-Tivadar L, Gallos P, Bolboacă SD. Editorial. Appl Med Inform. 2024;46(Suppl 2):S1. Available from: https://ami.info.umfcluj.ro/index.php/AMI/article/view/1071
51. Applied Medical Informatics [Internet]. Cluj-Napoca: Universitatea de Medicină şi Farmacie "Iuliu Haţieganu", Department of Medical Informatics and Biostatistics; 2026 [cited 2026 Sep 27]. Available from: https://ami.info.umfcluj.ro/index.php/AMI
52. Asociația de Standardizare din România. ASRO/CT 319 - Informatică medicală [Internet]. Bucureşti: ASRO; 2026 [cited 2026 Sep 27]. Available from: https://www.asro.ro/asro-ct-319-informatica-medicala/
53. Hayes CN, Nakahara H, Ono A, Tsuge M, Oka S. From Omics to Multi-Omics: A Review of Advantages and Tradeoffs. Genes (Basel). 2024;15(12):1551. doi:10.3390/genes15121551
54. Tiase VL, Hull W, McFarland MM, Sward KA, Del Fiol G, Staes C, et al. Patient-generated health data and electronic health record integration: a scoping review. JAMIA Open. 2020;3(4):619-27. doi:10.1093/jamiaopen/ooaa052
55. Lewis AE, Weiskopf N, Abrams ZB, Foraker R, Lai AM, Payne PRO, et al. Electronic health record data quality assessment and tools: a systematic review. J Am Med Inform Assoc. 2023;30(10):1730-40. doi:10.1093/jamia/ocad120
56. Syed R, Eden R, Makasi T, Chukwudi I, Mamudu A, Kamalpour M, et al. Digital Health Data Quality Issues: Systematic Review. J Med Internet Res. 2023;25:e42615. doi:10.2196/42615
57. Benish WA. A Review of the Application of Information Theory to Clinical Diagnostic Testing. Entropy (Basel). 2020;22(1):97. doi:10.3390/e22010097
58. Owens DK, Goldhaber-Fiebert JD, Sox HC. Biomedical Decision Making: Probabilistic Clinical Reasoning. In: Shortliffe EH, Cimino JJ, editors. Biomedical Informatics: Computer Applications in Health Care and Biomedicine. 5th ed. Cham: Springer; 2021. p. 77-120. doi:10.1007/978-3-030-58721-5_3
59. Deane KD, Van Hoovels L, Joy VE, Olschowka N, Bossuyt X. From autoantibody test results to decision making: Incorporating likelihood ratios within medical practice. Autoimmun Rev. 2024;23(5):103537. doi:10.1016/j.autrev.2024.103537
60. Tsalatsanis A, Hozo I, Djulbegovic B. Research synthesis of information theory measures of uncertainty: Meta-analysis of entropy and mutual information of diagnostic tests. J Eval Clin Pract. 2021;27(2):246-55. doi:10.1111/jep.13475
61. Rohlfsen C, Shannon K, Parsons AS. Entropy in Clinical Decision-Making: A Narrative Review Through the Lens of Decision Theory. J Gen Intern Med. 2025;40(16):4033-9. doi:10.1007/s11606-025-09868-x
62. Shortliffe EH, Cimino JJ, editors. Biomedical Informatics: Computer Applications in Health Care and Biomedicine. 5th ed. Cham: Springer; 2021. doi:10.1007/978-3-030-58721-5

63. The Unicode Consortium. The Unicode Standard, Version 18.0.0 [Internet]. South San Francisco (CA): The Unicode Consortium; 2026 [cited 2026 Sep 27]. Available from: https://www.unicode.org/versions/Unicode18.0.0/
64. The Unicode Consortium. FAQ: UTF-8, UTF-16, UTF-32 and BOM [Internet]. South San Francisco (CA): The Unicode Consortium; c2026 [cited 2026 Sep 27]. Available from: https://www.unicode.org/faq/utf_bom.html
65. The Unicode Consortium. Latin Extended-B, Range: 0180-024F. The Unicode Standard, Version 18.0 [Internet]. South San Francisco (CA): The Unicode Consortium; 2026 [cited 2026 Sep 27]. Available from: https://www.unicode.org/charts/PDF/U0180.pdf
66. Guardiolle V, Bazoge A, Morin E, Daille B, Toublant D, Bouzillé G, et al. Linking Biomedical Data Warehouse Records With the National Mortality Database in France: Large-scale Matching Algorithm. JMIR Med Inform. 2022;10(11):e36711. doi:10.2196/36711
67. Wagner P, Strodthoff N, Bousseljot RD, Kreiseler D, Lunze FI, Samek W, et al. PTB-XL, a large publicly available electrocardiography dataset. Sci Data. 2020;7(1):154. doi:10.1038/s41597-020-0495-6
68. Erickson BJ. Imaging Systems in Radiology. In: Shortliffe EH, Cimino JJ, editors. Biomedical Informatics: Computer Applications in Health Care and Biomedicine. 5th ed. Cham: Springer; 2021. p. 733-53. doi:10.1007/978-3-030-58721-5_22
69. Nurk S, Koren S, Rhie A, Rautiainen M, Bzikadze AV, Mikheenko A, et al. The complete sequence of a human genome. Science. 2022;376(6588):44-53. doi:10.1126/science.abj6987
70. Betschart RO, Sandberg F, Blankenberg S, Zoche M, Zeller T, Ziegler A. A benchmark study of compression software for human short-read sequence data. Sci Rep. 2025;15(1):15358. doi:10.1038/s41598-025-00491-8
71. Choi BCK, Barengo NC, Diaz PA. Public health surveillance and the data, information, knowledge, intelligence and wisdom paradigm. Rev Panam Salud Publica. 2024;48:e9. doi:10.26633/RPSP.2024.9
72. Hänsel K, Dudgeon SN, Cheung KH, Durant TJS, Schulz WL. From Data to Wisdom: Biomedical Knowledge Graphs for Real-World Data Insights. J Med Syst. 2023;47(1):65. doi:10.1007/s10916-023-01951-2
73. Serrano-Zamago AB, Altamirano-Bustamante MM. Appealing to Tacit Knowledge and Axiology to Enhance Medical Practice in the COVID-19 Pandemic: A Systematic Review and Hermeneutic Bioethical Analysis. Front Public Health. 2021;9:686773. doi:10.3389/fpubh.2021.686773
74. Charlet J, Cui L, Section Editors for the IMIA Yearbook Section on Knowledge Representation and Management. Knowledge Representation and Management 2022: Findings in Ontology Development and Applications. Yearb Med Inform. 2023;32(1):225-9. doi:10.1055/s-0043-1768747
75. Iglesias N, Juarez JM, Campos M. Comprehensive analysis of rule formalisms to represent clinical guidelines: Selection criteria and case study on antibiotic clinical guidelines. Artif Intell Med. 2020;103:101741. doi:10.1016/j.artmed.2019.101741
76. Ahn JW, Jang HY, Son YJ. Critical care nurses' communication challenges during handovers: A systematic review and qualitative meta-synthesis. J Nurs Manag. 2021;29(4):623-34. doi:10.1111/jonm.13207
77. Troyer L, Brady W. Barriers to effective EMS to emergency department information transfer at patient handover: A systematic review. Am J Emerg Med. 2020;38(7):1494-503. doi:10.1016/j.ajem.2020.04.036
78. Humphrey KE, Sundberg M, Milliren CE, Graham DA, Landrigan CP. Frequency and Nature of Communication and Handoff Failures in Medical Malpractice Claims. J Patient Saf. 2022;18(2):130-7. doi:10.1097/PTS.0000000000000937
79. Starmer AJ, Spector ND, O'Toole JK, Bismilla Z, Calaman S, Campos ML, et al. Implementation of the I-PASS handoff program in diverse clinical environments: A multicenter prospective effectiveness implementation study. J Hosp Med. 2023;18(1):5-14. doi:10.1002/jhm.12979
80. McCarthy S, Motala A, Lawson E, Shekelle PG. Use of structured handoff protocols for within-hospital unit transitions: a systematic review from Making Healthcare Safer IV. BMJ Qual Saf. 2025;34(10):680-90. doi:10.1136/bmjqs-2024-018385
81. Lo L, Rotteau L, Shojania K. Can SBAR be implemented with high fidelity and does it improve communication between healthcare workers? A systematic review. BMJ Open. 2021;11(12):e055247. doi:10.1136/bmjopen-2021-055247
82. Delardes B, McLeod L, Chakraborty S, Bowles KA. What is the effect of electronic clinical handovers on patient outcomes? A systematic review. Health Informatics J. 2020;26(4):2422-34. doi:10.1177/1460458220905162

83. Guyatt G, Agoritsas T, Brignardello-Petersen R, Mustafa RA, Rylance J, Foroutan F, et al. Core GRADE 1: overview of the Core GRADE approach. BMJ. 2025;389:e081903. doi:10.1136/bmj-2024-081903
84. Nussbaumer-Streit B, Booth A, Garritty C, Hamel C, Munn Z, Tricco AC, et al. Overview of evidence synthesis types and modes. J Clin Epidemiol. 2025;187:111970. doi:10.1016/j.jclinepi.2025.111970
85. Simmonds M, Elliott JH, Synnot A, Turner T. Living Systematic Reviews. Methods Mol Biol. 2022;2345:121-34. doi:10.1007/978-1-0716-1566-9_7
86. Cheyne S, Fraile Navarro D, Hill K, McDonald S, Tunnicliffe D, White H, et al. Methods for living guidelines: early guidance based on practical experience. Paper 1: Introduction. J Clin Epidemiol. 2023;155:84-96. doi:10.1016/j.jclinepi.2022.12.024
87. Lichtner G, Alper BS, Jurth C, Spies C, Boeker M, Meerpohl JJ, et al. Representation of evidence-based clinical practice guideline recommendations on FHIR. J Biomed Inform. 2023;139:104305. doi:10.1016/j.jbi.2023.104305
88. Rim SH, Colbert J, Kobernik E, Richards TB, Miller JW, Coleman King S, et al. Building Computable Guidelines and Clinical Decision Support Tools for Cervical, Breast and Colorectal Cancer Screening and Risk Management: A CDC Initiative Update. J Womens Health (Larchmt). 2026:15409996261489070. doi:10.1177/15409996261489070
89. Adam GP, Davies MC, George J, Caputo E, Htun JM, Coppola EL, et al. Machine Learning Tools To (Semi-)Automate Evidence Synthesis: A Rapid Review and Evidence Map. Research White Paper. AHRQ Publication No. 25-EHC006. Rockville (MD): Agency for Healthcare Research and Quality; 2024 Dec. doi:10.23970/AHRQEPCWHITEPAPERMACHINE. Available from: https://www.ncbi.nlm.nih.gov/books/NBK614977/
90. Khraisha Q, Put S, Kappenberg J, Warraitch A, Hadfield K. Can large language models replace humans in systematic reviews? Evaluating GPT-4's efficacy in screening and extracting data from peer-reviewed and grey literature in multiple languages. Res Synth Methods. 2024;15(4):616-26. doi:10.1002/jrsm.1715
91. Buko T, Tuczko N, Ishikawa T. DNA Data Storage. BioTech (Basel). 2023;12(2):44. doi:10.3390/biotech12020044
92. Chow JCL. Quantum Computing in Medicine. Med Sci (Basel). 2024;12(4):67. doi:10.3390/medsci12040067
93. HL7 International. FHIR Release 5 (v5.0.0): Resource Formats [Internet]. Ann Arbor (MI): Health Level Seven International; 2023 [cited 2026 Sep 27]. Available from: https://hl7.org/fhir/formats.html
94. Dilling TJ. Artificial Intelligence Research: The Utility and Design of a Relational Database System. Adv Radiat Oncol. 2020;5(6):1280-5. doi:10.1016/j.adro.2020.06.027
95. The PostgreSQL Global Development Group. PostgreSQL 18 Documentation. Appendix M. Glossary [Internet]. The PostgreSQL Global Development Group; 2025 [cited 2026 Sep 27]. Available from: https://www.postgresql.org/docs/current/glossary.html
96. Gamal A, Barakat S, Rezk A. Standardized electronic health record data modeling and persistence: A comparative review. J Biomed Inform. 2021;114:103670. doi:10.1016/j.jbi.2020.103670
97. Griffier R, Mougin F, Jouhet V. Integrating Health Care Data in an Informatics for Integrating Biology & the Bedside (i2b2) Model Persisted Through Elasticsearch: Design, Implementation, and Evaluation in a French University Hospital. JMIR Med Inform. 2025;13:e65753. doi:10.2196/65753
98. Idowu EAA, Teo J, Salih S, Valverde J, Yeung JA. Streams, rivers and data lakes: an introduction to understanding modern electronic healthcare records. Clin Med (Lond). 2023;23(4):409. doi:10.7861/clinmed.2022-0325
99. Padoan A, Plebani M. Flowing through laboratory clinical data: the role of artificial intelligence and big data. Clin Chem Lab Med. 2022;60(12):1875-80. doi:10.1515/cclm-2022-0653
100. Li X, Wood AR, Yuan Y, Zhang M, Huang Y, Hawkes G, et al. Streamlining large-scale genomic data management: Insights from the UK Biobank whole-genome sequencing data. Cell Genom. 2025;5(12):101009. doi:10.1016/j.xgen.2025.101009
101. Reich C, Ostropolets A, Ryan P, Rijnbeek P, Schuemie M, Davydov A, et al. OHDSI Standardized Vocabularies-a large-scale centralized reference ontology for international data harmonization. J Am Med Inform Assoc. 2024;31(3):583-90. doi:10.1093/jamia/ocad247
102. Santos HP, Colaço Júnior M, Valentim R, Dos Santos JPQ, da Silva MEM, de Oliveira LAHG, et al. Scientific and technological mapping of big data architectures in healthcare. Front Public Health. 2026;14:1775715. doi:10.3389/fpubh.2026.1775715
103. Observational Health Data Sciences and Informatics (OHDSI). OMOP Common Data Model [Internet]. OHDSI; 2026 [cited 2026 Sep 27]. Available from: https://ohdsi.github.io/CommonDataModel/

104. Quiroz JC, Chard T, Sa Z, Ritchie A, Jorm L, Gallego B. Extract, transform, load framework for the conversion of health databases to OMOP. PLoS One. 2022;17(4):e0266911. doi:10.1371/journal.pone.0266911
105. Voss EA, Blacketer C, van Sandijk S, Moinat M, Kallfelz M, van Speybroeck M, et al. European Health Data & Evidence Network-learnings from building out a standardized international health data network. J Am Med Inform Assoc. 2023;31(1):209-19. doi:10.1093/jamia/ocad214
106. Hripcsak G, Schuemie MJ, Madigan D, Ryan PB, Suchard MA. Drawing Reproducible Conclusions from Observational Clinical Data with OHDSI. Yearb Med Inform. 2021;30(1):283-9. doi:10.1055/s-0041-1726481
107. Reinecke I, Zoch M, Reich C, Sedlmayr M, Bathelt F. The Usage of OHDSI OMOP - A Scoping Review. Stud Health Technol Inform. 2021;283:95-103. doi:10.3233/SHTI210546
108. i2b2 tranSMART Foundation. i2b2: Informatics for Integrating Biology and the Bedside [Internet]. i2b2 tranSMART Foundation; 2026 [cited 2026 Sep 27]. Available from: https://www.i2b2.org/
109. GO FAIR Foundation. The FAIR Guiding Principles [Internet]. GO FAIR Foundation; 2026 [cited 2026 Sep 27]. Available from: https://www.gofair.foundation/fair-principles
110. Sinaci AA, Núñez-Benjumea FJ, Gencturk M, Jauer ML, Deserno T, Chronaki C, et al. From Raw Data to FAIR Data: The FAIRification Workflow for Health Research. Methods Inf Med. 2020;59(S 01):e21-32. doi:10.1055/s-0040-1713684
111. Inau ET, Sack J, Waltemath D, Zeleke AA. Initiatives, Concepts, and Implementation Practices of the Findable, Accessible, Interoperable, and Reusable Data Principles in Health Data Stewardship: Scoping Review. J Med Internet Res. 2023;25:e45013. doi:10.2196/45013
112. Schmidt CO, Struckmann S, Enzenbach C, Reineke A, Stausberg J, Damerow S, et al. Facilitating harmonized data quality assessments. A data quality framework for observational health research data collections with software implementations in R. BMC Med Res Methodol. 2021;21(1):63. doi:10.1186/s12874-021-01252-7
113. Parums DV. Editorial: Registries and Population Databases in Clinical Research and Practice. Med Sci Monit. 2021;27:e933554. doi:10.12659/MSM.933554
114. Malenka DJ, Bhatt DL, Bradley SM, Shahian DM, Draoui J, Segawa CA, et al. The National Cardiovascular Data Registry Data Quality Program 2020: JACC State-of-the-Art Review. J Am Coll Cardiol. 2022;79(17):1704-12. doi:10.1016/j.jacc.2022.02.034
115. Feier CVI, Santoro RR, Faur AM, Muntean C, Olariu S. Assessing Changes in Colon Cancer Care during the COVID-19 Pandemic: A Four-Year Analysis at a Romanian University Hospital. J Clin Med. 2023;12(20):6558. doi:10.3390/jcm12206558
116. Feier CVI, Bardan R, Muntean C, Olariu A, Olariu S. Impact of the COVID-19 Pandemic on the Elective Surgery for Colorectal Cancer: Lessons to Be Learned. Medicina (Kaunas). 2022;58(10):1322. doi:10.3390/medicina58101322
117. International Council for Harmonisation of Technical Requirements for Pharmaceuticals for Human Use (ICH). ICH Harmonised Guideline. Guideline for Good Clinical Practice E6(R3). Final version adopted on 6 January 2025 [Internet]. Geneva: ICH; 2025 [cited 2026 Sep 27]. Available from: https://database.ich.org/sites/default/files/ICH_E6%28R3%29_Step4_FinalGuideline_2025_0106.pdf
118. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). Official Journal of the European Union. 2016 May 4;L 119:1.
119. Vanderbilt University Libraries. REDCap at Vanderbilt University [Internet]. Nashville (TN): Vanderbilt University; 2026 [cited 2026 Sep 27]. Available from: https://www.library.vanderbilt.edu/redcap/
120. Garcia KKS, Abrahão AA. Research Development Using REDCap Software. Healthc Inform Res. 2021;27(4):341-9. doi:10.4258/hir.2021.27.4.341
121. Van Bulck L, Wampers M, Moons P. Research Electronic Data Capture (REDCap): tackling data collection, management, storage, and privacy challenges. Eur J Cardiovasc Nurs. 2022;21(1):85-91. doi:10.1093/eurjcn/zvab104
122. Maré IA, Kramer B, Hazelhurst S, Nhlapho MD, Zent R, Harris PA, et al. Electronic Data Capture System (REDCap) for Health Care Research and Training in a Resource-Constrained Environment: Technology Adoption Case Study. JMIR Med Inform. 2022;10(8):e33402. doi:10.2196/33402
123. Google. How to set rules for your form. Google Docs Editors Help [Internet]. Mountain View (CA): Google; 2026 [cited 2026 Sep 27]. Available from: https://support.google.com/docs/answer/3378864

