American Journal of
Biomedical Science & Research

www.biomedgrid.com
ISSN: 2642-1747

Review Article

Copyright © All rights are reserved by Mari Manjgaladze

From Definitions to Clinical Implementation: Introduction to Artificial Intelligence for Health Professions – A Combined Analytical Review of the Textbooks

Mari Manjgaladze[1,2*] and Ketevan Aptarashvili[2]

[1]*East-West University, Tbilisi, Georgia*

[2]*East European University, Tbilisi, Georgia*

***Corresponding author:** Mari Manjgaladze, East-West University; East European University, Tbilisi, Georgia.

To Cite This article: *Mari Manjgaladze* and Ketevan Aptarashvili, From Definitions to Clinical Implementation: Introduction to Artificial Intelligence for Health Professions – A Combined Analytical Review of the Textbooks. Am J Biomed Sci & Res. 2026 32(6) AJBSR.MS.ID.004216,*

DOI: 10.34297/AJBSR.2026.32.004216

Received: November 01, 2026 **Published:** November 12, 2026

Abstract

Artificial Intelligence (AI) is developing rapidly and has already become part of everyday clinical practice. At the same time, medical curricula do not yet adequately build the AI literacy that future clinicians will need to work safely with AI. Purpose-built textbooks are gradually emerging to fill this gap, and their pedagogical quality now has a direct bearing on patient care. This paper reviews and evaluates two editions of a textbook by A. Gagnidze – the English-language Introduction to Artificial Intelligence for Health Professions and the Georgian-language Introduction to Artificial Intelligence for Medical Professionals (East-West University, 2026) – as a learning resource for medical students and practising health professionals, and discusses how they can be integrated into the curriculum. Both editions share a common core of six modules, while the Georgian edition adds a seventh, integrative module that reviews specific medical AI systems already deployed in clinical practice. The paper is not limited to these textbooks: it also offers a comparative analysis with other textbooks in the field. The review also takes into account the companion journal article in which the textbook's author and the first author of this review condensed the textbook's clinical content into a peer-reviewed narrative review, Artificial Intelligence in Medicine: Opportunities, Risks, and Prospects for Georgia (Am J Biomed Sci & Res, 2026). The thematic analysis follows the six common modules, evaluates the seventh module of the Georgian edition separately, compares the structure and pedagogical apparatus of the editions, examines how the companion article tests the textbook against independent evidence, identifies strengths and opportunities for development, and offers practical conclusions. The common modules provide an accessible, deliberately plain-language, medicine-oriented introduction covering definitions of and approaches to AI, its history, machine learning, generative AI, ethical and social challenges, and the future of the health professions. Their particular strengths are plain language and analogies, analysis of clinical cases, a balanced treatment of opportunities and risks, practical tools (for example, the DOCTOR checklist), a consistent emphasis on equity, and a clear Georgian context, which is rare in this field. The seventh module of the Georgian edition carries this foundation into clinical reality: it discusses diagnostic systems in radiology, pathology, ophthalmology and cardiology, clinical decision support systems and LLM-based tools, proposes eight criteria for selecting a system, and introduces a consistent system of "warnings" that links each tool to the limits of its reliability. The main directions for future development include linking empirical claims to additional sources, deepening the technical discussion of model validation, and periodically updating long-term forecasts. The seventh module already responds to some of these directions; however, by its nature it is the fastest-ageing part of the book, and every quantitative claim in it needs to be supported by a source and a date.

This work is licensed under Creative Commons Attribution 4.0 License

Abstract

The companion article goes a considerable way towards this: it links the key clinical figures to peer-reviewed and regulatory sources, develops the DeepFAN, Aidoc and IDx-DR cases, applies the analytical tools of Module 7 to Georgian health-system data through a diabetic-retinopathy screening scenario for Kakheti, and restores a phased roadmap for Georgia. The comparison also reveals several figures that differ between the two texts – for example, Aidoc's real-world performance and the carbon footprint of training GPT-3 – which future editions should harmonize. There is also an asymmetry in the pedagogical apparatus: the English edition is more complete in terms of quizzes, cases and tests, whereas the Georgian edition is broader in content. Both editions provide a solid foundation for developing AI literacy in the health professions. Harmonizing them – adding the seventh module to the English edition and restoring quizzes, cases and tests to the Georgian one – and aligning them with the sourced figures of the companion article would make the textbook the basis of a thorough bilingual introductory course in medical AI.

Keywords: Artificial intelligence; Medical education; AI literacy; Healthcare; Generative AI; Clinical AI systems; Ethics; Evidence traceability; Georgia

Introduction

For medicine, artificial intelligence is no longer a distant prospect – it is already a present-day reality, embedded in diagnostics, documentation, medication selection and clinical decision support. AI-based systems today, for example, detect diabetic retinopathy and tumours on medical images with accuracy comparable to that of specialists [1], predict the onset of sepsis before overt clinical signs appear, and accelerate the selection of effective medicines [2,21]. The COVID-19 pandemic accelerated this process considerably: telemedicine, triage tools and predictive models entered everyday practice faster than any planned reform could have achieved. The textbook opens by emphasizing exactly this and concludes that AI in medicine "is not the future – it is already reality". Understanding its capabilities and limitations is therefore a professional necessity rather than a matter of personal curiosity. Against this background, the resulting educational gap is becoming ever harder to ignore. Medical curricula, overloaded with content and constrained by strict accreditation cycles, have been relatively slow to recognize the importance of AI literacy [33]. At the same time, it is clear that future physicians, whatever their specialty, will work with AI systems. This gap in AI literacy carries concrete risks: blind trust in opaque outputs, misuse of tools whose limitations are poorly understood, and a lack of ability to critically evaluate machine-generated recommendations at the bedside. A clinician who cannot assess an algorithm's reliability, validity or bias is not a neutral user – they become vulnerable. Building basic AI literacy is therefore a matter not only of physicians' professional development but also of patient safety [19]. A still fairly small but growing body of textbooks has emerged in response to this need. For example, Eric Topol's Deep Medicine [28] makes the case for the humanistic value of AI-enhanced care for a broad audience; Shortliffe and Cimino's Biomedical Informatics [25] remains the technical reference work of the field; and regulatory and policy documents such as the World Health Organization's Ethics and Governance of Artificial Intelligence for Health [35] set out principles of governance.

The contribution of A. Gagnidze's textbooks stands out in four respects: they are written for readers without technical training; every chapter puts medicine in the foreground and uses everyday examples to make the text easy to follow; they are clearly oriented towards the Georgian context, which is almost absent from English-language materials and publications; and, in the Georgian edition, they move from general concepts to specific tools already deployed in clinics, which is rare in introductory textbooks. The textbook exists in two editions. The English-language edition, Introduction to Artificial Intelligence for Health Professions [6], comprises six thematic modules and stands out for its rich pedagogical apparatus. The Georgian-language edition, Introduction to Artificial Intelligence for Medical Professionals [7], retains the content core of Modules 1–6 and adds a seventh module, "A Practical Review of Medical AI Systems". The present review therefore treats both editions as a single work: Sections 3–7 deal with the common modules, and their conclusions apply to both editions; Section 8 is devoted to the seventh module of the Georgian edition; and the structural differences between the editions are summarized in Section 2. In October 2026, the clinical content of the textbook – above all that of Module 7 and the Georgian material – was also condensed into a peer-reviewed narrative review in this journal, written by the textbook's author together with the first author of the present review [8]. Because the article re-examines the textbook's claims against independent sources, Section 9 uses it as a benchmark for judging which of the textbook's figures are traceable and where the two texts diverge.

The review has seven aims:

a) To provide an analytical reading of both editions;

b) To evaluate separately the seventh module of the Georgian edition as a new integrative component and to identify the differences between the editions;

c) To assess the textbook's strengths, opportunities for development and academic contribution;

d) To determine its suitability as a resource for medical education;

e) To propose how it might be integrated into the curriculum;

f) To identify areas of development that future editions and supplementary materials should address;

g) To compare the textbook with its companion journal article [8] and to identify which of the directions for development the article has already addressed.

The review aims at a balanced assessment, and its significance extends beyond the discussion of a single textbook. As AI literacy is transformed from a desirable add-on to medical education into a core competence, a review of this textbook contributes to the growing scholarly discussion of how health professionals should be taught to use artificial intelligence and to think alongside it.

Two Editions: A Common Core and Differences

A comparison of the two editions shows that they represent two partly different versions of the same course. The main text of Modules 1–6 – definitions, analogies, medical examples, tables and discussion topics – is the same in content, so the assessments made in Sections 3–7 apply equally to both editions. The differences appear mainly on three levels: the number of modules, the pedagogical apparatus, and the placement of the Georgian context (Table 1).