124. Google. Show questions based on answers. Google Docs Editors Help [Internet]. Mountain View (CA): Google; 2026 [cited 2026 Sep 27]. Available from: https://support.google.com/docs/answer/141062
125. Google. Choose where to save form responses. Google Docs Editors Help [Internet]. Mountain View (CA): Google; 2026 [cited 2026 Sep 27]. Available from: https://support.google.com/docs/answer/2917686
126. Zimba O, Gasparyan AY. Designing, Conducting, and Reporting Survey Studies: A Primer for Researchers. J Korean Med Sci. 2023;38(48):e403. doi:10.3346/jkms.2023.38.e403
127. Centers for Disease Control and Prevention. Epi Info [Internet]. Atlanta (GA): CDC; 2026 [updated 2026 Jun 22; cited 2026 Sep 27]. Available from: https://www.cdc.gov/epiinfo/index.html
128. Google. View and manage form responses. Google Docs Editors Help [Internet]. Mountain View (CA): Google; 2026 [cited 2026 Sep 27]. Available from: https://support.google.com/docs/answer/139706
129. Harrison JE, Weber S, Jakob R, Chute CG. ICD-11: an international classification of diseases for the twenty-first century. BMC Med Inform Decis Mak. 2021;21(Suppl 6):206. doi:10.1186/s12911-021-01534-6
130. Riley M, Lee J, Richardson S, Gjorgioski S, Robinson K. The applications of Australian-coded ICD-10 and ICD-10-AM data in research: A scoping review of the literature. Health Inf Manag. 2024;53(1):41-50. doi:10.1177/18333583231198592
131. Winter A, Ammenwerth E, Haux R, Marschollek M, Steiner B, Jahn F. Health Information Systems: Technological and Management Perspectives. 3rd ed. Cham: Springer; 2023. Chapter 3, Technological Perspective: Architecture, Integration, and Standards; p. 51-152. doi:10.1007/978-3-031-12310-8_3
132. Priyadharshini U, Vijayan R. Enhancing healthcare through ontology: a systematic review of challenges and future directions. Front Artif Intell. 2026;9:1872024. doi:10.3389/frai.2026.1872024
133. World Health Organization. International Statistical Classification of Diseases and Related Health Problems 10th Revision (ICD-10), version 2019, and Volume 2 Instruction manual, 6th ed. [Internet]. Geneva: WHO; 2019 [cited 2026 Sep 27]. Available from: https://icd.who.int/browse10/2019/en
134. World Health Organization. ICD-11 Reference Guide [Internet]. Geneva: WHO; 2022 [cited 2026 Sep 27]. Available from: https://icdcdn.who.int/icd11referenceguide/en/html/index.html
135. World Health Organization. ICD-11 Implementation: frequently asked questions [Internet]. Geneva: WHO; 2026 [cited 2026 Sep 27]. Available from: https://www.who.int/standards/classifications/frequently-asked-questions/icd-11-implementation
136. Fung KW, Xu J, Bodenreider O. The new International Classification of Diseases 11th edition: a comparative analysis with ICD-10 and ICD-10-CM. J Am Med Inform Assoc. 2020;27(5):738-46. doi:10.1093/jamia/ocaa030
137. Eastwood CA, Southern DA, Khair S, Doktorchik C, Cullen D, Ghali WA, et al. Field testing a new ICD coding system: methods and early experiences with ICD-11 Beta Version 2018. BMC Res Notes. 2022;15(1):343. doi:10.1186/s13104-022-06238-2
138. Lee H, Kim S. Impact of the ICD-11 on the accuracy of clinical coding in Korea. Health Inf Manag. 2023;52(3):221-8. doi:10.1177/18333583221095147
139. Institutul Național de Management al Serviciilor de Sănătate. Grupuri înrudite de diagnostic - clasificarea RO DRG v.1.1, versiunea martie 2024. Manual de definiții, volumul 1 [Internet]. București: INMSS; 2024 [cited 2026 Sep 27]. Available from: https://drg.inmss.ro/DocDRG/download.php?fi=28
140. World Health Organization. International Classification of Diseases (ICD) [Internet]. Geneva: WHO; 2026 [cited 2026 Sep 27]. Available from: https://www.who.int/standards/classifications/classification-of-diseases
141. World Health Organization. ICD-11 for Mortality and Morbidity Statistics, release 2026-01 [Internet]. Geneva: WHO; 2026 [cited 2026 Sep 27]. Available from: https://icd.who.int/browse/latest-release/mms/en
142. World Health Organization. ICD-11 Fact Sheet [Internet]. Geneva: WHO; 2026 May [cited 2026 Sep 27]. Available from: https://icd.who.int/en/docs/ICD11_Fact_Sheet_2026.pdf
143. Fung KW, Xu J, McConnell-Lamptey S, Pickett D, Bodenreider O. A practical strategy to use the ICD-11 for morbidity coding in the United States without a clinical modification. J Am Med Inform Assoc. 2023;30(10):1614-21. doi:10.1093/jamia/ocad128
144. Doktorchik C, Southern DA, King JA, Quan H. Applying an ICD-10 to ICD-11 mapping tool to identify causes of death codes in an Alberta dataset. Popul Health Metr. 2024;22(1):38. doi:10.1186/s12963-024-00358-6

145. Lyall K, Roberge I, Yang W, Baker S, Denny K. ICD-11 in Canada: Leveraging crosswalks to evaluate adoption, impact and transition strategies. Health Inf Manag. 2025:18333583251392355. doi:10.1177/18333583251392355
146. World Health Organization. International Classification of Functioning, Disability and Health (ICF) [Internet]. Geneva: WHO; 2026 [cited 2026 Sep 27]. Available from: https://www.who.int/standards/classifications/international-classification-of-functioning-disability-and-health
147. World Health Organization. International Classification of Health Interventions (ICHI) [Internet]. Geneva: WHO; 2026 [cited 2026 Sep 27]. Available from: https://www.who.int/standards/classifications/international-classification-of-health-interventions
148. SNOMED International. What is SNOMED CT? [Internet]. London: SNOMED International; 2026 [cited 2026 Sep 27]. Available from: https://www.snomed.org/what-is-snomed-ct
149. SNOMED International. SNOMED CT Starter Guide [Internet]. London: SNOMED International; 2026 [cited 2026 Sep 27]. Available from: https://docs.snomed.org/snomed-ct-practical-guides/snomed-ct-starter-guide
150. Chang E, Mostafa J. Use of SNOMED CT, 2020-2025: literature review. J Am Med Inform Assoc. 2026;33(9):1798-810. doi:10.1093/jamia/ocag094
151. SNOMED International. Our Members [Internet]. London: SNOMED International; 2026 [cited 2026 Sep 27]. Available from: https://www.snomed.org/members
152. Regenstrief Institute. LOINC Users Guide: 2 Major Parts of a LOINC term [Internet]. Indianapolis (IN): Regenstrief Institute; 2026 Aug 17 [cited 2026 Sep 27]. Available from: https://loinc.org/kb/users-guide/major-parts-of-a-loinc-term/
153. Regenstrief Institute. LOINC 2345-7 Glucose [Mass/volume] in Serum or Plasma [Internet]. Indianapolis (IN): Regenstrief Institute; 2026 [cited 2026 Sep 27]. Available from: https://loinc.org/2345-7
154. WHO Collaborating Centre for Drug Statistics Methodology. ATC: structure and principles [Internet]. Oslo: Norwegian Institute of Public Health; 2026 [cited 2026 Sep 27]. Available from: https://atcddd.fhi.no/atc/structure_and_principles/
155. WHO Collaborating Centre for Drug Statistics Methodology. DDD: definition and general considerations [Internet]. Oslo: Norwegian Institute of Public Health; 2026 [cited 2026 Sep 27]. Available from: https://atcddd.fhi.no/ddd/definition_and_general_considera/
156. National Library of Medicine (US). Medical Subject Headings (MeSH) [Internet]. Bethesda (MD): NLM; 2026 [cited 2026 Sep 27]. Available from: https://www.nlm.nih.gov/mesh/meshhome.html
157. National Library of Medicine (US). Unified Medical Language System (UMLS) [Internet]. Bethesda (MD): NLM; 2026 [cited 2026 Sep 27]. Available from: https://www.nlm.nih.gov/research/umls/index.html
158. Pi S, Masterson J, Ma SP, Corbin CK, Milstein A, Chen JH. Using Case Mix Index within Diagnosis-Related Groups to Evaluate Variation in Hospitalization Costs at a Large Academic Medical Center. AMIA Annu Symp Proc. 2023;2023:1201-8.
159. Liu J, Capurro D, Nguyen A, Verspoor K. Early prediction of diagnostic-related groups and estimation of hospital cost by processing clinical notes. NPJ Digit Med. 2021;4(1):103. doi:10.1038/s41746-021-00474-9
160. Chen YJ, Zhang XY, Yan JQ, Xue-Tang, Qian MC, Ying XH. Impact of Diagnosis-Related Groups on Inpatient Quality of Health Care: A Systematic Review and Meta-Analysis. Inquiry. 2023;60:469580231167011. doi:10.1177/00469580231167011
161. Campbell AC, Tyler J, Shuttleworth RR, Pearce LA, Heerde JA, Borschmann R, et al. Accuracy of administrative data in ascertaining health conditions: a systematic review. JAMIA Open. 2026;9(3):ooag109. doi:10.1093/jamiaopen/ooag109
162. Gabel E, Gal J, Grogan T, Hofer I. A retrospective analysis using comorbidity detecting algorithmic software to determine the incidence of International Classification of Diseases (ICD) code omissions and appropriateness of Diagnosis-Related Group (DRG) code modifiers. BMC Med Inform Decis Mak. 2024;24(1):309. doi:10.1186/s12911-024-02724-8
163. Crespin D, Dworsky M, Levin J, Ruder T, Whaley CM. Upcoding Linked To Up To Two-Thirds Of Growth In Highest-Intensity Hospital Discharges In 5 States, 2011-19. Health Aff (Millwood). 2024;43(12):1619-27. doi:10.1377/hlthaff.2024.00596
164. Khalid AR, Ali H, Fathima K, Sonia Carole K, Kim HC. Recent Advances in AI for Automated ICD Coding: A Systematic Literature Review. J Med Syst. 2026;50(1):113. doi:10.1007/s10916-026-02429-7

165. Soroush A, Glicksberg BS, Zimlichman E, Barash Y, Freeman R, Charney AW, et al. Large Language Models Are Poor Medical Coders - Benchmarking of Medical Code Querying. NEJM AI. 2024;1(5). doi:10.1056/aidbp2300040
166. Naliyatthaliyazchayil P, Muthyala R, Gichoya JW, Purkayastha S. Evaluating the Reasoning Capabilities of Large Language Models for Medical Coding and Hospital Readmission Risk Stratification: Zero-Shot Prompting Approach. J Med Internet Res. 2025;27:e74142. doi:10.2196/74142
167. Galland-Decker C, Ursenbacher M, Nunes C, Bouche F, Blanchard M, Carra G, et al. Improving Meditron for Medical Coding Through Fine-Tuning: A Comparative Evaluation Against GPT-4. Stud Health Technol Inform. 2026;336:1072-3. doi:10.3233/SHTI260356
168. Chomutare T, Svenning TO, Hernández MÁT, Ngo PD, Budrionis A, Markljung K, et al. Artificial Intelligence to Improve Clinical Coding Practice in Scandinavia: Crossover Randomized Controlled Trial. J Med Internet Res. 2025;27:e71904. doi:10.2196/71904
169. Friedman Y, Smajdor A. Synthetic Medical Data vs. The Four Principles: Ethical Trade-Offs and Zombie Data. Health Care Anal. 2026. doi:10.1007/s10728-026-00596-9
170. Goodman KW, Miller RA. Ethics in Biomedical and Health Informatics: Users, Standards, and Outcomes. In: Shortliffe EH, Cimino JJ, editors. Biomedical Informatics: Computer Applications in Health Care and Biomedicine. 5th ed. Cham: Springer; 2021. p. 391-423. doi:10.1007/978-3-030-58721-5_12
171. Legea nr. 46/2003 a drepturilor pacientului. Monitorul Oficial al României, Partea I. 2003 Jan 29;(51).
172. Cohen E, Byrom B, Becher A, Jörntén-Karlsson M, Mackenzie AK. Comparative Effectiveness of eConsent: Systematic Review. J Med Internet Res. 2023;25:e43883. doi:10.2196/43883
173. Zenker S, Strech D, Ihrig K, Jahns R, Müller G, Schickhardt C, et al. Data protection-compliant broad consent for secondary use of health care data and human biosamples for (bio)medical research: Towards a new German national standard. J Biomed Inform. 2022;131:104096. doi:10.1016/j.jbi.2022.104096
174. Lee AR, Koo D, Kim IK, Lee E, Kim HH, Yoo S, et al. Identifying facilitators of and barriers to the adoption of dynamic consent in digital health ecosystems: a scoping review. BMC Med Ethics. 2023;24(1):107. doi:10.1186/s12910-023-00988-9
175. Phuyal S, Bhandari M, Correia Bezerra R, Bista R, Ferreira JC. Dynamic Consent for Secondary Use of Health Data: Challenges and Opportunities Under European Law. JMIR Med Inform. 2026;14:e93348. doi:10.2196/93348
176. Legea nr. 190/2018 privind măsuri de punere în aplicare a Regulamentului (UE) 2016/679 al Parlamentului European și al Consiliului din 27 aprilie 2016 privind protecția persoanelor fizice în ceea ce privește prelucrarea datelor cu caracter personal și privind libera circulație a acestor date și de abrogare a Directivei 95/46/CE (Regulamentul general privind protecția datelor). Monitorul Oficial al României, Partea I. 2018 Jul 26;(651).
177. Olatunji IE, Rauch J, Katzensteiner M, Khosla M. A Review of Anonymization for Healthcare Data. Big Data. 2024;12(6):538-55. doi:10.1089/big.2021.0169
178. Packhäuser K, Gündel S, Münster N, Syben C, Christlein V, Maier A. Deep learning-based patient re-identification is able to exploit the biometric nature of medical chest X-ray data. Sci Rep. 2022;12(1):14851. doi:10.1038/s41598-022-19045-3
179. Xia W, Liu Y, Wan Z, Vorobeychik Y, Kantacioglu M, Nyemba S, et al. Enabling realistic health data re-identification risk assessment through adversarial modeling. J Am Med Inform Assoc. 2021;28(4):744-52. doi:10.1093/jamia/ocaa327
180. European Union Agency for Cybersecurity (ENISA). ENISA Threat Landscape: Health Sector (January 2021 to March 2023) [Internet]. Athens: ENISA; 2023 Jul [cited 2026 Sep 27]. Available from: https://www.enisa.europa.eu/publications/health-threat-landscape
181. Jalali MS, Bruckes M, Westmattelmann D, Schewe G. Why Employees (Still) Click on Phishing Links: Investigation in Hospitals. J Med Internet Res. 2020;22(1):e16775. doi:10.2196/16775
182. Herrera Montano I, Góngora Alonso S, Sañudo García S, García Aranda JJ, Rodrígues JJPC, de la Torre Díez I. Security practices and insider threats in Spanish healthcare centers: a survey-based risk assessment. Int J Med Inform. 2026;205:106107. doi:10.1016/j.ijmedinf.2025.106107
183. Neprash HT, McGlave CC, Cross DA, Virnig BA, Puskarich MA, Huling JD, et al. Trends in Ransomware Attacks on US Hospitals, Clinics, and Other Health Care Delivery Organizations, 2016-2021. JAMA Health Forum. 2022;3(12):e224873. doi:10.1001/jamahealthforum.2022.4873
184. Dameff C, Tully J, Chan TC, Castillo EM, Savage S, Maysent P, et al. Ransomware Attack Associated With Disruptions at Adjacent Emergency Departments in the US. JAMA Netw Open. 2023;6(5):e2312270. doi:10.1001/jamanetworkopen.2023.12270