Table 1: Structural comparison of the English and Georgian editions.

Feature	English Edition	Georgian Edition
Number of modules	6 thematic modules	7 modules: 6 common + integrative Module 7
Main text of Modules 1-6	Common content core	Common content core
Quick quizzes	At the end of every module	None
Practical cases	In every module (Cases 1.1-6.3)	Only in Module 6 (6.1-6.3); announced in Module 7 but absent from the text
Tests, practical exercises	Present (tests with answers; 7 exercises)	No tests in the text
Georgian-context inserts	Distributed across modules: "The Georgian Connection" (Mod. 2), Georgian healthcare equity and Case 5.2 (Mod. 5), 6.7 and the 2027-2037 vision (Mod. 6)	Mostly gathered at the end of Module 7 ("Medical AI in Georgia"); the Georgian NLP example and discussion questions are retained
Specific clinical systems	Individual illustrative examples	Module 7: more than 20 systems, a typology of assistance levels, 8 selection criteria
Supplementary material	"A Dialogue Across Time" and a 20-question self-assessment	"Frequently Asked Questions" (FAQ, 40 questions and answers)

The table reveals a certain "asymmetry" between the editions. The English edition is pedagogically more complete: each module ends with a quick quiz, practical cases and discussion topics; the course has pre- and post-tests, seven practical exercises and a supplementary reading, "A Dialogue Across Time", in which Hippocrates, Alan Turing and Elon Musk, in an imagined conversation, set three different visions of medicine, the machine and the future against one another. This is an original teaching device that restates the core ideas of the course in dramatic form and opens with the author's clear disclaimer that the characters' words are invented and are not quotations. The Georgian edition, by contrast, is broader in content – the seventh module and the 40-question FAQ enrich it with elements of a clinical reference – but it does not include the cases, quizzes, tests and exercises of Modules 1–5. The placement of the Georgian context also differs. In the English edition it is distributed across the modules and links almost every topic to local healthcare. In the Georgian edition this material is mostly gathered in one place – at the end of the seventh module. As a result, a Georgian-language reader will find several analytical inserts on Georgian healthcare – for example, the Georgian case of algorithmic bias and the vision for 2027–2037 – only in the English edition. Harmonizing the editions is a natural task for future reprints (see 11.3).

The Textbook's Conceptual Framework

The first module of the textbook establishes the conceptual foundation on which the rest of the book rests. Its central idea is to present intelligence not as a binary property (present/absent) but as a graded capacity. This is a highly effective pedagogical approach, because it protects the novice reader from two common errors: perceiving every automated system as AI, and imagining AI as a single, monolithic and absolutely infallible technology.

The Four Faces of AI

The textbook organizes definitions of AI by means of a well-known matrix that distinguishes systems on two dimensions: whether they aim to reproduce human behaviour or to approximate ideal rationality, and whether the target is thinking or acting. The reader is thus introduced to four well-known research traditions: (Table 2).

Table 2: The four faces of AI [6,7].

	Human-Like	Ideal / Rational
Thinking	Cognitive modelling (systems that imitate human thought)	Laws of thought (logic, in the Aristotelian tradition)
Acting	Turing test (behaviour indistinguishable from a human's)	Rational agent (acting to achieve the best outcome)

This framework is standard, and its inclusion in an introductory text is appropriate. It reflects the organizing principle used by Russell and Norvig [22] in their well-known textbook. Two observations are important here:

a) First, the matrix is deliberately simplified and does not capture the paradigms the reader will encounter later (for example, situated cognition and the synthesis of modern deep learning, which fits neatly into none of the cells).

b) Second, the textbook explicitly names the rational-agent principle as its main model.

This choice is pedagogically justified, but it is nonetheless a choice. The text would gain additional value by explaining clearly that the rational-agent model is an analytical framework for describing clinical reality, not a claim of functional equivalence between physician and machine.

The Rational Agent in Medicine

Having chosen the rational agent as its guiding thread, the textbook applies it directly to the clinical setting. An agent is defined as anything that perceives its environment through sensors and acts upon it through actuators. In medical terms, sensors correspond to electronic health records, imaging and laboratory data, and actuators to alerts, recommendations and orders. The physician is presented as a rational agent who perceives information through history-taking and examination and acts through diagnosis and treatment. Moreover, the physician-agent operates under conditions of bounded rationality – that is, seeks the best available decision within the real limits of information, time and cognitive load. This is an elegant, unifying approach. The concept of bounded rationality is especially well suited to medicine, where the idealized model of a fully informed decision-maker does not work in real clinical settings. Clinical reasoning, however, is not merely constrained optimization: it includes pattern recognition developed through experience, narrative interpretation and communication with the patient. Representing both physician and machine as "agents" should therefore not erase the essential differences between them.

The Spectrum of Intelligence

One of the module's most memorable teaching devices is a scale from 0 to 10 on which systems are placed according to their adaptive capabilities. According to the textbook, a simple calculator sits at about 0, a traditional search engine at 3, an ELIZA-type chatbot at 6, a modern large language model such as ChatGPT at 8, human-level general intelligence at about 9.5, and hypothetical superintelligence at 10 (the last two do not yet exist). As a mnemonic device, the scale is useful: it makes visible the book's core idea that intelligence is not a binary concept. The scale is not, however, a precise measuring instrument – the numbers are illustrative only and not scientifically measured. Against the background of the wider literature, this framework is closest in spirit to the rigorous taxonomy of Russell and Norvig [22], although it presents technical issues at a more intuitive level. It is more schematic than Topol's clinically grounded arguments [28] and far from the policy register of the WHO document [35]. For the intended audience, this is the right calibration. The book does not try to train computer scientists: its aim is to prepare clinicians who can reason about and use these technologies, and its conceptual framework serves precisely this aim.

Historical and Technical Foundations

The second and third modules carry the intellectual load of the book: the second gives the reader historical context, and the third the technical foundations of machine learning. It is they that determine whether the reader acquires a durable mental model or merely absorbs terminology. Overall, these modules succeed through a simple but clear technical treatment that is particularly strong in its intuitive explanations, while the discussion of the quantitative mechanisms of validation retains potential for further deepening.

Historical Foundations

The historical periodization presented in the textbook is widely accepted and accurate: the McCulloch–Pitts neuron (1943), Turing's imitation game (1950) [29], the 1956 Dartmouth workshop that gave the field the name "artificial intelligence", the first chatbot ELIZA (1966), the era of expert systems and MYCIN in the 1970s, the subsequent "AI winters", Deep Blue's victory over Kasparov (1997), AlphaGo's victory over Lee Sedol (2016), and, in 2022, the arrival of ChatGPT as the first mass-market generative system. The textbook's decision to emphasize the "AI winters" deserves praise, since the cyclical pattern of overpromising followed by retreat is one of the most important lessons history offers a newcomer to a field prone to hype. Two features raise this chapter above a simple chronology. The first is a timeline of specifically medical milestones and a short "Georgian context" that links global history to the reader's own region. The second, and pedagogically more valuable, is the discussion of the AI implementation gap using the example of MYCIN: this 1970s expert system, with about five hundred rules, performed at or above the level of junior physicians in studies but

was never deployed in practice. The system failed not for lack of accuracy but because of problems of workflow integration, unclear responsibility and clinicians' reasonable refusal to trust blindly a system they could neither question nor verify. Presenting MYCIN as a lesson gives the reader an analytical tool: the awareness that clinical value depends on integration and trust, not only on performance metrics. This proposition remains fully relevant to modern models. When comparing the editions, it is worth noting that "The Georgian Connection" insert and the extended discussion of MYCIN's implementation gap (Case 2.1) appear in the English edition. In the Georgian edition, the MYCIN lesson is retained as a short "historical lesson" with discussion questions, and its logic is fully developed in the seventh module through the examples of IBM Watson and Epic's sepsis model (see 8.3). A minor inconsistency can also be noted in the Georgian edition: in the timeline, MYCIN's non-adoption is attributed to the high cost of computers, whereas the subsequent "historical lesson" attributes it to integration problems; the two explanations should be reconciled.