185. Pham TT, Loo TM, Malhotra A, Longhurst CA, Hylton D, Dameff C, et al. Ransomware Cyberattack Associated With Cardiac Arrest Incidence and Outcomes at Untargeted, Adjacent Hospitals. Crit Care Explor. 2024;6(4):e1079. doi:10.1097/CCE.0000000000001079
186. Directoratul Național de Securitate Cibernetică. UPDATE: Un atac cibernetic de tip ransomware a afectat spitale din România [Internet]. București: DNSC; 2024 Feb 15 [updated 2024 Feb 16; cited 2026 Sep 27]. Available from: https://www.dnsc.ro/citeste/atac-cibernetic-ransomware-spitale-Romania. Archived at: http://web.archive.org/web/20240419005822/https://dnsc.ro/citeste/atac-cibernetic-ransomware-spitale-Romania
187. Ordonanța de urgență a Guvernului nr. 155/2024 privind instituirea unui cadru pentru securitatea cibernetică a rețelelor și sistemelor informatice din spațiul cibernetic național civil. Monitorul Oficial al României, Partea I. 2024 Dec 31;(1332).
188. Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union, amending Regulation (EU) No 910/2014 and Directive (EU) 2018/1972, and repealing Directive (EU) 2016/1148 (NIS 2 Directive). Official Journal of the European Union. 2022 Dec 27;L 333:80.
189. ISO 27799:2025. Health informatics - Information security controls in health based on ISO/IEC 27002. 3rd ed. Geneva: International Organization for Standardization; 2025.
190. ISO/DIS 27799:2025. Health informatics - Information security management in health using ISO/IEC 27002. Draft International Standard (voting 2025 Jan 20 - 2025 Apr 14). Geneva: International Organization for Standardization; 2025.
191. Chen IY, Pierson E, Rose S, Joshi S, Ferryman K, Ghassemi M. Ethical Machine Learning in Healthcare. Annu Rev Biomed Data Sci. 2021;4:123-44. doi:10.1146/annurev-biodatasci-092820-114757
192. Sjoding MW, Dickson RP, Iwashyna TJ, Gay SE, Valley TS. Racial Bias in Pulse Oximetry Measurement. N Engl J Med. 2020;383(25):2477-8. doi:10.1056/NEJMc2029240
193. Parr NJ, Beech EH, Young S, Valley TS. Racial and Ethnic Disparities in Occult Hypoxemia Prevalence and Clinical Outcomes Among Hospitalized Patients: A Systematic Review and Meta-analysis. J Gen Intern Med. 2024;39(13):2543-53. doi:10.1007/s11606-024-08852-1
194. Seyyed-Kalantari L, Zhang H, McDermott MBA, Chen IY, Ghassemi M. Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. Nat Med. 2021;27(12):2176-82. doi:10.1038/s41591-021-01595-0
195. Inker LA, Eneanya ND, Coresh J, Tighiouart H, Wang D, Sang Y, et al. New Creatinine- and Cystatin C-Based Equations to Estimate GFR without Race. N Engl J Med. 2021;385(19):1737-49. doi:10.1056/NEJMoa2102953
196. Alderman JE, Palmer J, Laws E, McCradden MD, Ordish J, Ghassemi M, et al. Tackling algorithmic bias and promoting transparency in health datasets: the STANDING Together consensus recommendations. Lancet Digit Health. 2025;7(1):e64-88. doi:10.1016/S2589-7500(24)00224-3
197. Regulation (EU) 2026/1744 of the European Parliament and of the Council of 8 July 2026 amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI). Official Journal of the European Union. 2026 Jul 24;L 2026/1744.
198. Legea nr. 199/2023 a învățământului superior. Monitorul Oficial al României, Partea I. 2023 Jul 5;(614).
199. Fadlalmola HA, Elhusein AM, Swamy DSV, Hussein MK, Mamanao DM, Mohamedsalih WE. Plagiarism among nursing students: A systematic review and meta-analysis. Int Nurs Rev. 2022;69(4):492-502. doi:10.1111/inr.12755
200. Habibzadeh F. Plagiarism: A Bird's Eye View. J Korean Med Sci. 2023;38(45):e373. doi:10.3346/jkms.2023.38.e373
201. National Library of Medicine (US). Samples of Formatted References for Authors of Journal Articles [Internet]. Bethesda (MD): National Library of Medicine; [last reviewed 2026 Jun 2; cited 2026 Sep 27]. Available from: https://www.nlm.nih.gov/bsd/uniform_requirements.html
202. International Committee of Medical Journal Editors. Recommendations for the Conduct, Reporting, Editing, and Publication of Scholarly Work in Medical Journals [Internet]. [place unknown]: ICMJE; [updated 2026 Jan; cited 2026 Sep 27]. Available from: https://www.icmje.org/recommendations/
203. Committee on Publication Ethics. Authorship and AI tools. COPE position statement [Internet]. COPE; 2023 Feb 13 [cited 2026 Sep 27]. Available from: https://publicationethics.org/guidance/cope-position/authorship-and-ai-tools
204. Ganjavi C, Eppler MB, Pekcan A, Biedermann B, Abreu A, Collins GS, et al. Publishers' and journals' instructions to authors on use of generative artificial intelligence in academic and scientific publishing: bibliometric analysis. BMJ. 2024;384:e077192. doi:10.1136/bmj-2023-077192

205. Liang W, Yuksekgonul M, Mao Y, Wu E, Zou J. GPT detectors are biased against non-native English writers. Patterns (N Y). 2023;4(7):100779. doi:10.1016/j.patter.2023.100779
206. Herasevich V, Pickering BW, Clemmer TP, Mark RG. Patient Monitoring Systems. In: Shortliffe EH, Cimino JJ, editors. Biomedical Informatics: Computer Applications in Health Care and Biomedicine. 5th ed. Cham: Springer; 2021. p. 693-732. doi:10.1007/978-3-030-58721-5_21
207. Barmpagiannos K, Malik Z, Whyte S, Stavrakis S. Quantifying high-pass filter-induced ST-segment distortion. J Electrocardiol. 2026;99:154442. doi:10.1016/j.jelectrocard.2026.154442
208. Peltola ME, Leitinger M, Halford JJ, Vinayan KP, Kobayashi K, Pressler RM, et al. Routine and sleep EEG: Minimum recording standards of the International Federation of Clinical Neurophysiology and the International League Against Epilepsy. Epilepsia. 2023;64(3):602-18. doi:10.1111/epi.17448
209. Nayak CS, Anilkumar AC. Normal EEG Waveforms. In: StatPearls [Internet]. Treasure Island (FL): StatPearls Publishing; 2025 [updated 2025 Aug 3; cited 2026 Sep 27]. Available from: https://www.ncbi.nlm.nih.gov/books/NBK539805/
210. Sul JH, Piyathilaka L, Moratuwage D, Dunu Arachchige S, Jayawardena A, Kahandawa G, et al. Electromyography Signal Acquisition, Filtering, and Data Analysis for Exoskeleton Development. Sensors (Basel). 2025;25(13):4004. doi:10.3390/s25134004
211. Park J, Seok HS, Kim SS, Shin H. Photoplethysmogram Analysis and Applications: An Integrative Review. Front Physiol. 2021;12:808451. doi:10.3389/fphys.2021.808451
212. Quigley KS, Gianaros PJ, Norman GJ, Jennings JR, Berntson GG, de Geus EJC. Publication guidelines for human heart rate and heart rate variability studies in psychophysiology-Part 1: Physiological underpinnings and foundations of measurement. Psychophysiology. 2024;61(9):e14604. doi:10.1111/psyp.14604
213. PhysioNet. MIT-BIH Arrhythmia Database, version 1.0.0 [Internet]. Cambridge (MA): MIT Laboratory for Computational Physiology; 2026 [cited 2026 Sep 27]. Available from: https://physionet.org/content/mitdb/1.0.0/
214. Rahman S, Karmakar C, Natgunanathan I, Yearwood J, Palaniswami M. Robustness of electrocardiogram signal quality indices. J R Soc Interface. 2022;19(189):20220012. doi:10.1098/rsif.2022.0012
215. Babiloni C, Barry RJ, Başar E, Blinowska KJ, Cichocki A, Drinkenburg WHIM, et al. International Federation of Clinical Neurophysiology (IFCN) - EEG research workgroup: Recommendations on frequency and topographic analysis of resting state EEG rhythms. Part 1: Applications in clinical research studies. Clin Neurophysiol. 2020;131(1):285-307. doi:10.1016/j.clinph.2019.06.234
216. EDFbrowser, version 2.15. Free, open-source viewer for EDF/EDF+/BDF files [Internet]. Teuniz.net; 2026 [cited 2026 Sep 27]. Available from: https://www.teuniz.net/edfbrowser/
217. Weinstein A, Rodiño J, Otero M. Effects of electrocardiogram qrs detection algorithms in heart rate variability metrics. Sci Rep. 2026;16(1):18589. doi:10.1038/s41598-026-49215-6
218. Littmann L. Electrocardiographic artifact. J Electrocardiol. 2021;64:23-9. doi:10.1016/j.jelectrocard.2020.11.006
219. Amin U, Nascimento FA, Karakis I, Schomer D, Benbadis SR. Normal variants and artifacts: Importance in EEG interpretation. Epileptic Disord. 2023;25(5):591-648. doi:10.1002/epd2.20040
220. van der Bijl K, Elgendi M, Menon C. Automatic ECG Quality Assessment Techniques: A Systematic Review. Diagnostics (Basel). 2022;12(11):2578. doi:10.3390/diagnostics12112578
221. European Data Format. EDF and EDF+ [Internet]. 2026 [cited 2026 Sep 27]. Available from: https://www.edfplus.info/
222. WFDB: Waveform Database. Open standards and software for physiologic signal and annotation data [Internet]. 2026 [cited 2026 Sep 27]. Available from: https://wfdb.io/
223. National Electrical Manufacturers Association. DICOM PS3.3 2026d - Information Object Definitions. Section A.34: Waveform IODs [Internet]. Rosslyn (VA): NEMA; 2026 [cited 2026 Sep 27]. Available from: https://dicom.nema.org/medical/dicom/current/output/chtml/part03/sect_A.34.html
224. Health Level Seven International. HL7 Version 3 Implementation Guide: Annotated ECG (aECG) R1, Release 2 - US Realm. Product brief [Internet]. Ann Arbor (MI): HL7 International; c2007-2026 [cited 2026 Sep 27]. Available from: https://www.hl7.org/implement/standards/product_brief.cfm?product_id=102
225. ISO/IEEE 11073-10103:2025. Health informatics - Device interoperability - Part 10103: Nomenclature - Implantable device, cardiac. Geneva: International Organization for Standardization; 2025.
226. PhysioNet. The research resource for complex physiologic signals [Internet]. Cambridge (MA): MIT Laboratory for Computational Physiology; 2026 [cited 2026 Sep 27]. Available from: https://physionet.org/

227. Pollard T, Moody BE, Lehman LwH, Gow BJ, Fernandes C, Xie C, et al. PhysioNet as a global platform for biomedical research. Nat Health. 2026;1(8):792-5. doi:10.1038/s44360-026-00096-z
228. PhysioNet. LightWAVE 0.71: waveform and annotation viewer [Internet]. Cambridge (MA): MIT Laboratory for Computational Physiology; 2026 [cited 2026 Sep 27]. Available from: https://physionet.org/lightwave/
229. Ose B, Sattar Z, Gupta A, Toquica C, Harvey C, Noheria A. Artificial Intelligence Interpretation of the Electrocardiogram: A State-of-the-Art Review. Curr Cardiol Rep. 2024;26(6):561-80. doi:10.1007/s11886-024-02062-1
230. Yao X, Rushlow DR, Inselman JW, McCoy RG, Thacher TD, Behnken EM, et al. Artificial intelligence-enabled electrocardiograms for identification of patients with low ejection fraction: a pragmatic, randomized clinical trial. Nat Med. 2021;27(5):815-9. doi:10.1038/s41591-021-01335-4
231. Tsai DJ, Lin C, Liu WT, Lee CC, Chang CH, Lin WY, et al. Artificial intelligence-assisted diagnosis and prognostication in low ejection fraction using electrocardiograms in inpatient department: a pragmatic randomized controlled trial. BMC Med. 2025;23(1):342. doi:10.1186/s12916-025-04190-z
232. Adedinsewo DA, Morales-Lara AC, Afolabi BB, Kushimo OA, Mbakwem AC, Ibiyemi KF, et al. Artificial intelligence guided screening for cardiomyopathies in an obstetric population: a pragmatic randomized clinical trial. Nat Med. 2024;30(10):2897-906. doi:10.1038/s41591-024-03243-9
233. Christopoulos G, Graff-Radford J, Lopez CL, Yao X, Attia ZI, Rabinstein AA, et al. Artificial Intelligence-Electrocardiography to Predict Incident Atrial Fibrillation: A Population-Based Study. Circ Arrhythm Electrophysiol. 2020;13(12):e009355. doi:10.1161/CIRCEP.120.009355
234. Charlton PH, Allen J, Bailón R, Baker S, Behar JA, Chen F, et al. The 2023 wearable photoplethysmography roadmap. Physiol Meas. 2023;44(11):111001. doi:10.1088/1361-6579/acead2
235. Lubitz SA, Faranesh AZ, Selvaggi C, Atlas SJ, McManus DD, Singer DE, et al. Detection of Atrial Fibrillation in a Large Population Using Wearable Devices: The Fitbit Heart Study. Circulation. 2022;146(19):1415-24. doi:10.1161/CIRCULATIONAHA.122.060291
236. Lubitz SA, McConnell MV, Selvaggi C, Krishnamoorthy A, Atlas SJ, McManus DD, et al. Wearable Irregular Heart Rhythm Detection Recurrences and Electrocardiographic Atrial Fibrillation Confirmation: The Fitbit Heart Study. Circ Arrhythm Electrophysiol. 2025;18(7):e013565. doi:10.1161/CIRCEP.124.013565
237. van Steijn NJ, Blommestijn IS, Blok S, Pepplinkhuizen S, Somsen GA, Knops RE, et al. Enhanced Detection and Prompt Diagnosis of Atrial Fibrillation Using Apple Watch: A Randomized Controlled Trial. J Am Coll Cardiol. 2026;87(14):1714-28. doi:10.1016/j.jacc.2025.11.032
238. Barrera N, Solorzano M, Jimenez Y, Kushnir Y, Gallegos-Koyner F, Dagostin de Carvalho G. Accuracy of Smartwatches in the Detection of Atrial Fibrillation: A Systematic Review and Diagnostic Meta-Analysis. JACC Adv. 2025;4(11 Pt 1):102133. doi:10.1016/j.jacadv.2025.102133
239. Brandes A, Stavrakis S, Freedman B, Antoniou S, Boriani G, Camm AJ, et al. Consumer-Led Screening for Atrial Fibrillation: Frontier Review of the AF-SCREEN International Collaboration. Circulation. 2022;146(19):1461-74. doi:10.1161/CIRCULATIONAHA.121.058911
240. Haase CB, Jensen AE, Modin FA, Siersma V, Brodersen JB. Overdiagnosis in atrial fibrillation screening with wearables. Scand J Prim Health Care. 2026;44(1):2656694. doi:10.1080/02813432.2026.2656694
241. Noseworthy PA, Attia ZI, Brewer LC, Hayes SN, Yao X, Kapa S, et al. Assessing and Mitigating Bias in Medical Artificial Intelligence: The Effects of Race and Ethnicity on a Deep Learning Model for ECG Analysis. Circ Arrhythm Electrophysiol. 2020;13(3):e007988. doi:10.1161/CIRCEP.119.007988
242. Singh S, Bennett MR, Chen C, Shin S, Ghanbari H, Nelson BW. Impact of Skin Pigmentation on Pulse Oximetry Blood Oxygenation and Wearable Pulse Rate Accuracy: Systematic Review and Meta-Analysis. J Med Internet Res. 2024;26:e62769. doi:10.2196/62769
243. Muntean C, Gaborean V, Faur AM, Faur IF, Prodan-Bărbulescu C, Feier CVI. Continuous Wearable-Sensor Monitoring After Colorectal Surgery: A Systematic Review of Clinical Outcomes and Predictive Analytics. Diagnostics (Basel). 2025;15(17):2194. doi:10.3390/diagnostics15172194
244. Ministerul Sănătății. Caiet de sarcini pentru servicii de dezvoltare și implementare sistem informatic integrat pentru evidența clinică a secțiilor A.T.I. (S.I.E.C.-A.T.I.) [Internet]. București: Ministerul Sănătății; 2021 [cited 2026 Sep 27]. Available from: https://ms.ro/media/documents/Caiet-de-Sarcini-Sist-integrat-SIEC-ATI.pdf
245. National Electrical Manufacturers Association. Digital Imaging and Communications in Medicine (DICOM) Standard, PS3.1-PS3.22, 2026d edition [Internet]. Rosslyn (VA): NEMA; 2026 [cited 2026 Sep 27]. Available from: https://www.dicomstandard.org/current
246. Dhamija E, Gulati M, Deo SVS, Gogia A, Hari S. Digital Breast Tomosynthesis: an Overview. Indian J Surg Oncol. 2021;12(2):315-29. doi:10.1007/s13193-021-01310-y

247. Cornish TC, Dayal S, Ferber D, Chiweshe J, Monroe R. The evolving role of DICOM in digital pathology. J Pathol Inform. 2026;21:100658. doi:10.1016/j.jpi.2026.100658
248. Evans AJ, Brown RW, Bui MM, Chlipala EA, Lacchetti C, Milner DA, et al. Validating Whole Slide Imaging Systems for Diagnostic Purposes in Pathology. Arch Pathol Lab Med. 2022;146(4):440-50. doi:10.5858/arpa.2020-0723-CP
249. Jo YY, Choi YS, Park HW, Lee JH, Jung H, Kim HE, et al. Impact of image compression on deep learning-based mammogram classification. Sci Rep. 2021;11(1):7924. doi:10.1038/s41598-021-86726-w
250. Trieu PDY, Barron M, Lewis SJ. Use of Full-quality DICOM Images Compared to Minimally Compressed Mammograms in JPEG Format for Radiology Training: A Study From Radiologist and Radiographer Perspectives. Acad Radiol. 2023;30(8):1748-55. doi:10.1016/j.acra.2022.11.012
251. Fischer M, Neher P, Schüffler P, Ziegler S, Xiao S, Peretzke R, et al. Unlocking the potential of digital pathology: Novel baselines for compression. J Pathol Inform. 2025;17:100421. doi:10.1016/j.jpi.2025.100421
252. ISO/DIS 12052:2025. Health informatics - Digital imaging and communication in medicine (DICOM) including workflow and data management. Draft International Standard. Geneva: International Organization for Standardization; 2025.
253. ISO 12052:2026. Health informatics - Digital imaging and communication in medicine (DICOM) including workflow and data management. 3rd ed. Geneva: International Organization for Standardization; 2026.
254. National Electrical Manufacturers Association. DICOMweb [Internet]. Rosslyn (VA): NEMA; 2026 [cited 2026 Sep 27]. Available from: https://www.dicomstandard.org/using/dicomweb
255. IHE International. Scheduled Workflow (SWF) [Internet]. IHE Wiki; [cited 2026 Sep 27]. Available from: https://wiki.ihe.net/index.php/Scheduled_Workflow
256. Tang ST, Tjia V, Noga T, Febri J, Lien CY, Chu WC, et al. Creating a Medical Imaging Workflow Based on FHIR, DICOMweb, and SVG. J Digit Imaging. 2023;36(3):794-803. doi:10.1007/s10278-021-00522-6
257. Clunie DA, Napora B, Groth J, Zerbe N, Yousif M, Shelar G, et al. Report on the 2025 DICOM WSI Connectathon. J Pathol Inform. 2026;21:100657. doi:10.1016/j.jpi.2026.100657
258. Clunie D, Prior F, Rutherford M, Moore S, Parker W, Kondylakis H, et al. Summary of the National Cancer Institute 2023 Virtual Workshop on Medical Image De-identification-Part 1: Report of the MIDI Task Group - Best Practices and Recommendations, Tools for Conventional Approaches to De-identification, International Approaches to De-identification, and Industry Panel on Image De-identification. J Imaging Inform Med. 2025;38(1):1-15. doi:10.1007/s10278-024-01182-y
259. Eichelberg M, Kleber K, Kämmerer M. Cybersecurity in PACS and Medical Imaging: an Overview. J Digit Imaging. 2020;33(6):1527-42. doi:10.1007/s10278-020-00393-3
260. Ghaderi H, April A. Assessing Available Open-Source PACS Options. J Digit Imaging. 2023;36(6):2323-8. doi:10.1007/s10278-021-00435-4
261. IHE International. IHE Radiology (RAD) Technical Framework, Revision 23.0 [Internet]. IHE International; 2025 Aug 8 [cited 2026 Sep 27]. Available from: https://profiles.ihe.net/RAD/index.html
262. Kawa J, Pyciński B, Smoliński M, Bożek P, Kwasecki M, Pietrzyk B, et al. Design and Implementation of a Cloud PACS Architecture. Sensors (Basel). 2022;22(21):8569. doi:10.3390/s22218569
263. Isensee F, Jaeger PF, Kohl SAA, Petersen J, Maier-Hein KH. nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. Nat Methods. 2021;18(2):203-11. doi:10.1038/s41592-020-01008-z
264. Wasserthal J, Breit HC, Meyer MT, Pradella M, Hinck D, Sauter AW, et al. TotalSegmentator: Robust Segmentation of 104 Anatomic Structures in CT Images. Radiol Artif Intell. 2023;5(5):e230024. doi:10.1148/ryai.230024
265. van Timmeren JE, Cester D, Tanadini-Lang S, Alkadhi H, Baessler B. Radiomics in medical imaging-"how-to" guide and critical reflection. Insights Imaging. 2020;11(1):91. doi:10.1186/s13244-020-00887-2
266. Zwanenburg A, Vallières M, Abdalah MA, Aerts HJWL, Andrearczyk V, Apte A, et al. The Image Biomarker Standardization Initiative: Standardized Quantitative Radiomics for High-Throughput Image-based Phenotyping. Radiology. 2020;295(2):328-38. doi:10.1148/radiol.2020191145
267. U.S. Food and Drug Administration. Artificial Intelligence-Enabled Medical Devices [Internet]. Silver Spring (MD): FDA; 2026 Sep 4 [cited 2026 Sep 27]. Available from: https://www.fda.gov/medical-devices/artificial-intelligence-enabled-medical-devices/list-artificial-intelligence-enabled-medical-devices
268. Lee YS, Youn BY. Clinical Specialty Expansion of AI-Enabled and Machine Learning-Enabled Medical Devices Authorized by the US Food and Drug Administration From 1995 to 2025: Longitudinal Content Analysis. J Med Internet Res. 2026;28:e103040. doi:10.2196/103040