The Turing Test, the Chinese Room and Moravec's Paradox

The book takes the Turing test seriously: it presents its standard configuration, sets out John Searle's well-known "Chinese room" argument [23], and poses the pointed clinical question of whether a chatbot can genuinely perform the role of a physician. In the era of fluent large language models, this question is even more pressing: modern AI systems produce clinically convincing text without any confirmed evidence, which gives philosophical arguments practical significance. Moravec's paradox is used well in the book: tasks that humans find hard (chess, memorization, abstract reasoning) are relatively easy for machines, whereas tasks that humans find easy (perception, movement, common sense, empathy, emotion) remain difficult for machines. The medical interpretation is clear to the reader from the outset: AI succeeds, for example, in image analysis but struggles with physical examination and interaction with patients. The chapter would gain further value if a future edition reflected the development of multimodal systems as well: these systems have already begun to partly overcome the paradox, which indicates that the boundary described in the textbook is dynamic.

Machine Learning and its Three Types

The third module explains machine learning through two well-chosen analogies and then presents the standard three-part taxonomy with medical examples (Table 3).

Table 3: The three types of machine learning [6,7], Module 3.

Type	How it Learns	Medical Example
Supervised	From labelled examples	Classification of labelled X-ray images or pathology slides
Unsupervised	From unlabelled data, by discovering structure	Grouping patients into previously unrecognized subgroups
Reinforcement	Through reward and penalty over time	Optimizing drug dosing in the intensive care unit

The analogies used in the textbook are apt and effective, and the medical examples are appropriate. The room for improvement lies not in the classification itself but in the level of detail of the surrounding technical issues. Validation, overfitting and generalization are discussed briefly, while semi-supervised, self-supervised and transfer learning are mentioned and need further elaboration. As noted above, the textbook deliberately chooses a simplified presentation of technical details, which is appropriate for its target audience.

Deep Learning, the Black Box and the Data Pipeline

The discussion of neural networks in the textbook is intuitive and easy to follow. Input, hidden and output layers are explained clearly, and a "deep" network is defined as one with many hidden layers. The textbook addresses the black-box problem directly: a model may have billions of parameters whose combined logic no human can read, so the model can answer the question "what?" without answering the question "why?". The textbook contrasts this with the decision tree as a transparent but often less accurate alternative, and describes the data-processing chain – from raw data through preprocessing, feature extraction, training, validation and deployment to clinical use. The trade-off between accuracy and transparency is exactly what should be established early in a clinician's awareness. The focus of the discussion here is on identifying the problem, while the tools for solving it are a natural direction for future expansion. Explainable AI techniques – SHAP, LIME, attention maps and other methods – are practical ways of responding to the black-box problem. Mentioning them briefly already makes the reader aware of the problem; a more detailed explanation would also equip the reader with practical tools.

Generative Artificial Intelligence in Medicine

The fourth module addresses the technology that, more than any other, has made AI part of the discussion about clinicians' everyday work, and it is one of the most topical and practically useful chapters of the textbook. Its main merit is that it presents generative AI not as a magical technology but as a statistical system with specific capabilities and characteristic errors.

What Generative AI is and how Large Language Models Work

The chapter distinguishes generative systems from earlier AI

by their ability to create new content – text, images, code – and illustrates the distinction with a memorable analogy comparing a librarian and a writer: a search engine points you to existing sources, whereas a generative model creates new content. It then explains that large language models are trained on enormous internet-scale corpora, work by predicting the next token and are based on the transformer architecture, before describing the evolution of the GPT family (Table 4).

Table 4: Evolution of the GPT family [6,7].

Version	Year	Parameters	Main Innovation
GPT-1	2018	117 million	The first GPT
GPT-2	2019	1.5 billion	Text generation
GPT-3	2020	175 billion	Few-shot learning
GPT-4	2023	~1 trillion (estimated)	Multimodality
GPT-4o	2024	-	Full multimodality, speed
GPT-5	2025	-	Improved reasoning, accuracy

The textbook rightly notes that ChatGPT's decisive innovation was no less social than technical: a free, simple interface made the technology accessible to the general public. It also compares three widely used AI assistants (ChatGPT, Claude and DeepSeek) and, valuably, advises cross-checking the answer to a clinical question across several assistants rather than relying on one. A reader seeking more depth on multimodal generation – the fusion of text, image and audio that is changing reporting in radiology and pathology – will find only a general outline here; this topic could naturally become a separate chapter in a future edition of the textbook.

Hallucinations: The Central Clinical Risk

The discussion of hallucinations is one of the book's most important contributions from the standpoint of safe practice, because hallucinations are presented not simply as random glitches but as a fundamental risk of generative systems. The textbook distinguishes several types of medical hallucination – factual, source, diagnostic, drug-related and guideline-related – and proposes practical measures: requesting and verifying sources, and checking answers against external, reliable sources (Table 5).