269. Lång K, Josefsson V, Larsson AM, Larsson S, Högberg C, Sartor H, et al. Artificial intelligence-supported screen reading versus standard double reading in the Mammography Screening with Artificial Intelligence trial (MASAI): a clinical safety analysis of a randomised, controlled, non-inferiority, single-blinded, screening accuracy study. Lancet Oncol. 2023;24(8):936-44. doi:10.1016/S1470-2045(23)00298-X
270. Hernström V, Josefsson V, Sartor H, Schmidt D, Larsson AM, Hofvind S, et al. Screening performance and characteristics of breast cancer detected in the Mammography Screening with Artificial Intelligence trial (MASAI): a randomised, controlled, parallel-group, non-inferiority, single-blinded, screening accuracy study. Lancet Digit Health. 2025;7(3):e175-83. doi:10.1016/S2589-7500(24)00267-X
271. Gommers J, Hernström V, Josefsson V, Sartor H, Schmidt D, Hjelmgren A, et al. Interval cancer, sensitivity, and specificity comparing AI-supported mammography screening with standard double reading without AI in the MASAI study: a randomised, controlled, non-inferiority, single-blinded, population-based, screening-accuracy trial. Lancet. 2026;407(10527):505-14. doi:10.1016/S0140-6736(25)02464-X
272. Eisemann N, Bunk S, Mukama T, Baltus H, Elsner SA, Gomille T, et al. Nationwide real-world implementation of AI for cancer detection in population-based mammography screening. Nat Med. 2025;31(3):917-24. doi:10.1038/s41591-024-03408-6
273. Geppert J, Asgharzadeh A, Brown A, Stinton C, Helm EJ, Jayakody S, et al. Software using artificial intelligence for nodule and cancer detection in CT lung cancer screening: systematic review of test accuracy studies. Thorax. 2024;79(11):1040-9. doi:10.1136/thorax-2024-221662
274. Geppert J, Auguste P, Asgharzadeh A, Ghiasvand H, Patel M, Brown A, et al. Software with artificial intelligence-derived algorithms for detecting and analysing lung nodules in CT scans: systematic review and economic evaluation. Health Technol Assess. 2025;29(14):1-234. doi:10.3310/JYTW8921
275. Yu AC, Mohajer B, Eng J. External Validation of Deep Learning Algorithms for Radiologic Diagnosis: A Systematic Review. Radiol Artif Intell. 2022;4(3):e210064. doi:10.1148/ryai.210064
276. Kelly BS, Judge C, Bollard SM, Clifford SM, Healy GM, Aziz A, et al. Radiology artificial intelligence: a systematic review and evaluation of methods (RAISE). Eur Radiol. 2022;32(11):7998-8007. doi:10.1007/s00330-022-08784-6
277. McGenity C, Clarke EL, Jennings C, Matthews G, Cartlidge C, Freduah-Agyemang H, et al. Artificial intelligence in digital pathology: a systematic review and meta-analysis of diagnostic test accuracy. NPJ Digit Med. 2024;7(1):114. doi:10.1038/s41746-024-01106-8
278. Komura D, Ochi M, Ishikawa S. Machine learning methods for histopathological image analysis: Updates in 2024. Comput Struct Biotechnol J. 2025;27:383-400. doi:10.1016/j.csbj.2024.12.033
279. Varoquaux G, Cheplygina V. Machine learning for medical imaging: methodological failures and recommendations for the future. NPJ Digit Med. 2022;5(1):48. doi:10.1038/s41746-022-00592-y
280. Tejani AS, Ng YS, Xi Y, Rayan JC. Understanding and Mitigating Bias in Imaging Artificial Intelligence. Radiographics. 2024;44(5):e230067. doi:10.1148/rg.230067
281. Dratsch T, Chen X, Rezazade Mehrizi M, Kloeckner R, Mähringer-Kunz A, Püsken M, et al. Automation Bias in Mammography: The Impact of Artificial Intelligence BI-RADS Suggestions on Reader Performance. Radiology. 2023;307(4):e222176. doi:10.1148/radiol.222176
282. Bernstein MH, Atalay MK, Dibble EH, Maxwell AWP, Karam AR, Agarwal S, et al. Can incorrect artificial intelligence (AI) results impact radiologists, and if so, what can we do about it? A multi-reader pilot study of lung cancer detection with chest radiography. Eur Radiol. 2023;33(11):8263-9. doi:10.1007/s00330-023-09747-1
283. Yu F, Moehring A, Banerjee O, Salz T, Agarwal N, Rajpurkar P. Heterogeneity and predictors of the effects of AI assistance on radiologists. Nat Med. 2024;30(3):837-49. doi:10.1038/s41591-024-02850-w
284. Regulation (EU) 2017/745 of the European Parliament and of the Council of 5 April 2017 on medical devices, amending Directive 2001/83/EC, Regulation (EC) No 178/2002 and Regulation (EC) No 1223/2009 and repealing Council Directives 90/385/EEC and 93/42/EEC. Official Journal of the European Union. 2017 May 5;L 117:1.
285. Melton GB, McDonald CJ, Tang PC, Hripcsak G. Electronic Health Records. In: Shortliffe EH, Cimino JJ, editors. Biomedical Informatics: Computer Applications in Health Care and Biomedicine. 5th ed. Cham: Springer; 2021. p. 467-509. doi:10.1007/978-3-030-58721-5_14
286. Brands MR, Gouw SC, Beestrum M, Cronin RM, Fijnvandraat K, Badawy SM. Patient-Centered Digital Health Records and Their Effects on Health Outcomes: Systematic Review. J Med Internet Res. 2022;24(12):e43086. doi:10.2196/43086
287. Wright A, Schreiber R, Bates DW, Aaron S, Ai A, Cholan RA, et al. A multi-site randomized trial of a clinical decision support intervention to improve problem list completeness. J Am Med Inform Assoc. 2023;30(5):899-906. doi:10.1093/jamia/ocad020

288. Tsang JY, Planner C, Low CN, Stewart S, Morris C, Wong B, et al. Digital patient safety interventions in primary care: a systematic review and meta-analysis. BMC Med. 2026;24(1):407. doi:10.1186/s12916-026-05048-8
289. Pinevich Y, Clark KJ, Harrison AM, Pickering BW, Herasevich V. Interaction Time with Electronic Health Records: A Systematic Review. Appl Clin Inform. 2021;12(4):788-99. doi:10.1055/s-0041-1733909
290. Overhage JM, McCallie D. Physician Time Spent Using the Electronic Health Record During Outpatient Encounters: A Descriptive Study. Ann Intern Med. 2020;172(3):169-74. doi:10.7326/M18-3684
291. Holmgren AJ, Downing NL, Bates DW, Shanafelt TD, Milstein A, Sharp CD, et al. Assessment of Electronic Health Record Use Between US and Non-US Health Systems. JAMA Intern Med. 2021;181(2):251-9. doi:10.1001/jamainternmed.2020.7071
292. Arndt BG, Micek MA, Rule A, Shafer CM, Baltus JJ, Sinsky CA. More Tethered to the EHR: EHR Workload Trends Among Academic Primary Care Physicians, 2019-2023. Ann Fam Med. 2024;22(1):12-8. doi:10.1370/afm.3047
293. Moy AJ, Schwartz JM, Elias J, Imran S, Lucas E, Cato KD, et al. Time-motion examination of electronic health record utilization and clinician workflows indicate frequent task switching and documentation burden. AMIA Annu Symp Proc. 2020;2020:886-95.
294. Wu Y, Wu M, Wang C, Lin J, Liu J, Liu S. Evaluating the Prevalence of Burnout Among Health Care Professionals Related to Electronic Health Record Use: Systematic Review and Meta-Analysis. JMIR Med Inform. 2024;12:e54811. doi:10.2196/54811
295. Holmgren AJ, Hendrix N, Maisel N, Everson J, Bazemore A, Rotenstein L, et al. Electronic Health Record Usability, Satisfaction, and Burnout for Family Physicians. JAMA Netw Open. 2024;7(8):e2426956. doi:10.1001/jamanetworkopen.2024.26956
296. Rotenstein LS, Holmgren AJ, Thombley R, Sriram A, Dbouk RH, Jost M, et al. Changes in Clinician Time Expenditure and Visit Quantity With Adoption of Artificial Intelligence-Powered Scribes: A Multisite Study. JAMA. 2026;335(16):1408-17. doi:10.1001/jama.2026.2253
297. Eurostat. Individuals - internet activities (isoc_ci_ac_i) [Internet]. Luxembourg: Eurostat; 2026 Apr 17 [cited 2026 Sep 27]. doi:10.2908/ISOC_CI_AC_I. Available from: https://ec.europa.eu/eurostat/databrowser/view/isoc_ci_ac_i/default/table?lang=en
298. El-Yafouri R, Klieb L. A scoping review of electronic health records interoperability levels, expectations, approaches, and problems. Health Informatics J. 2025;31(4):14604582251385986. doi:10.1177/14604582251385986
299. Luu HS, Campbell WS, Cholan RA, Edgerton ME, Englund A, Keller A, et al. Analysis of laboratory data transmission between two healthcare institutions using a widely used point-to-point health information exchange platform: a case report. JAMIA Open. 2024;7(2):ooae032. doi:10.1093/jamiaopen/ooae032
300. HL7 International. HL7 Version 2 Product Suite [Internet]. Ann Arbor (MI): HL7 International; 2026 [cited 2026 Sep 27]. Available from: https://www.hl7.org/implement/standards/product_brief.cfm?product_id=185
301. HL7 International. HL7 Clinical Document Architecture (CDA) Release 2 [Internet]. Ann Arbor (MI): HL7 International; 2026 [cited 2026 Sep 27]. Available from: https://www.hl7.org/implement/standards/product_brief.cfm?product_id=7
302. openEHR International. What is openEHR [Internet]. openEHR International; 2026 [cited 2026 Sep 27]. Available from: https://openehr.org/about/what_is_openehr
303. Pedrera-Jiménez M, García-Barrio N, Frid S, Moner D, Boscá-Tomás D, Lozano-Rubí R, et al. Can OpenEHR, ISO 13606, and HL7 FHIR Work Together? An Agnostic Approach for the Selection and Application of Electronic Health Record Standards to the Next-Generation Health Data Spaces. J Med Internet Res. 2023;25:e48702. doi:10.2196/48702
304. Marfoglia A, Arcobelli VA, Moscato S, La Mattina AA, Mellone S, Carbonaro A. Challenges of health data standard adoption and usage: a systematic review. J Biomed Inform. 2026;178:105022. doi:10.1016/j.jbi.2026.105022
305. HL7 International. FHIR Release 5 (v5.0.0) [Internet]. Ann Arbor (MI): HL7 International; 2023 [cited 2026 Sep 27]. Available from: https://hl7.org/fhir/R5/
306. HL7 International. SMART App Launch Implementation Guide, release 2.2.0 (STU 2.2) [Internet]. Ann Arbor (MI): HL7 International; 2024 [cited 2026 Sep 27]. Available from: https://hl7.org/fhir/smart-app-launch/
307. HL7 International. FHIR publication (version) history [Internet]. Ann Arbor (MI): HL7 International; 2026 [cited 2026 Sep 27]. Available from: https://hl7.org/fhir/directory.html

308. Wesley DB, Blumenthal J, Shah S, Littlejohn RA, Pruitt Z, Dixit R, et al. A novel application of SMART on FHIR architecture for interoperable and scalable integration of patient-reported outcome data with electronic health records. J Am Med Inform Assoc. 2021;28(10):2220-5. doi:10.1093/jamia/ocab110
309. HL7 International. CDS Hooks Specification, STU 2 Release (v2.0.1) [Internet]. Ann Arbor (MI): HL7 International; 2025 [cited 2026 Sep 27]. Available from: https://cds-hooks.hl7.org/
310. Morgan KL, Kukhareva PV, Warner PB, Wilkof J, Snyder M, Horton D, et al. Using CDS Hooks to increase SMART on FHIR app utilization: a cluster-randomized trial. J Am Med Inform Assoc. 2022;29(9):1461-70. doi:10.1093/jamia/ocac085
311. ISO 27269:2025. Health informatics - International patient summary. 2nd ed. Geneva: International Organization for Standardization; 2025.
312. ISO 17523:2025. Health informatics - Requirements for electronic prescriptions. 2nd ed. Geneva: International Organization for Standardization; 2025.
313. HL7 International. International Patient Summary Implementation Guide, release 2.0.1 (STU 2) [Internet]. Ann Arbor (MI): HL7 International; 2026 [cited 2026 Sep 27]. Available from: https://hl7.org/fhir/uv/ips/
314. D'Amore JD, Vallapu ER, Samal L, Karras BT, Lober WB. From Pandemic Response to Kill the Clipboard: Patient-Controlled Sharing of Health Data Using International Patient Summary (IPS) and QR codes. Appl Clin Inform. 2026;17(3):559-67. doi:10.1055/a-2902-8515
315. Laleci Erturkmen GB, Sinaci AA, Namli T, Boonstra MJ, Lekadir K, Gkontra P, et al. An assessment of the European Patient Summary for clinical research: a case study in cardiology. Front Med (Lausanne). 2024;11:1481551. doi:10.3389/fmed.2024.1481551
316. European Commission. Electronic cross-border health services (MyHealth@EU) [Internet]. Brussels: European Commission; 2026 [cited 2026 Sep 27]. Available from: https://health.ec.europa.eu/ehealth-digital-health-and-care/electronic-cross-border-health-services_en
317. Bruthans J, Jiráková K. The Current State and Usage of European Electronic Cross-border Health Services (eHDSI). J Med Syst. 2023;47(1):21. doi:10.1007/s10916-023-01920-9
318. Jõgi R, Timonen J, Saastamoinen L, Laius O, Volmer D. Implementation of European Cross-border Electronic Prescription and Electronic Dispensing Service: Cross-sectional Survey. J Med Internet Res. 2023;25:e42453. doi:10.2196/42453
319. European Commission. European Health Data Space Regulation (EHDS) [Internet]. Brussels: European Commission; 2026 [cited 2026 Sep 27]. Available from: https://health.ec.europa.eu/ehealth-digital-health-and-care/european-health-data-space-regulation-ehds_en
320. Donia J, Marelli L. Anticipating ethical and social dimensions of the European Health Data Space: A rapid systematic review. Health Policy. 2025;162:105443. doi:10.1016/j.healthpol.2025.105443
321. Legea nr. 45/2019 pentru modificarea și completarea Legii nr. 95/2006 privind reforma în domeniul sănătății. Monitorul Oficial al României, Partea I. 2019 Mar 11;(192).
322. Casa Națională de Asigurări de Sănătate. Portal CNAS-PIAS: Platforma Integrată a Asigurărilor de Sănătate [Internet]. București: CNAS; 2026 [cited 2026 Sep 27]. Available from: https://portal.cnas.ro/cnas/
323. Casa Națională de Asigurări de Sănătate. Comunicat de presă: Finalizare proiect "Redimensionarea, standardizarea și optimizarea Platformei Informatice din Asigurările de Sănătate - PIAS" [Internet]. București: CNAS; 2026 Aug 31 [cited 2026 Sep 27]. Available from: https://cnas.ro/wp-content/uploads/2026/09/Comunicat-PNRR-01092026b-final.pdf
324. Casa Națională de Asigurări de Sănătate. Comunicat: CNAS a lansat portalul eSănătateaMea [Internet]. București: CNAS; 2026 Sep 2 [cited 2026 Sep 27]. Available from: https://cnas.ro/2026/09/02/comunicat-cnas-a-lansat-portalul-esanatateamea/
325. Casa Națională de Asigurări de Sănătate. Comunicat: Precizări privind utilizarea CEAS/CEI în sistemul asigurărilor sociale de sănătate [Internet]. București: CNAS; 2026 Aug 31 [cited 2026 Sep 27]. Available from: https://cnas.ro/2026/08/31/comunicat-precizari-privind-utilizarea-ceas-cei-in-sistemul-asigurarilor-sociale-de-sanatate/
326. Ministerul Sănătății. Componenta 7 - Transformare digitală. Investiția specifică I3.3 - Investiții în sistemele informatice și în infrastructura digitală a unităților sanitare publice [Internet]. București: Ministerul Sănătății; 2023 [cited 2026 Sep 27]. Available from: https://ms.ro/ro/de-interes/apeluri-pnrr-c12-sanatate/componenta-7-transformare-digital%C4%83-investitia-specific%C4%83-i33-investi%C8%9Bii-%C3%AEn-sistemele-informatice-%C8%99i-%C3%AEn-infrastructura-digital%C4%83-a-unit%C4%83%C8%9Bilor-sanitare-publice/
327. Spitalul Clinic Județean de Urgență "Pius Brînzeu" Timișoara. Sănătatea, siguranța și confortul pacientului prin servicii integrate digitale la SCJU Timișoara - PNRR, Componenta C7 - Transformare digitală, PNRR/C7/MS-733/I3.3 [Internet]. Timișoara: SCJUPBT; 2024 Aug 17 [cited 2026 Sep 27]. Available

from: https://www.hosptm.ro/sanatatea-siguranta-si-confortul-pacientului-prin-servicii-integrate-digitale-la-scju-timisoara-pnrr-componenta-c7-transformare-digitala-pnrr-c7-ms-733-i3-3/

328. Adams KT, Pruitt Z, Kazi S, Hettinger AZ, Howe JL, Fong A, et al. Identifying Health Information Technology Usability Issues Contributing to Medication Errors Across Medication Process Stages. J Patient Saf. 2021;17(8):e988-94. doi:10.1097/PTS.0000000000000868
329. Schwappach D, Hautz W, Krummrey G, Pfeiffer Y, Ratwani RM. EMR usability and patient safety: a national survey of physicians. NPJ Digit Med. 2025;8(1):282. doi:10.1038/s41746-025-01657-4
330. Edrees H, Amato MG, Wong A, Seger DL, Bates DW. High-priority drug-drug interaction clinical decision support overrides in a newly implemented commercial computerized provider order-entry system: Override appropriateness and adverse drug events. J Am Med Inform Assoc. 2020;27(6):893-900. doi:10.1093/jamia/ocaa034
331. Ray CE, Wilson GM, Hughes AM, Cunningham Goedken C, Liu EP, Fitzpatrick MA, et al. Alert fatigue measurement in clinical decision support: a systematic review. J Am Med Inform Assoc. 2026;33(8):1523-31. doi:10.1093/jamia/ocag064
332. Steinkamp J, Kantrowitz JJ, Airan-Javia S. Prevalence and Sources of Duplicate Information in the Electronic Medical Record. JAMA Netw Open. 2022;5(9):e2233348. doi:10.1001/jamanetworkopen.2022.33348
333. Sittig DF, Flanagan T, Sengstack P, Cholankeril RT, Ehsan S, Heidemann A, et al. Revisions to the Safety Assurance Factors for Electronic Health Record Resilience (SAFER) Guides to update national recommendations for safe use of electronic health records. J Am Med Inform Assoc. 2025;32(4):755-60. doi:10.1093/jamia/ocaf018
334. World Health Organization. Consolidated telemedicine implementation guide [Internet]. Geneva: WHO; 2022 [cited 2026 Sep 27]. Available from: https://www.who.int/publications/i/item/9789240059184
335. World Health Organization. Classification of digital interventions, services and applications in health: a shared language to describe the uses of digital technology for health. 2nd ed [Internet]. Geneva: WHO; 2023 [cited 2026 Sep 27]. Available from: https://www.who.int/publications/i/item/9789240081949
336. O'Shea AMJ, Howren MB, Mulligan K, Haraldsson B, Shahnazi A, Kaboli PJ. Quantifying the Digital Divide: Associations of Broadband Internet with Tele-mental Health Access Before and During the COVID-19 Pandemic. J Gen Intern Med. 2023;38(Suppl 3):832-40. doi:10.1007/s11606-023-08120-8
337. Fan M, Wang Q, Liu J, Su L, Wang B, Xu H, et al. Real-world evaluation of the Stemoscope electronic tele-auscultation system. Biomed Eng Online. 2022;21(1):63. doi:10.1186/s12938-022-01032-4
338. Martyin K, Meznerics FA, Bokor LA, Szabó B, Hegyi P, Kiss N, et al. Diagnostic accuracy of teledermatology for skin diseases: a systematic review and meta-analysis. Front Med (Lausanne). 2026;13:1739592. doi:10.3389/fmed.2026.1739592
339. Tong SYK, Jackson TM, Lau AYS. Virtual physical examination in teleconsultation: A scoping review. Int J Med Inform. 2024;191:105561. doi:10.1016/j.ijmedinf.2024.105561
340. Jitsi. Security and privacy in Jitsi Meet [Internet]. 2026 [cited 2026 Sep 27]. Available from: https://jitsi.org/security/
341. Tukur M, Saad G, AlShagathrh FM, Househ M, Agus M. Telehealth interventions during COVID-19 pandemic: a scoping review of applications, challenges, privacy and security issues. BMJ Health Care Inform. 2023;30(1):e100676. doi:10.1136/bmjhci-2022-100676
342. Demaerschalk BM, Pines A, Butterfield R, Haglin JM, Haddad TC, Yiannias J, et al. Assessment of Clinician Diagnostic Concordance With Video Telemedicine in the Integrated Multispecialty Practice at Mayo Clinic During the Beginning of COVID-19 Pandemic From March to June 2020. JAMA Netw Open. 2022;5(9):e2229958. doi:10.1001/jamanetworkopen.2022.29958
343. Patel SY, Mehrotra A, Huskamp HA, Uscher-Pines L, Ganguli I, Barnett ML. Variation In Telemedicine Use And Outpatient Care During The COVID-19 Pandemic In The United States. Health Aff (Millwood). 2021;40(2):349-58. doi:10.1377/hlthaff.2020.01786
344. Kirk UB, Pellin MCG, Hoffmann K, Stummer FO, Van Poel E, Vestergaard CH, et al. Trends in the use of video consultation in general practice during COVID-19: impact of practice and country characteristics based on the international, cross-sectional PRICOV-19 study. BMC Prim Care. 2025;26(1):205. doi:10.1186/s12875-025-02892-0
345. McBain RK, Schuler MS, Breslau J, Kofner A, Wang L, Cantor JH. Telehealth Availability for Mental Health Care During and After the COVID-19 Public Health Emergency. JAMA Netw Open. 2024;7(7):e2420853. doi:10.1001/jamanetworkopen.2024.20853
346. Scholte NTB, Gürgöze MT, Aydin D, Theuns DAMJ, Manintveld OC, Ronner E, et al. Telemonitoring for heart failure: a meta-analysis. Eur Heart J. 2023;44(31):2911-26. doi:10.1093/eurheartj/ehad280

347. Kuan PX, Chan WK, Fern Ying DK, Rahman MAA, Peariasamy KM, Lai NM, et al. Efficacy of telemedicine for the management of cardiovascular disease: a systematic review and meta-analysis. Lancet Digit Health. 2022;4(9):e676-91. doi:10.1016/S2589-7500(22)00124-8
348. Kaihara T, Intan-Goey V, Scherrenberg M, Falter M, Kario K, Akashi Y, et al. Automatic transmission of home blood pressure data can be effective in managing hypertension: a systematic review and meta-analysis. Eur Heart J Digit Health. 2022;3(4):638-53. doi:10.1093/ehjdh/ztac049
349. Grover S, Mishra HP, Gupta R, Gupta LK. Effect of telemonitoring and home blood pressure monitoring on blood pressure reduction in hypertensive adults: a network meta-analysis. J Hypertens. 2025;43(7):1091-8. doi:10.1097/HJH.0000000000004008
350. Hangaard S, Laursen SH, Andersen JD, Kronborg T, Vestergaard P, Hejlesen O, et al. The Effectiveness of Telemedicine Solutions for the Management of Type 2 Diabetes: A Systematic Review, Meta-Analysis, and Meta-Regression. J Diabetes Sci Technol. 2023;17(3):794-825. doi:10.1177/19322968211064633
351. Zhang J, Ji X, Xie J, Lin K, Yao M, Chi C. Effectiveness of synchronous teleconsultation for patients with type 2 diabetes mellitus: a systematic review and meta-analysis. BMJ Open Diabetes Res Care. 2023;11(1):e003180. doi:10.1136/bmjdrc-2022-003180
352. Janjua S, Carter D, Threapleton CJ, Prigmore S, Disler RT. Telehealth interventions: remote monitoring and consultations for people with chronic obstructive pulmonary disease (COPD). Cochrane Database Syst Rev. 2021;7(7):CD013196. doi:10.1002/14651858.CD013196.pub2
353. Cox NS, Dal Corso S, Hansen H, McDonald CF, Hill CJ, Zanaboni P, et al. Telerehabilitation for chronic respiratory disease. Cochrane Database Syst Rev. 2021;1(1):CD013040. doi:10.1002/14651858.CD013040.pub2
354. Hagi K, Kurokawa S, Takamiya A, Fujikawa M, Kinoshita S, Iizuka M, et al. Telepsychiatry versus face-to-face treatment: systematic review and meta-analysis of randomised controlled trials. Br J Psychiatry. 2023;223(3):407-14. doi:10.1192/bjp.2023.86
355. Mohamed A, Elsherif S, Legere B, Fatima N, Shuaib A, Saqqur M. Is telestroke more effective than conventional treatment for acute ischemic stroke? A systematic review and meta-analysis of patient outcomes and thrombolysis rates. Int J Stroke. 2024;19(3):280-92. doi:10.1177/17474930231206066
356. Behrens JR, Kaffes M, Aigner A, Herm J, Erdur H, Siepmann T, et al. Teleneurology vs On-Site Neurology Consultation for Postadmission Hospital Care of Stroke. JAMA Neurol. 2026;83(6):554-62. doi:10.1001/jamaneurol.2026.0615
357. Storjohann S, Kirsch M, Rosenberg B, Rosenberg C, Lange S, Syperek A, et al. The Accuracy of On-Call CT Reporting in Teleradiology Networks in Comparison to In-House Reporting. Healthcare (Basel). 2021;9(4):405. doi:10.3390/healthcare9040405
358. Wittman SR, Hoberman A, Mehrotra A, Sabik LM, Yabes JG, Ray KN. Antibiotic Receipt for Pediatric Telemedicine Visits With Primary Care vs Direct-to-Consumer Vendors. JAMA Netw Open. 2024;7(3):e242359. doi:10.1001/jamanetworkopen.2024.2359
359. Lam K, Lu AD, Shi Y, Covinsky KE. Assessing Telemedicine Unreadiness Among Older Adults in the United States During the COVID-19 Pandemic. JAMA Intern Med. 2020;180(10):1389-91. doi:10.1001/jamainternmed.2020.2671
360. Solimini R, Busardò FP, Gibelli F, Sirignano A, Ricci G. Ethical and Legal Challenges of Telemedicine in the Era of the COVID-19 Pandemic. Medicina (Kaunas). 2021;57(12):1314. doi:10.3390/medicina57121314
361. Elberg JT, Adashi EY. Fraud, Access, and the Future of Telemedicine. J Am Board Fam Med. 2023;36(5):864-6. doi:10.3122/jabfm.2023.230044R1
362. Harris Y, Goldwater JC. Lack of evidence for telehealth fraud. J Telemed Telecare. 2025;31(2):301-5. doi:10.1177/1357633X231177739
363. Ordonanța de urgență a Guvernului nr. 196/2020 pentru modificarea și completarea Legii nr. 95/2006 privind reforma în domeniul sănătății. Monitorul Oficial al României, Partea I. 2020 Nov 19;(1108).
364. Hotărârea Guvernului nr. 1.133/2022 privind aprobarea Normelor metodologice de punere în aplicare a prevederilor Ordonanței de urgență a Guvernului nr. 196/2020 pentru modificarea și completarea Legii nr. 95/2006 privind reforma în domeniul sănătății. Monitorul Oficial al României, Partea I. 2022 Sep 15;(909).
365. Hotărârea Guvernului nr. 521/2023 pentru aprobarea pachetelor de servicii și a Contractului-cadru care reglementează condițiile acordării asistenței medicale, a medicamentelor și a dispozitivelor medicale, în cadrul sistemului de asigurări sociale de sănătate. Monitorul Oficial al României, Partea I. 2023 May 26;(466).
366. Ordinul ministrului sănătății și al președintelui Casei Naționale de Asigurări de Sănătate nr. 1.857/441/2023 privind aprobarea Normelor metodologice de aplicare a Hotărârii Guvernului nr.

521/2023. Monitorul Oficial al României, Partea I. 2023 May 31;(484). Forma consolidată publicată de CNAS, august 2025 [cited 2026 Sep 27]. Available from: https://cnas.ro/wp-content/uploads/2025/08/Norme-1-si-2.pdf

367. Casa Națională de Asigurări de Sănătate. Comunicat: Consultațiile medicale și rețetele electronice pot fi acordate în continuare la distanță [Internet]. București: CNAS; 2023 Jul 6 [cited 2026 Sep 27]. Available from: https://cnas.ro/2023/07/06/comunicat-consultatiile-medicale-si-retetele-electronice-pot-fi-acordate-in-continuare-la-distanta/
368. Ministerul Sănătății. Apeluri PNRR: Componenta 12 Sănătate și Componenta 7 Transformare digitală [Internet]. București: Ministerul Sănătății; 2026 [cited 2026 Sep 27]. Available from: https://www.ms.ro/ro/de-interes/apeluri-pnrr-c12-sanatate/
369. Directive 2011/24/EU of the European Parliament and of the Council of 9 March 2011 on the application of patients' rights in cross-border healthcare. Official Journal of the European Union. 2011 Apr 4;L 88:45-65.
370. Caffery LJ, Taylor ML, Abbott LM, Janda M, Guitera P, Mar V, et al. Practice Guidelines for Teledermatology in Australia: 2025 Update. Australas J Dermatol. 2025;66(8):e553-64. doi:10.1111/ajd.14599
371. Edwards PJ, Bennett-Britton I, Ridd MJ, Booker M, Barnes RK. Factors affecting the documentation of spoken safety-netting advice in routine GP consultations: a cross-sectional study. Br J Gen Pract. 2021;71(712):e869-76. doi:10.3399/BJGP.2021.0195
372. Medical Device Coordination Group. MDCG 2019-11 rev.1. Qualification and classification of software - Regulation (EU) 2017/745 and Regulation (EU) 2017/746 [Internet]. Brussels: European Commission; 2025 Jun [cited 2026 Sep 27]. Available from: https://health.ec.europa.eu/document/download/b45335c5-1679-4c71-a91c-fc7a4d37f12b_en?filename=mdcg_2019_11_en.pdf
373. Vasiloglou MF, Kotzakioulafi E, Mainardi F, Ahmed M, Staiano AE, Dimidi E, et al. Effectiveness of app-based interventions for glucose management, cardiometabolic and mental health, and related risk factors: an umbrella review with meta-analysis. Lancet Digit Health. 2026;8(8):101009. doi:10.1016/j.landig.2026.101009
374. Sun S, Simonsson O, McGarvey S, Torous J, Goldberg SB. Mobile phone interventions to improve health outcomes among patients with chronic diseases: an umbrella review and evidence synthesis from 34 meta-analyses. Lancet Digit Health. 2024;6(11):e857-70. doi:10.1016/S2589-7500(24)00119-5
375. Moschonis G, Siopis G, Jung J, Eweka E, Willems R, Kwasnicka D, et al. Effectiveness, reach, uptake, and feasibility of digital health interventions for adults with type 2 diabetes: a systematic review and meta-analysis of randomised controlled trials. Lancet Digit Health. 2023;5(3):e125-43. doi:10.1016/S2589-7500(22)00233-3
376. Yap HJ, Lim JJJ, Tan SD, Ang CS. Effectiveness of digital health interventions on adherence and control of hypertension: a systematic review and meta-analysis. J Hypertens. 2024;42(9):1490-504. doi:10.1097/HJH.0000000000003793
377. Chu S, Feng L, Jing H, Zhang D, Li J, Meng X, et al. Efficacy of smartphone apps used alone or with traditional interventions for smoking cessation: a systematic review and meta-analysis. BMJ Evid Based Med. 2026;31(4):242-53. doi:10.1136/bmjebm-2025-113971
378. Linardon J, Torous J, Firth J, Cuijpers P, Messer M, Fuller-Tyszkiewicz M. Current evidence on the efficacy of mental health smartphone apps for symptoms of depression and anxiety. A meta-analysis of 176 randomized controlled trials. World Psychiatry. 2024;23(1):139-49. doi:10.1002/wps.21183
379. Terhorst Y, Philippi P, Sander LB, Schultchen D, Paganini S, Bardus M, et al. Validation of the Mobile Application Rating Scale (MARS). PLoS One. 2020;15(11):e0241480. doi:10.1371/journal.pone.0241480
380. Zych MM, Bond R, Mulvenna M, Martinez Carracedo J, Bai L, Leigh S. Quality Assessment of Digital Health Apps: Umbrella Review. J Med Internet Res. 2024;26:e58616. doi:10.2196/58616
381. Jakob R, Harperink S, Rudolf AM, Fleisch E, Haug S, Mair JL, et al. Factors Influencing Adherence to mHealth Apps for Prevention or Management of Noncommunicable Diseases: Systematic Review. J Med Internet Res. 2022;24(5):e35371. doi:10.2196/35371
382. Liu C, Torous J, Fuller-Tyszkiewicz M, Messer M, Anderson C, Soliman OM, et al. Uptake, Adherence, and Attrition in Clinical Trials of Depression and Anxiety Apps: A Systematic Review and Meta-Analysis. JAMA Psychiatry. 2026;83(1):43-50. doi:10.1001/jamapsychiatry.2025.3439
383. Bluetooth Special Interest Group. Generic Health Sensor Profile, version 1.0 [Internet]. Kirkland (WA): Bluetooth SIG; 2023 Jun 13 [cited 2026 Sep 27]. Available from: https://www.bluetooth.com/specifications/specs/generic-health-sensor-profile/