Table 5: Types of medical hallucination [6,7].

Type of Hallucination	What it Looks Like in Medicine
Factual	A clinical fact stated with confidence but incorrect
Source	Citation of a plausible but non-existent paper
Diagnostic	A specific diagnosis not supported by the data provided
Drug-related	An incorrect dose, interaction or contraindication
Guideline-related	A recommendation wrongly attributed to an authoritative guideline

This taxonomy is well constructed and directly useful in practice. In particular, the category of source hallucination addresses precisely the error that poses the greatest danger to a student writing a literature review with the help of AI.

Prompt Engineering, Gans and Documentation

The chapter offers practical guidance on writing prompts – be specific, provide examples, assign a role, ask for step-by-step reasoning – and illustrates the difference between weak and strong prompts with medical examples, including assigning the model the role of a specialist and asking it to cite guidelines such as those of the European Society of Cardiology. It discusses Generative Adversarial Networks (GANs) and deepfakes, addresses disinformation, fraud and non-consensual content with appropriate concern, and extends the discussion to falsified medical images. It also covers the use of generative AI in clinical documentation (e.g. generating notes and patient instructions), weighing real efficiency gains against risks to accuracy, privacy and accountability, and treats verification by the physician as an essential safeguard. The open question posed in the chapter – whether prompt engineering is a durable profession or a transitional skill – is very apt. A natural continuation of this section would be a discussion of the regulatory framework that the use of AI in documentation requires: HIPAA, GDPR and their Georgian counterparts.

Ethical and Social Challenges

The fifth module is the ethical core of the book and one of its strongest chapters precisely because ethics is not isolated here

but woven into every case and paragraph. This reflects how ethical questions arise in real practice – intertwined with technical choices rather than bolted on afterwards [3]. For an introductory text, the thematic scope of the chapter is impressive.

The Black-Box Dilemma and Explainability

The module returns to the problem of the "black box", or opacity – this time as an ethical rather than merely technical issue. High-stakes decisions require justification: a system that recommends chemotherapy without interpretable reasoning places both physician and patient in an untenable position. In response, the textbook proposes explainable AI and uses the metaphor of a "glass box" to describe the goal. The metaphor is useful, and the chapter's potential would be increased further by a detailed explanation of the specific explainable AI techniques described in the literature and by linking the "right to explanation" discussed under the GDPR to a clinical scenario. It is understandable, however, that this would make the textbook considerably more technical, so the author's choice to preserve the accessibility of the text is pedagogically justified.

Algorithmic Bias

The discussion of bias is concrete and the examples are well chosen: the text moves from general technological cases to medical ones and thereby shows that bias is not an abstraction but can be a matter of life and death (Table 6).

Table 6: Examples of algorithmic bias [6,7].

Example	Nature of the Bias
Amazon's hiring algorithm	Discriminated against female candidates; learned from historical hiring data
UK A-level algorithm (2020)	Lowered the grades of pupils from less well-resourced schools
Pulse oximetry	Less accurate for patients with darker skin
Kidney disease algorithm (2019)	A race-based correction distorted the assessment of access to treatment

The pulse oximetry and kidney disease examples are particularly valuable because they are directly medical, recent and associated with serious consequences. They echo the argument presented by Obermeyer and colleagues in Science about racial bias in a population health algorithm [17]. The textbook explains that bias arises mainly from unrepresentative training data – the model absorbs historical inequality and then reproduces it faster and at scale. A natural extension of the discussion of bias would be ways of mitigating it: fairness metrics, deliberately diverse datasets, reporting of results by subgroup, and independent audit. These are the countermeasures a future clinician needs to know, and they deserve fuller discussion in the future, although it is understandable that this would make the textbook more technical.

Digital Poverty and Equity in the Georgian Context

Here the textbook offers one of its most original contributions. It develops the concept of digital poverty along three dimensions – access, quality and competence – and grounds it in a comparison between a well-equipped Tbilisi hospital and a resource-poor rural clinic. The discussion of equity is not mechanically transferred from a Western context but stems from the health system in which most of its readers actually work. The textbook's recommendations – digitization, the creation of local databases, investment in infrastructure, and training – are sensible, and the framework usefully links individual applications of clinical AI to the structural factors that determine who benefits from it. A natural extension that the textbook leaves to the reader is the role of international organizations and transnational cooperation in reducing this inequality beyond the borders of a single country. In the Georgian edition, the three-dimensional framework of digital poverty is fully retained, while the comparison between the Tbilisi hospital and the rural clinic and the separate subsection on equity in Georgian healthcare are found in the English edition; in the Georgian edition, this theme is partly echoed in the concluding subsection of the seventh module (see 8.7).