384. HL7 International. Personal Health Device FHIR Implementation Guide, version 2.0.0 (STU 2) [Internet]. Ann Arbor (MI): HL7 International; 2026 Apr 1 [cited 2026 Sep 27]. Available from: https://hl7.org/fhir/uv/phd/
385. Abedian S, Yesakov E, Ostrovskiy S, Hussein R. Integrating Garmin Wearable Data into FHIR-Based Health Systems for Improved Interoperability. Stud Health Technol Inform. 2025;332:185-9. doi:10.3233/SHTI251523
386. Doherty C, Baldwin M, Keogh A, Caulfield B, Argent R. Keeping Pace with Wearables: A Living Umbrella Review of Systematic Reviews Evaluating the Accuracy of Consumer Wearable Technologies in Health Measurement. Sports Med. 2024;54(11):2907-26. doi:10.1007/s40279-024-02077-2
387. Islam SMS, Chow CK, Daryabeygikhotbehsara R, Subedi N, Rawstorn J, Tegegne T, et al. Wearable cuffless blood pressure monitoring devices: a systematic review and meta-analysis. Eur Heart J Digit Health. 2022;3(2):323-37. doi:10.1093/ehjdh/ztac021
388. Rizos EC, Markozannes G, Charitakis N, Filis P, Stoimeni AE, Nørgaard K, et al. Continuous Glucose Monitoring in Type 1 Diabetes, Type 2 Diabetes, and Diabetes During Pregnancy: A Systematic Review with Meta-Analysis of Randomized Controlled Trials. Diabetes Technol Ther. 2025;27(7):537-52. doi:10.1089/dia.2024.0599
389. Bundesinstitut für Arzneimittel und Medizinprodukte (BfArM). Häufig gestellte Fragen (FAQ): Änderungen durch das DigiG [Internet]. Bonn: BfArM; 2026 [cited 2026 Sep 27]. Available from: https://www.bfarm.de/DE/Medizinprodukte/_FAQ/DigiG/_artikel.html
390. Kolominsky-Rabas PL, Tauscher M, Gerlach R, Perleth M, Dietzel N. [How robust are studies of currently permanently included digital health applications (DiGA)? Methodological quality of studies demonstrating positive health care effects of DiGA]. Z Evid Fortbild Qual Gesundhwes. 2022;175:1-16. doi:10.1016/j.zefq.2022.09.008
391. Federal Institute for Drugs and Medical Devices (BfArM). Digital Health Applications (DiGA): interesting facts [Internet]. Bonn: BfArM; 2026 [cited 2026 Sep 27]. Available from: https://www.bfarm.de/EN/Medical-devices/Tasks/DiGA-and-DiPA/Digital-Health-Applications/Interesting-facts/_artikel.html
392. Goeldner M, Gehder S. Digital Health Applications (DiGAs) on a Fast Track: Insights From a Data-Driven Analysis of Prescribable Digital Therapeutics in Germany From 2020 to Mid-2024. J Med Internet Res. 2024;26:e59013. doi:10.2196/59013
393. Kawu AA, Hederman L, O'Sullivan D. Exchanging Context and Provenance Based Standardized Patient Generated Health Data with an Electronic Health Record. AMIA Annu Symp Proc. 2024;2024:575-84.
394. Ungurean I, Gherman OI, Lavric A, Dimian M. A privacy-centric microservice framework for secure FHIR-based integration of heterogeneous medical data. Int J Med Inform. 2026;216:106477. doi:10.1016/j.ijmedinf.2026.106477
395. Liu M, Okuhara T, Chang X, Shirabe R, Nishiie Y, Okada H, et al. Performance of ChatGPT Across Different Versions in Medical Licensing Examinations Worldwide: Systematic Review and Meta-Analysis. J Med Internet Res. 2024;26:e60807. doi:10.2196/60807
396. Singhal K, Tu T, Gottweis J, Sayres R, Wulczyn E, Amin M, et al. Toward expert-level medical question answering with large language models. Nat Med. 2025;31(3):943-50. doi:10.1038/s41591-024-03423-7
397. Gong EJ, Bang CS, Lee JJ, Baik GH. Knowledge-Practice Performance Gap in Clinical Large Language Models: Systematic Review of 39 Benchmarks. J Med Internet Res. 2025;27:e84120. doi:10.2196/84120
398. Ayers JW, Poliak A, Dredze M, Leas EC, Zhu Z, Kelley JB, et al. Comparing Physician and Artificial Intelligence Chatbot Responses to Patient Questions Posted to a Public Social Media Forum. JAMA Intern Med. 2023;183(6):589-96. doi:10.1001/jamainternmed.2023.1838
399. Goh E, Gallo R, Hom J, Strong E, Weng Y, Kerman H, et al. Large Language Model Influence on Diagnostic Reasoning: A Randomized Clinical Trial. JAMA Netw Open. 2024;7(10):e2440969. doi:10.1001/jamanetworkopen.2024.40969
400. Goh E, Gallo RJ, Strong E, Weng Y, Kerman H, Freed JA, et al. GPT-4 assistance for improvement of physician performance on patient care tasks: a randomized controlled trial. Nat Med. 2025;31(4):1233-8. doi:10.1038/s41591-024-03456-y
401. Tu T, Schaekermann M, Palepu A, Saab K, Freyberg J, Tanno R, et al. Towards conversational diagnostic artificial intelligence. Nature. 2025;642(8067):442-50. doi:10.1038/s41586-025-08866-7
402. Hager P, Jungmann F, Holland R, Bhagat K, Hubrecht I, Knauer M, et al. Evaluation and mitigation of the limitations of large language models in clinical decision-making. Nat Med. 2024;30(9):2613-22. doi:10.1038/s41591-024-03097-1

403. Bean AM, Payne RE, Parsons G, Kirk HR, Ciro J, Mosquera-Gómez R, et al. Reliability of LLMs as medical assistants for the general public: a randomized preregistered study. Nat Med. 2026;32(2):609-15. doi:10.1038/s41591-025-04074-y
404. Ramaswamy A, Tyagi A, Hugo H, Jiang J, Jayaraman P, Jangda M, et al. ChatGPT Health performance in a structured test of triage recommendations. Nat Med. 2026;32(5):1671-5. doi:10.1038/s41591-026-04297-7
405. Tiller NB, Marcon AR, Zenone M, Kidd KE, Jeukendrup AE, Master Z, et al. Generative artificial intelligence-driven chatbots and medical misinformation: an accuracy, referencing and readability audit. BMJ Open. 2026;16(4):e112695. doi:10.1136/bmjopen-2025-112695
406. Asgari E, Montaña-Brown N, Dubois M, Khalil S, Balloch J, Yeung JA, et al. A framework to assess clinical safety and hallucination rates of LLMs for medical text summarisation. NPJ Digit Med. 2025;8(1):274. doi:10.1038/s41746-025-01670-7
407. Chen S, Gao M, Sasse K, Hartvigsen T, Anthony B, Fan L, et al. When helpfulness backfires: LLMs and the risk of false medical information due to sycophantic behavior. NPJ Digit Med. 2025;8(1):605. doi:10.1038/s41746-025-02008-z
408. Walters WH, Wilder EI. Fabrication and errors in the bibliographic citations generated by ChatGPT. Sci Rep. 2023;13(1):14045. doi:10.1038/s41598-023-41032-5
409. Haltaufderheide J, Ranisch R. The ethics of ChatGPT in medicine and healthcare: a systematic review on Large Language Models (LLMs). NPJ Digit Med. 2024;7(1):183. doi:10.1038/s41746-024-01157-x
410. World Health Organization. WHO releases AI ethics and governance guidance for large multi-modal models [Internet]. Geneva: WHO; 2024 Jan 18 [cited 2026 Sep 27]. Available from: https://www.who.int/news/item/18-01-2024-who-releases-ai-ethics-and-governance-guidance-for-large-multi-modal-models
411. World Health Organization. Ethics and governance of artificial intelligence for health: guidance on large multi-modal models [Internet]. Geneva: WHO; 2025 [cited 2026 Sep 27]. Available from: https://www.who.int/publications/i/item/9789240084759
412. Agenția Națională a Medicamentului și a Dispozitivelor Medicale din România. Dispozitive medicale [Internet]. București: ANMDMR; 2026 [cited 2026 Sep 27]. Available from: https://www.anm.ro/dispozitive-medicale/
413. Winter A, Ammenwerth E, Haux R, Marschollek M, Steiner B, Jahn F. Health Information Systems: Technological and Management Perspectives. 3rd ed. Cham: Springer; 2023. Chapter 2, Basic Concepts and Terms; p. 13-49. doi:10.1007/978-3-031-12310-8_2
414. Winter A, Ammenwerth E, Haux R, Marschollek M, Steiner B, Jahn F. Health Information Systems: Technological and Management Perspectives. 3rd ed. Cham: Springer; 2023. Chapter 4, Management Perspective: Scopes and Tasks of Managing Health Information Systems; p. 153-88. doi:10.1007/978-3-031-12310-8_4
415. Krause T, Jolkver E, Mc Kevitt P, Kramer M, Hemmje M. A Systematic Approach to Diagnostic Laboratory Software Requirements Analysis. Bioengineering (Basel). 2022;9(4):144. doi:10.3390/bioengineering9040144
416. Finnegan H, Mountford N. 25 Years of Electronic Health Record Implementation Processes: Scoping Review. J Med Internet Res. 2025;27:e60077. doi:10.2196/60077
417. Zharima C, Mhlanga S, Abdulla S, Goudge J, Griffiths F. What engagement strategies are useful in facilitating the implementation of electronic health records in health care settings? A rapid review of qualitative evidence synthesis using the normalization process theory. Digit Health. 2024;10:20552076241291286. doi:10.1177/20552076241291286
418. Casa Națională de Asigurări de Sănătate. Comunicat: Aplicațiile informatice ale furnizorilor trebuie actualizate pentru utilizarea noilor funcționalități disponibile prin eSănătateaMea [Internet]. București: CNAS; 2026 Aug 29 [cited 2026 Sep 27]. Available from: https://cnas.ro/2026/08/29/comunicat-aplicatiile-informatice-ale-furnizorilor-trebuie-actualizate-pentru-utilizarea-noilor-functionalitati-disponibile-prin-esanatateamea/
419. ISO/IEC 27001:2022. Information security, cybersecurity and privacy protection - Information security management systems - Requirements. 3rd ed. Geneva: International Organization for Standardization; 2022 (amended by ISO/IEC 27001:2022/Amd 1:2024).
420. ISO/IEC 27005:2022. Information security, cybersecurity and privacy protection - Guidance on managing information security risks. 4th ed. Geneva: International Organization for Standardization; 2022.
421. Wasserman L, Wasserman Y. Hospital cybersecurity risks and gaps: Review (for the non-cyber professional). Front Digit Health. 2022;4:862221. doi:10.3389/fdgth.2022.862221

422. Rule A, Chiang MF, Hribar MR. Using electronic health record audit logs to study clinical activity: a systematic review of aims, measures, and methods. J Am Med Inform Assoc. 2020;27(3):480-90. doi:10.1093/jamia/ocz196
423. Sittig DF, Wright A. A guide to mitigating audit log-related risk in medical professional liability cases. J Healthc Risk Manag. 2023;43(2):37-47. doi:10.1002/jhrm.21553
424. Larsen EP, Rao AH, Sasangohar F. Understanding the scope of downtime threats: A scoping review of downtime-focused literature and news media. Health Informatics J. 2020;26(4):2660-72. doi:10.1177/1460458220918539
425. Tully JL, Rao S, Straw I, Gabriel RA, Longhurst CA, Savage S, et al. Patient Care Technology Disruptions Associated With the CrowdStrike Outage. JAMA Netw Open. 2025;8(7):e2530226. doi:10.1001/jamanetworkopen.2025.30226
426. National Institute of Standards and Technology. Advanced Encryption Standard (AES). FIPS PUB 197, updated 2023 May 9. Gaithersburg (MD): NIST; 2023. doi:10.6028/NIST.FIPS.197-upd1
427. National Institute of Standards and Technology. Digital Signature Standard (DSS). FIPS PUB 186-5. Gaithersburg (MD): NIST; 2023 Feb. doi:10.6028/NIST.FIPS.186-5
428. ISO/DIS 17090-4:2025. Health informatics - Public key infrastructure - Part 4: Digital signatures for healthcare documents. Draft International Standard (voting 2025 Jan 27 - 2025 Apr 21). Geneva: International Organization for Standardization; 2025.
429. ISO 17090-1:2021. Health informatics - Public key infrastructure - Part 1: Overview of digital certificate services. Geneva: International Organization for Standardization; 2021.
430. ISO 17090-4:2026. Health informatics - Public key infrastructure - Part 4: Digital signatures for healthcare documents. 3rd ed. Geneva: International Organization for Standardization; 2026.
431. Rescorla E. The Transport Layer Security (TLS) Protocol Version 1.3. RFC 9846 [Internet]. Internet Engineering Task Force; 2026 Jul [cited 2026 Sep 27]. doi:10.17487/RFC9846
432. National Institute of Standards and Technology. Module-Lattice-Based Key-Encapsulation Mechanism Standard. FIPS PUB 203. Gaithersburg (MD): NIST; 2024 Aug. doi:10.6028/NIST.FIPS.203
433. Regulation (EU) No 910/2014 of the European Parliament and of the Council of 23 July 2014 on electronic identification and trust services for electronic transactions in the internal market and repealing Directive 1999/93/EC. Official Journal of the European Union. 2014 Aug 28;L 257:73-114.
434. Regulation (EU) 2024/1183 of the European Parliament and of the Council of 11 April 2024 amending Regulation (EU) No 910/2014 as regards establishing the European Digital Identity Framework. Official Journal of the European Union. 2024 Apr 30;L 2024/1183.
435. Ordinul ministrului sănătății și al președintelui CNAS nr. 674/252/2012 privind aprobarea formularului de prescripție medicală electronică pentru medicamente cu și fără contribuție personală în tratamentul ambulatoriu și a Normelor metodologice privind utilizarea și modul de completare a acestuia. Monitorul Oficial al României, Partea I. 2012 Jul 2;(439). Forma consolidată la 29.07.2026.
436. Legea nr. 214/2024 privind utilizarea semnăturii electronice, a mărcii temporale și prestarea serviciilor de încredere bazate pe acestea. Monitorul Oficial al României, Partea I. 2024 Jul 8;(647).
437. European Commission. The legal and technical road to EU Digital Identity Wallets [Internet]. Brussels: European Commission; 2026 [cited 2026 Sep 27]. Available from: https://ec.europa.eu/digital-building-blocks/sites/spaces/EUDIGITALIDENTITYWALLET/pages/915931903/The+legal+and+technical+road+to+EU+Digital+Identity+Wallets
438. International Organization for Standardization. ISO Open Data: ISO deliverables metadata (iso_deliverables_metadata) [Internet]. Geneva: ISO; 2026 [updated 2026 Sep 27; cited 2026 Sep 27]. Licensed under ODC-By 1.0. Available from: https://www.iso.org/open-data.html
439. International Organization for Standardization. ISO/TC 215 Health informatics [Internet]. Geneva: ISO; 2026 [cited 2026 Sep 27]. Available from: https://www.iso.org/committee/54960.html
440. European Committee for Standardization. CEN/TC 251 Health informatics: Business Plan [Internet]. Brussels: CEN; 2022 Dec 13 [cited 2026 Sep 27]. Available from: https://standards.cencenelec.eu/BPCEN/6232.pdf
441. Asociația de Standardizare din România. Buletinul Standardizării, decembrie 2025 [Internet]. București: ASRO; 2025 [cited 2026 Sep 27]. Available from: https://www.asro.ro/wp-content/uploads/2025/12/BS-12-2025.pdf
442. Nifakos S, Chandramouli K, Nikolaou CK, Papachristou P, Koch S, Panaousis E, et al. Influence of Human Factors on Cyber Security within Healthcare Organisations: A Systematic Review. Sensors (Basel). 2021;21(15):5119. doi:10.3390/s21155119

443. Rizzoni F, Magalini S, Casaroli A, Mari P, Dixon M, Coventry L. Phishing simulation exercise in a large hospital: A case study. Digit Health. 2022;8:20552076221081716. doi:10.1177/20552076221081716
444. Yeng PK, Fauzi MA, Vestad A, Yang B, De Moor K, Jacobsen C, et al. Cognitive Dissonance-Based Priming Intervention: Randomized Encouragement With in-the-Wild Phishing Simulation Attack in Health Care. J Med Internet Res. 2026;28:e68051. doi:10.2196/68051
445. Temoshok D, Fenton JL, Choong YY, Lefkovitz N, Regenscheid A, Galluzzo R, et al. Digital Identity Guidelines: Authentication and Lifecycle Management. NIST Special Publication 800-63B-4. Gaithersburg (MD): National Institute of Standards and Technology; 2025. doi:10.6028/NIST.SP.800-63B-4
446. European Union Agency for Cybersecurity (ENISA). Cyber Hygiene in the Health Sector [Internet]. Athens: ENISA; 2025 Sep [cited 2026 Sep 27]. doi:10.2824/3818935
447. Directoratul Național de Securitate Cibernetică. PNRISC - Platforma națională pentru raportarea incidentelor de securitate cibernetică [Internet]. București: DNSC; 2026 [cited 2026 Sep 27]. Available from: https://pnrisc.dnsc.ro/
448. Winter A, Ammenwerth E, Haux R, Marschollek M, Steiner B, Jahn F. Health Information Systems: Technological and Management Perspectives. 3rd ed. Cham: Springer; 2023. doi:10.1007/978-3-031-12310-8
449. Backes J, Geissler A, Subelack J, Ehlig D. Does System Choice Matter? Influence of Hospital Information Systems on Digital Maturity. J Med Syst. 2026;50(1):124. doi:10.1007/s10916-026-02450-w
450. Devin J, Cleary BJ, Cullinan S. The impact of health information technology on prescribing errors in hospitals: a systematic review and behaviour change technique analysis. Syst Rev. 2020;9(1):275. doi:10.1186/s13643-020-01510-7
451. Srinivasamurthy SK, Ashokkumar R, Kodidela S, Howard SC, Samer CF, Chakradhara Rao US. Impact of computerised physician order entry (CPOE) on the incidence of chemotherapy-related medication errors: a systematic review. Eur J Clin Pharmacol. 2021;77(8):1123-31. doi:10.1007/s00228-021-03099-9
452. Taheri Moghadam S, Sadoughi F, Velayati F, Ehsanzadeh SJ, Poursharif S. The effects of clinical decision support system for prescribing medication on patient outcomes and physician practice performance: a systematic review and meta-analysis. BMC Med Inform Decis Mak. 2021;21(1):98. doi:10.1186/s12911-020-01376-8
453. Kalvelage C, Rademacher S, Dohmen S, Marx G, Benstoem C. Decision-Making Authority During Tele-ICU Care Reduces Mortality and Length of Stay-A Systematic Review and Meta-Analysis. Crit Care Med. 2021;49(7):1169-81. doi:10.1097/CCM.0000000000004943
454. Felisberto M, Lima GDS, Celuppi IC, Fantonelli MDS, Zanotto WL, Dias de Oliveira JM, et al. Override rate of drug-drug interaction alerts in clinical decision support systems: A brief systematic review and meta-analysis. Health Informatics J. 2024;30(2):14604582241263242. doi:10.1177/14604582241263242
455. Van Dort BA, Zheng WY, Sundar V, Baysari MT. Optimizing clinical decision support alerts in electronic medical records: a systematic review of reported strategies adopted by hospitals. J Am Med Inform Assoc. 2021;28(1):177-83. doi:10.1093/jamia/ocaa279
456. Chromik J, Klopfenstein SAI, Pfitzner B, Sinno ZC, Arnrich B, Balzer F, et al. Computational approaches to alleviate alarm fatigue in intensive care medicine: A systematic literature review. Front Digit Health. 2022;4:843747. doi:10.3389/fdgth.2022.843747
457. Zhu L, Wei S, An Y, Hu W, Xie X. Severity and Associated Factors of Alarm Fatigue Among Intensive Care Unit Nurses: A Systematic Review and Meta-Analysis. Nurs Crit Care. 2026;31(5):e70666. doi:10.1111/nicc.70666
458. Ministerul Sănătății. Sistem informatic pentru evidența clinică a secțiilor A.T.I. (S.I.E.C.-A.T.I.), cod SMIS 131065 [Internet]. București: Ministerul Sănătății; 2021 [cited 2026 Sep 27]. Available from: https://ms.ro/ro/unitatea-de-implementare-si-coordonare-programe/siec-ati/
459. Toma D, Ghenciu LA, Bedreag OH, Băloi A, Gizea CA, Rițiu SA, et al. Digital Transformation in Critical Care: Implications for Quality of Care, Infection Control, and Clinical Outcomes. J Clin Med. 2025;14(24):8964. doi:10.3390/jcm14248964
460. Vernic C, Ștefănigă SA, Apostol AS, Muntean C. The impact of digitization of intensive care services on patient safety in the Pius Brânzeu Emergency County Hospital in Timișoara [abstract]. Appl Med Inform. 2024;46(Suppl 1):S20. Available from: https://ami.info.umfcluj.ro/index.php/AMI/article/view/1046
461. Vernic C, Bedreag O, Suliţanu DC, Petre I, Muntean LI, Ursoniu S. Digital Hesitancy Among Nurses and Physicians in Anesthesia and Intensive Care: A Cross-Sectional Study from Romania. Nurs Rep. 2026;16(8):290. doi:10.3390/nursrep16080290

462. Lichtner G, Schiefenhövel F, Gashi B, Martin I, Jurth C, Vasiljewa L, et al. Multicenter Evaluation of an Interoperable System for Automated Guideline Adherence Monitoring in ICUs. Crit Care Med. 2026;54(1):119-28. doi:10.1097/CCM.0000000000006961
463. de Andrade JBC, de Medeiros Cavalcante MAO, Lopes TLM, Treigher JMS, Balsells MD, Vasconcelos JL, et al. Discovery of data quality issues in electronic health records: profound consequences for critical care medicine applications - a systematized review. Crit Care. 2026;30(1):19. doi:10.1186/s13054-025-05677-0
464. Vernic C, Tamas TP, Petre I, Ursoniu S. Transforming critical care: the digital revolution's impact on intensive care units. Front Digit Health. 2025;7:1664382. doi:10.3389/fdgth.2025.1664382
465. Institutul Național de Management al Serviciilor de Sănătate. DRG.ro: clasificare DRG, grouper RO DRG v1, noutăți [Internet]. București: INMSS; 2026 [cited 2026 Sep 27]. Available from: https://drg.inmss.ro/
466. Casa Națională de Asigurări de Sănătate. Comunicat: Platforma PIAS va fi înlocuită cu un sistem informatic modern, sigur și accesibil [Internet]. București: CNAS; 2025 Oct 15 [cited 2026 Sep 27]. Available from: https://cnas.ro/2025/10/15/comunicat-platforma-pias-va-fi-inlocuita-cu-un-sistem-informatic-modern-sigur-si-accesibil/
467. Ministerul Sănătății. Componenta 7 - Transformare digitală. Investiția specifică I3.2 - Digitalizarea instituțiilor cu atribuții în domeniul sanitar aflate în subordinea MS [Internet]. București: Ministerul Sănătății; 2023 [cited 2026 Sep 27]. Available from: https://www.ms.ro/ro/de-interes/apeluri-pnrr-c12-sanatate/componenta-7-transformare-digital%C4%83-investitia-specific%C4%83-i32-digitalizarea-institu%C8%9Biilor-cu-atribu%C8%9Bii-%C3%AEn-domeniul-sanitar-aflate-%C3%AEn-subordinea-ms/
468. Regulation (EU) 2021/241 of the European Parliament and of the Council of 12 February 2021 establishing the Recovery and Resilience Facility. Official Journal of the European Union. 2021 Feb 18;L 57:17.
469. Healthcare Information and Management Systems Society (HIMSS). Electronic Medical Record Adoption Model (EMRAM) [Internet]. Chicago: HIMSS; 2026 [cited 2026 Sep 27]. Available from: https://www.himss.org/maturity-models/emram/
470. Snowdon A, Hussein A, Danforth M, Wright A, Oakes R. Digital Maturity as a Predictor of Quality and Safety Outcomes in US Hospitals: Cross-Sectional Observational Study. J Med Internet Res. 2024;26:e56316. doi:10.2196/56316
471. Duncan R, Eden R, Woods L, Wong I, Sullivan C. Synthesizing Dimensions of Digital Maturity in Hospitals: Systematic Review. J Med Internet Res. 2022;24(3):e32994. doi:10.2196/32994
472. Vogel J, Haering A, Kuklinski D, Geissler A. Assessing the Relationship between Hospital Process Digitalization and Hospital Quality - Evidence from Germany. J Med Syst. 2024;48(1):85. doi:10.1007/s10916-024-02101-y
473. Kücking F, Hering A, Strate J, Naumann L, Liebe JD, Hübner UH. Observing the Long Struggle for Digital Transformation in German Hospitals: Results from a Panel Study 2017 to 2025. Stud Health Technol Inform. 2026;340:87-92. doi:10.3233/SHTI260989
474. Eurostat. Hospital discharges and length of stay statistics [Internet]. Luxembourg: Eurostat; 2025 [cited 2026 Sep 27]. Available from: https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Hospital_discharges_and_length_of_stay_statistics
475. Eurostat. Healthcare resource statistics - beds [Internet]. Luxembourg: Eurostat; 2026 [cited 2026 Sep 27]. Available from: https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Healthcare_resource_statistics_-_beds
476. DeRienzo C, Levine D, Jain S, Harris A. Postpandemic Recovery of Case Mix Index and Risk-Adjusted Mortality in US Hospitals. JAMA Netw Open. 2025;8(11):e2543398. doi:10.1001/jamanetworkopen.2025.43398
477. Schipmann S, Suero Molina E, Windheuser J, Doods J, Schwake M, Wilbers E, et al. The 30-day readmission rate in neurosurgery-a useful indicator for quality assessment? Acta Neurochir (Wien). 2020;162(11):2659-69. doi:10.1007/s00701-020-04382-x
478. Rabiei R, Almasi S. Requirements and challenges of hospital dashboards: a systematic literature review. BMC Med Inform Decis Mak. 2022;22(1):287. doi:10.1186/s12911-022-02037-8
479. Coiera E, Chan A, Brooke-Cowden K, Rahimi-Ardabili H, Halim N, Tufanaru C. Clinical and economic impact of digital dashboards on hospital inpatient care: a systematic review. JAMIA Open. 2025;8(4):ooaf078. doi:10.1093/jamiaopen/ooaf078
480. Keogh RJ, Harvey H, Brady C, Hassett E, Costelloe SJ, O'Sullivan MJ, et al. Dealing with digital paralysis: Surviving a cyberattack in a National Cancer center. J Cancer Policy. 2024;39:100466. doi:10.1016/j.jcpo.2023.100466

481. Lyon R, Jones A, Burke R, Baysari MT. What Goes Up, Must Come Down: A State-of-the-Art Electronic Health Record Downtime and Uptime Procedure in a Metropolitan Health Setting. Appl Clin Inform. 2023;14(3):513-20. doi:10.1055/s-0043-1768995
482. Aliferis C, Simon G. Artificial Intelligence (AI) and Machine Learning (ML) for Healthcare and Health Sciences: The Need for Best Practices Enabling Trust in AI and ML. In: Simon GJ, Aliferis C, editors. Artificial Intelligence and Machine Learning in Health Care and Medical Sciences: Best Practices and Pitfalls. Cham: Springer; 2024. p. 1-31. doi:10.1007/978-3-031-39355-6_1
483. Aliferis C, Simon G. Foundations and Properties of AI/ML Systems. In: Simon GJ, Aliferis C, editors. Artificial Intelligence and Machine Learning in Health Care and Medical Sciences: Best Practices and Pitfalls. Cham: Springer; 2024. p. 33-94. doi:10.1007/978-3-031-39355-6_2
484. Choi RY, Coyner AS, Kalpathy-Cramer J, Chiang MF, Campbell JP. Introduction to Machine Learning, Neural Networks, and Deep Learning. Transl Vis Sci Technol. 2020;9(2):14. doi:10.1167/tvst.9.2.14
485. Simon G, Aliferis C. An Appraisal and Operating Characteristics of Major ML Methods Applicable in Healthcare and Health Science. In: Simon GJ, Aliferis C, editors. Artificial Intelligence and Machine Learning in Health Care and Medical Sciences: Best Practices and Pitfalls. Cham: Springer; 2024. p. 95-195. doi:10.1007/978-3-031-39355-6_3
486. Han R, Acosta JN, Shakeri Z, Ioannidis JPA, Topol EJ, Rajpurkar P. Randomised controlled trials evaluating artificial intelligence in clinical practice: a scoping review. Lancet Digit Health. 2024;6(5):e367-73. doi:10.1016/S2589-7500(24)00047-5
487. Moor M, Banerjee O, Abad ZSH, Krumholz HM, Leskovec J, Topol EJ, et al. Foundation models for generalist medical artificial intelligence. Nature. 2023;616(7956):259-65. doi:10.1038/s41586-023-05881-4
488. Schutte D, Zhang R. Considerations for Specialized Health AI & ML Modelling and Applications: NLP. In: Simon GJ, Aliferis C, editors. Artificial Intelligence and Machine Learning in Health Care and Medical Sciences: Best Practices and Pitfalls. Cham: Springer; 2024. p. 623-41. doi:10.1007/978-3-031-39355-6_14
489. Thirunavukarasu AJ, Ting DSJ, Elangovan K, Gutierrez L, Tan TF, Ting DSW. Large language models in medicine. Nat Med. 2023;29(8):1930-40. doi:10.1038/s41591-023-02448-8
490. Aliferis C, Simon G. Lessons Learned from Historical Failures, Limitations and Successes of AI/ML in Healthcare and the Health Sciences. Enduring Problems, and the Role of Best Practices. In: Simon GJ, Aliferis C, editors. Artificial Intelligence and Machine Learning in Health Care and Medical Sciences: Best Practices and Pitfalls. Cham: Springer; 2024. p. 543-606. doi:10.1007/978-3-031-39355-6_12
491. Aliferis C, Simon G. The Development Process and Lifecycle of Clinical Grade and Other Safety and Performance-Sensitive AI/ML Models. In: Simon GJ, Aliferis C, editors. Artificial Intelligence and Machine Learning in Health Care and Medical Sciences: Best Practices and Pitfalls. Cham: Springer; 2024. p. 289-340. doi:10.1007/978-3-031-39355-6_6
492. Collins GS, Dhiman P, Ma J, Schlussel MM, Archer L, Van Calster B, et al. Evaluation of clinical prediction models (part 1): from development to external validation. BMJ. 2024;384:e074819. doi:10.1136/bmj-2023-074819
493. Tejani AS, Klontzas ME, Gatti AA, Mongan JT, Moy L, Park SH, et al. Checklist for Artificial Intelligence in Medical Imaging (CLAIM): 2024 Update. Radiol Artif Intell. 2024;6(4):e240300. doi:10.1148/ryai.240300
494. Chen RJ, Wang JJ, Williamson DFK, Chen TY, Lipkova J, Lu MY, et al. Algorithmic fairness in artificial intelligence for medicine and healthcare. Nat Biomed Eng. 2023;7(6):719-42. doi:10.1038/s41551-023-01056-8
495. Van Calster B, Steyerberg EW, Wynants L, van Smeden M. There is no such thing as a validated prediction model. BMC Med. 2023;21(1):70. doi:10.1186/s12916-023-02779-w
496. Finlayson SG, Subbaswamy A, Singh K, Bowers J, Kupke A, Zittrain J, et al. The Clinician and Dataset Shift in Artificial Intelligence. N Engl J Med. 2021;385(3):283-6. doi:10.1056/NEJMc2104626
497. National Academy of Medicine; Krishnan S, Elliott A, editors. Generative Artificial Intelligence in Health and Medicine: Opportunities and Responsibilities for Transformative Innovation. Washington (DC): National Academies Press; 2025. doi:10.17226/28907
498. Kapoor S, Narayanan A. Leakage and the reproducibility crisis in machine-learning-based science. Patterns (N Y). 2023;4(9):100804. doi:10.1016/j.patter.2023.100804
499. Monaghan TF, Rahman SN, Agudelo CW, Wein AJ, Lazar JM, Everaert K, et al. Foundational Statistical Principles in Medical Research: Sensitivity, Specificity, Positive Predictive Value, and Negative Predictive Value. Medicina (Kaunas). 2021;57(5):503. doi:10.3390/medicina57050503

500. Van Calster B, Collins GS, Vickers AJ, Wynants L, Kerr KF, Barreñada L, et al. Evaluation of performance measures in predictive artificial intelligence models to support medical decisions: overview and guidance. Lancet Digit Health. 2025;7(12):100916. doi:10.1016/j.landig.2025.100916
501. Çorbacıoğlu ŞK, Aksel G. Receiver operating characteristic curve analysis in diagnostic accuracy studies: A guide to interpreting the area under the curve value. Turk J Emerg Med. 2023;23(4):195-8. doi:10.4103/tjem.tjem_182_23
502. Wong A, Otles E, Donnelly JP, Krumm A, McCullough J, DeTroyer-Cooley O, et al. External Validation of a Widely Implemented Proprietary Sepsis Prediction Model in Hospitalized Patients. JAMA Intern Med. 2021;181(8):1065-70. doi:10.1001/jamainternmed.2021.2626
503. Wong A, Currey D, Schwinne M, Park-Egan B, Meyer S, Gutting A, et al. Multicenter Prospective Validation of an Updated Proprietary Sepsis Prediction Model. JAMA Netw Open. 2026;9(2):e260181. doi:10.1001/jamanetworkopen.2026.0181
504. Vickers AJ, Holland F. Decision curve analysis to evaluate the clinical benefit of prediction models. Spine J. 2021;21(10):1643-8. doi:10.1016/j.spinee.2021.02.024
505. Collins GS, Moons KGM, Dhiman P, Riley RD, Beam AL, Van Calster B, et al. TRIPOD+AI statement: updated guidance for reporting clinical prediction models that use regression or machine learning methods. BMJ. 2024;385:e078378. doi:10.1136/bmj-2023-078378
506. Moons KGM, Damen JAA, Kaul T, Hooft L, Andaur Navarro C, Dhiman P, et al. PROBAST+AI: an updated quality, risk of bias, and applicability assessment tool for prediction models using regression or artificial intelligence methods. BMJ. 2025;388:e082505. doi:10.1136/bmj-2024-082505
507. Liu X, Cruz Rivera S, Moher D, Calvert MJ, Denniston AK, SPIRIT-AI and CONSORT-AI Working Group. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: the CONSORT-AI extension. Nat Med. 2020;26(9):1364-74. doi:10.1038/s41591-020-1034-x
508. Cruz Rivera S, Liu X, Chan AW, Denniston AK, Calvert MJ, SPIRIT-AI and CONSORT-AI Working Group, et al. Guidelines for clinical trial protocols for interventions involving artificial intelligence: the SPIRIT-AI extension. Nat Med. 2020;26(9):1351-63. doi:10.1038/s41591-020-1037-7
509. Vasey B, Nagendran M, Campbell B, Clifton DA, Collins GS, Denaxas S, et al. Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. Nat Med. 2022;28(5):924-33. doi:10.1038/s41591-022-01772-9
510. Sounderajah V, Guni A, Liu X, Collins GS, Karthikesalingam A, Markar SR, et al. The STARD-AI reporting guideline for diagnostic accuracy studies using artificial intelligence. Nat Med. 2025;31(10):3283-9. doi:10.1038/s41591-025-03953-8
511. Gallifant J, Afshar M, Ameen S, Aphinyanaphongs Y, Chen S, Cacciamani G, et al. The TRIPOD-LLM reporting guideline for studies using large language models. Nat Med. 2025;31(1):60-9. doi:10.1038/s41591-024-03425-5
512. Rajpurkar P, Chen E, Banerjee O, Topol EJ. AI in health and medicine. Nat Med. 2022;28(1):31-8. doi:10.1038/s41591-021-01614-0
513. Plana D, Shung DL, Grimshaw AA, Saraf A, Sung JJY, Kann BH. Randomized Clinical Trials of Machine Learning Interventions in Health Care: A Systematic Review. JAMA Netw Open. 2022;5(9):e2233946. doi:10.1001/jamanetworkopen.2022.33946
514. Makar J, Abdelmalak J, Con D, Hafeez B, Garg M. Use of artificial intelligence improves colonoscopy performance in adenoma detection: a systematic review and meta-analysis. Gastrointest Endosc. 2025;101(1):68-81.e8. doi:10.1016/j.gie.2024.08.033
515. Lukac PJ, Turner W, Vangala S, Chin AT, Khalili J, Shih YT, et al. Ambient AI Scribes in Clinical Practice: A Randomized Trial. NEJM AI. 2025;2(12). doi:10.1056/aioa2501000
516. Afshar M, Baumann MR, Resnik F, Hintzke J, Sullivan AG, Wills G, et al. A Pragmatic Randomized Controlled Trial of Ambient Artificial Intelligence to Improve Health Practitioner Well-Being. NEJM AI. 2025;2(12). doi:10.1056/aioa2500945
517. Chowdhury A, Casey M, Wilson J, Pollak KI, Goldstein BA, Bedoya A, et al. Comparing ambient scribes: a randomized crossover clinical trial addressing ambient scribe technologies' impact on physician burnout. J Am Med Inform Assoc. 2026;33(5):990-9. doi:10.1093/jamia/ocag018
518. Arun N, Gaw N, Singh P, Chang K, Aggarwal M, Chen B, et al. Assessing the Trustworthiness of Saliency Maps for Localizing Abnormalities in Medical Imaging. Radiol Artif Intell. 2021;3(6):e200267. doi:10.1148/ryai.2021200267
519. Ghassemi M, Oakden-Rayner L, Beam AL. The false hope of current approaches to explainable artificial intelligence in health care. Lancet Digit Health. 2021;3(11):e745-50. doi:10.1016/S2589-7500(21)00208-9

520. Gichoya JW, Banerjee I, Bhimireddy AR, Burns JL, Celi LA, Chen LC, et al. AI recognition of patient race in medical imaging: a modelling study. Lancet Digit Health. 2022;4(6):e406-14. doi:10.1016/S2589-7500(22)00063-2
521. National Academy of Medicine; Adams L, Fontaine E, Matheny M, Krishnan S, editors. An Artificial Intelligence Code of Conduct for Health and Medicine: Essential Guidance for Aligned Action. Washington (DC): National Academies Press; 2025. doi:10.17226/29087
522. Alber DA, Yang Z, Alyakin A, Yang E, Rai S, Valliani AA, et al. Medical large language models are vulnerable to data-poisoning attacks. Nat Med. 2025;31(2):618-26. doi:10.1038/s41591-024-03445-1
523. European Commission. AI Act: regulatory framework on artificial intelligence [Internet]. Brussels: European Commission; 2026 [cited 2026 Sep 27]. Available from: https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai
524. European Commission. Proposal for a Regulation of the European Parliament and of the Council amending Regulations (EU) 2017/745 and (EU) 2017/746 as regards simplifying and reducing the burden of the rules on medical devices and in vitro diagnostic medical devices, and amending Regulation (EU) 2022/123 and Regulation (EU) 2024/1689. COM(2025) 1023 final, 2025/0404 (COD) [Internet]. Strasbourg: European Commission; 2025 Dec 16 [cited 2026 Sep 27]. Available from: https://health.ec.europa.eu/document/download/25e7ea7c-cab3-40cf-86d9-d11f5e7744d8_en?filename=md_com_2025-1023_act_en.pdf
525. Ştefănigă SA, Cordoş AA, Ivascu T, Feier CVI, Muntean C, Stupinean CV, et al. Advancing Precision Oncology with Digital and Virtual Twins: A Scoping Review. Cancers (Basel). 2024;16(22):3817. doi:10.3390/cancers16223817
526. Council of the European Union. Artificial intelligence: Council and Parliament agree to simplify and streamline rules [Internet]. Brussels: Council of the EU; 2026 May 7 [updated 2026 May 18; cited 2026 Sep 27]. Available from: https://www.consilium.europa.eu/en/press/press-releases/2026/05/07/artificial-intelligence-council-and-parliament-agree-to-simplify-and-streamline-rules/
527. Regulation (EU) 2023/2854 of the European Parliament and of the Council of 13 December 2023 on harmonised rules on fair access to and use of data and amending Regulation (EU) 2017/2394 and Directive (EU) 2020/1828 (Data Act). Official Journal of the European Union. 2023 Dec 22;L 2023/2854.
528. European Parliament. Legislative Train Schedule: The Digital Omnibus Regulation proposal, 2025/0360(COD) [Internet]. Brussels: European Parliament; 2026 Aug 1 [cited 2026 Sep 27]. Available from: https://www.europarl.europa.eu/legislative-train/theme-a-new-plan-for-europe-s-sustainable-prosperity-and-competitiveness/file-digital-package
529. Regenstrief Institute. Download LOINC: LOINC version 2.83 [Internet]. Indianapolis (IN): Regenstrief Institute; 2026 Aug 19 [cited 2026 Sep 27]. Available from: https://loinc.org/downloads/
530. Google. Share files from Google Drive [Internet]. Google Drive Help; 2026 [cited 2026 Sep 27]. Available from: https://support.google.com/drive/answer/2494822
531. Google. Find what's changed in a file [Internet]. Google Docs Editors Help; 2026 [cited 2026 Sep 27]. Available from: https://support.google.com/docs/answer/190843
532. Google. Create a strong password & a more secure account [Internet]. Google Account Help; 2026 [cited 2026 Sep 27]. Available from: https://support.google.com/accounts/answer/32040
533. Google. Turn on 2-Step Verification [Internet]. Google Account Help; 2026 [cited 2026 Sep 27]. Available from: https://support.google.com/accounts/answer/185839
534. Cybersecurity and Infrastructure Security Agency. #StopRansomware Guide [Internet]. Washington (DC): CISA; 2023 Sep [cited 2026 Sep 27]. Available from: https://www.cisa.gov/stopransomware/ransomware-guide
535. National Library of Medicine (US). MeSH Browser [Internet]. Bethesda (MD): NLM; 2026 [cited 2026 Sep 27]. Available from: https://meshb.nlm.nih.gov/
536. Schiavenato M, Chu F. PICO: What it is and what it is not. Nurse Educ Pract. 2021;56:103194. doi:10.1016/j.nepr.2021.103194
537. Farič N, Potts HW, Heilman JM. Quality of Male and Female Medical Content on English-Language Wikipedia: Quantitative Content Analysis. J Med Internet Res. 2024;26:e47562. doi:10.2196/47562
538. Marquez-Romero JM, Lee Á, Soto-Cabrera E, Hernández-Curiel BC, Prado-Aguilar CA. Content and Quality of Information about Stroke in Wikipedia across Multiple Languages. Eur Neurol. 2022;85(4):308-12. doi:10.1159/000521938
539. Halboub E, Al-Ak'hali MS, Alqahtani AS, Abdulghani EA, Kamil MA, Alhajj MN, et al. Quality of web-based Arabic health information on dental implants: an infodemiological study. BMC Oral Health. 2023;23(1):232. doi:10.1186/s12903-023-02938-8

540. Cherrez-Ojeda I, Zuberbier T, Rodas-Valero G, Sanchez J, Rudenko M, Dramburg S, et al. Evaluation of the Quality and Reliability of ChatGPT-4's Responses on Allergen Immunotherapy Using Validated Instruments for Health Information Quality Assessment. Clin Transl Allergy. 2025;15(12):e70130. doi:10.1002/clt2.70130
541. Gerek ME, Önalan T, Çölkesen F, Arslan Ş. Evaluating large language models for WAO/EAACI guideline compliance in hereditary angioedema management. Allergol Immunopathol (Madr). 2025;53(4):51-9. doi:10.15586/aei.v53i4.1353
542. CHART Collaborative. Reporting guideline for chatbot health advice studies: the Chatbot Assessment Reporting Tool (CHART) statement. BMJ Med. 2025;4(1):e001632. doi:10.1136/bmjmed-2025-001632
543. World Health Organization. Hypertension. Fact sheet [Internet]. Geneva: WHO; 2025 Sep 25 [cited 2026 Sep 27]. Available from: https://www.who.int/news-room/fact-sheets/detail/hypertension
544. Google. Create and use pivot tables. Google Docs Editors Help [Internet]. Mountain View (CA): Google; 2026 [cited 2026 Sep 27]. Available from: https://support.google.com/docs/answer/1272900
545. World Health Organization. ICD-11 Coding Tool, Mortality and Morbidity Statistics (MMS), release 2026-01 [Internet]. Geneva: WHO; 2026 [cited 2026 Sep 27]. Available from: https://icd.who.int/ct
546. World Health Organization. ICD API [Internet]. Geneva: WHO; 2026 [cited 2026 Sep 27]. Available from: https://icd.who.int/icdapi
547. Digital Scholar. Zotero 10: download [Internet]. Digital Scholar; c2026 [cited 2026 Sep 27]. Available from: https://www.zotero.org/download/
548. Digital Scholar. Zotero documentation: Word processor plugin integration [Internet]. Digital Scholar; c2026 [cited 2026 Sep 27]. Available from: https://www.zotero.org/support/word_processor_integration
549. Digital Scholar. Zotero documentation: Citation styles [Internet]. Digital Scholar; c2026 [cited 2026 Sep 27]. Available from: https://www.zotero.org/support/styles
550. Elsevier. Mendeley Reference Manager for Windows, version 2.149.0 [Internet]. Elsevier; c2026 [cited 2026 Sep 27]. Available from: https://www.mendeley.com/download-reference-manager/windows
551. Digital Scholar. Zotero documentation: How do I export my Zotero library? [Internet]. Digital Scholar; [updated 2025 Nov 24; cited 2026 Sep 27]. Available from: https://www.zotero.org/support/kb/exporting
552. Elsevier. How can I add new references to my Library in Mendeley Reference Manager? [Internet]. Elsevier Support Center; c2026 [cited 2026 Sep 27]. Available from: https://www.elsevier.support/mendeley/answer/how-can-i-add-new-references-to-my-library-in-mendeley-reference-manager
553. PhysioNet. EEG Motor Movement/Imagery Dataset, version 1.0.0 [Internet]. Cambridge (MA): MIT Laboratory for Computational Physiology; 2026 [cited 2026 Sep 27]. Available from: https://physionet.org/content/eegmmidb/1.0.0/
554. MIT Laboratory for Computational Physiology. wfdb: the WFDB Python package, version 4.3.1 [Internet]. Python Package Index; 2026 [cited 2026 Sep 27]. Available from: https://pypi.org/project/wfdb/
555. WFDB Python package documentation, version 4.3.1: WFDB records and annotations (rdrecord, rdann) [Internet]. Read the Docs; 2026 [cited 2026 Sep 27]. Available from: https://wfdb.readthedocs.io/en/latest/io.html
556. Schroeder AB, Dobson ETA, Rueden CT, Tomancak P, Jug F, Eliceiri KW. The ImageJ ecosystem: Open-source software for image visualization, processing, and analysis. Protein Sci. 2021;30(1):234-49. doi:10.1002/pro.3993
557. ImageJ community. Fiji Downloads [Internet]. imagej.net; 2026 [cited 2026 Sep 27]. Available from: https://imagej.net/software/fiji/downloads
558. Rasband W. ImageJ Release Notes, version 1.54t [Internet]. Bethesda (MD): National Institutes of Health; 2026 May 16 [cited 2026 Sep 27]. Available from: https://wsr.imagej.net/notes.html
559. ImageJ. ImageJ Documentation: Menu commands [Internet]. Bethesda (MD): National Institutes of Health; [cited 2026 Sep 27]. Available from: https://imagej.net/ij/docs/menus/file.html
560. The Cancer Imaging Archive. Data Usage Policies and Restrictions [Internet]. TCIA; [cited 2026 Sep 27]. Available from: https://www.cancerimagingarchive.net/data-usage-policies-and-restrictions/
561. Gahm NA, Rueden CT, Evans EL, Selzer G, Hiner MC, Chacko JV, et al. New Extensibility and Scripting Tools in the ImageJ Ecosystem. Curr Protoc. 2021;1(8):e204. doi:10.1002/cpz1.204
562. Rueden CT, Hiner MC, Evans EL, Pinkert MA, Lucas AM, Carpenter AE, et al. PyImageJ: A library for integrating ImageJ and Python. Nat Methods. 2022;19(11):1326-7. doi:10.1038/s41592-022-01655-4
563. OpenEMR Project. OpenEMR 8.4.1 release [Internet]. GitHub; 2026 Sep 20 [cited 2026 Sep 27]. Available from: https://github.com/openemr/openemr/releases

564. OpenEMR Project. OpenEMR online demo (version 8.4.1) [Internet]. OpenEMR Foundation; 2026 [cited 2026 Sep 27]. Available from: https://www.open-emr.org/demo/
565. OpenEMR Project. OpenEMR FHIR API (FHIR_README) [Internet]. GitHub; 2026 [cited 2026 Sep 27]. Available from: https://github.com/openemr/openemr/blob/master/FHIR_README.md
566. HAPI FHIR Project. HAPI FHIR Test/Demo Server, R4 endpoint [Internet]. 2026 [cited 2026 Sep 27]. Available from: https://hapi.fhir.org/baseR4
567. HL7 International. FHIR R5 (v5.0.0). Observation Vital Signs Profile [Internet]. Ann Arbor (MI): HL7 International; 2023 [cited 2026 Sep 27]. Available from: https://hl7.org/fhir/R5/observation-vitalsigns.html
568. OpenEMR Project. HOWTO: Create a New Patient Record - OpenEMR v7 [Internet]. OpenEMR Project Wiki; 2026 [cited 2026 Sep 27]. Available from: https://www.open-emr.org/wiki/index.php/HOWTO:_Create_a_New_Patient_Record_-_OpenEMR_v7
569. Murphy EM, Stein A, Pahwa R, McGuire M, Kumra T. Difference in medical student performance in a standardized patient encounter between telemedicine and in-person environments. Med Educ Online. 2024;29(1):2388422. doi:10.1080/10872981.2024.2388422
570. Gunner CK, Eisner E, Watson AJ, Duncan JL. Teaching webside manner: development and initial evaluation of a video consultation skills training module for undergraduate medical students. Med Educ Online. 2021;26(1):1954492. doi:10.1080/10872981.2021.1954492
571. Jitsi. Jitsi Meet, stable release 2.0.11248 [software on the Internet]. 2026 Sep 14 [cited 2026 Sep 27]. Available from: https://github.com/jitsi/jitsi-meet/releases
572. Pozo P, Landino MC, Maga JM, Jorge LM, Araya R, Horn DB, et al. Situational Awareness in Telehealth: A Virtual Standardized Patient Case for Transitioning Preclinical to Clinical Medical Students. MedEdPORTAL. 2025;21:11517. doi:10.15766/mep_2374-8265.11517
573. Lesselroth B, Monkman H, Palmer R, Kuziemsky C, Liew A, Foulks K, et al. Assessing Telemedicine Competencies: Developing and Validating Learner Measures for Simulation-Based Telemedicine Training. AMIA Annu Symp Proc. 2023;2023:474-83.
574. Hart A, Romney D, Sarin R, Mechanic O, Hertelendy AJ, Larson D, et al. Developing Telemedicine Curriculum Competencies for Graduate Medical Education: Outcomes of a Modified Delphi Process. Acad Med. 2022;97(4):577-85. doi:10.1097/ACM.0000000000004463
575. World Health Organization. WHO guidelines [Internet]. Geneva: WHO; 2026 [cited 2026 Sep 27]. Available from: https://www.who.int/publications/who-guidelines
576. Adobe. Create self-signed digital IDs in Acrobat [Internet]. Adobe Help Center; 2025 Sep 23 [cited 2026 Sep 27]. Available from: https://helpx.adobe.com/acrobat/desktop/protect-documents/manage-digital-ids/create-self-signed-id.html
577. Adobe. Validate digital signatures in Adobe Acrobat [Internet]. Adobe Help Center; 2025 Sep 23 [cited 2026 Sep 27]. Available from: https://helpx.adobe.com/acrobat/desktop/e-sign-documents/manage-digital-signatures/validate-digital-sign.html
578. Google. Check if a site's connection is secure [Internet]. Google Chrome Help; 2026 [cited 2026 Sep 27]. Available from: https://support.google.com/chrome/answer/95617
579. European Commission. EU/EEA Trusted List Browser, eIDAS Dashboard [Internet]. Brussels: European Commission; 2026 [cited 2026 Sep 27]. Available from: https://eidas.ec.europa.eu/efda/trust-services/browse/eidas/tls
580. JGraph Ltd. draw.io (diagrams.net) [Internet]. Northampton: JGraph; 2026 [cited 2026 Sep 27]. Available from: https://www.drawio.com/
581. JGraph Ltd. drawio-desktop v31.5.3 [software]. GitHub; 2026 Sep 26 [cited 2026 Sep 27]. Available from: https://github.com/jgraph/drawio-desktop/releases/tag/v31.5.3
582. JGraph Ltd. Flowcharts with swimlanes [Internet]. draw.io documentation; 2026 [cited 2026 Sep 27]. Available from: https://www.drawio.com/docs/diagram-types/swimlane-diagrams/
583. Google. Google Sheets function list [Internet]. Mountain View (CA): Google; 2026 [cited 2026 Sep 27]. Available from: https://support.google.com/docs/table/25273?hl=en
584. Microsoft. SUMPRODUCT function [Internet]. Redmond (WA): Microsoft; 2026 [cited 2026 Sep 27]. Available from: https://support.microsoft.com/en-us/office/sumproduct-function-16753e75-9f68-4874-94ac-4d2145a2fd2e
585. Antonacci G, Lennox L, Barlow J, Evans L, Reed J. Process mapping in healthcare: a systematic review. BMC Health Serv Res. 2021;21(1):342. doi:10.1186/s12913-021-06254-1

586. Gartner JB, Landa P, Haren MT, Lemaire C, Tanfani E, Paquet C, et al. The Implementation of a Business Process Model and Notation for Modeling Patient Health Care Trajectories: Systematic Review. J Med Internet Res. 2026;28:e78506. doi:10.2196/78506
587. Google. Add & edit a chart or graph. Google Docs Editors Help [Internet]. Mountain View (CA): Google; 2026 [cited 2026 Sep 27]. Available from: https://support.google.com/docs/answer/63824
588. The Document Foundation. Download LibreOffice (version 26.8.0) [Internet]. Berlin: The Document Foundation; 2026 [cited 2026 Sep 27]. Available from: https://www.libreoffice.org/download/download-libreoffice/