Privacy, Content Pollution and Disinformation

The textbook discusses training data obtained without consent and its conflict with the right to erasure under the GDPR, and names the striking problem of content pollution, whereby AI-generated text increasingly contaminates the corpora on which future models will be trained. It treats medical confidentiality with due seriousness and warns against the illusion that anonymization can ever be complete. The discussion of disinformation and deepfakes extends to the "stochastic parrot" critique – the argument that a large language model assembles learned patterns without true understanding – which the book uses successfully to dispel the mystique surrounding the models' fluent speech. The environmental section offers memorable figures: a single training run of GPT-3 is compared to the annual carbon emissions of about five hundred cars, and, according to the book, a single ChatGPT query consumes about ten times more energy than an ordinary web search. These numbers show vividly that a technology presented as "clean" is anything but clean in physical terms. At the same time, it is precisely these empirical claims that need to be linked to additional sources, which defines the limits of their use for readers who wish to cite or update them. The comparison with the companion article shows why this matters: citing the original estimate [20], the article puts

the emissions of a single GPT-3 training run at about 550 tonnes of CO_2 equivalent – comparable to the annual emissions of somewhat more than one hundred cars rather than five hundred (see 9.2). This issue is discussed further in Section 11.

Principles of Responsible use and the DOCTOR Checklist

The module ends with core principles – human oversight, transparency, fairness, privacy, accountability – and, most usefully, with a memorable verification procedure. The DOCTOR checklist (Doubt – question the answer; Original source – find the primary source; Check guidelines; Talk to colleagues; Observe the patient – which AI cannot do; Reassess – re-evaluate the clinical picture) turns responsible use into a procedure the reader can apply even under time pressure. Its merit is that it turns an abstract principle into an everyday habit at the bedside. Compared with the WHO, UNESCO and EU frameworks [5,31,35], the principles are broadly consistent but presented in a less formalized way. The DOCTOR checklist, by contrast, is a genuinely original teaching tool that these higher-level documents do not offer.

AI and the Future of the Health Professions

The sixth module moves from the question of what AI is to the question of how it will affect the professions that use it. This is naturally the most debatable chapter: the textbook is most convincing when it reasons from history, while forecasts tied to specific dates need to be checked and updated over time.

The Labour Market and Medical Practice

Drawing on historical analogies (the printing press, the automated loom, the ATM), the module argues that technologies generally transform professions rather than simply eliminate them [16]. Applying this to medicine, the textbook predicts the strongest early impact on image- and signal-based specialties (radiology, pathology, ECG-based cardiology) and a smaller one on specialties based on manual procedures and human contact (surgery, psychiatry, paediatrics). It also anticipates the emergence of new roles – the clinical AI specialist, the medical data scientist, the AI ethics specialist – and, more broadly, new hybrid professions such as the prompt engineer and the data curator. The historical analogy is a reasonable teaching device that calls for some caution: the pace and character of the current transformation may not resemble those of earlier transformations, and the formula "transformation, not replacement" may understate the disruption experienced by people in mid-career. The specialty-level forecasts are plausible and consistent with views expressed in the field; at the same time, they are informed conjecture rather than established fact.

Cognitive Risk and the Delegation of Thinking

One of the module's most valuable warnings concerns cognitive atrophy: if a clinician constantly hands reasoning tasks over to the machine, basic skills may gradually erode – a danger the book captures in the vivid image of "AI slave syndrome". The proposed mitigations are sound: AI should be treated as a second opinion rather than the truth, one should always ask "why?", and one should continue developing one's own clinical skills. This is a genuinely important lesson, and the concern is well supported by analogy with skill erosion in other automated fields [18]. In a future edition it would be desirable to link it to the literature on deliberate practice and simulation-based learning, which offers concrete educational measures against exactly this kind of erosion.

Regulation, Collaboration and the Georgian Setting

The module provides a thorough overview of the regulatory landscape – the GDPR, the EU AI Act [5], China's more centralized approach and the US sectoral model – and explains logically why regulation lags behind technology, citing commercial secrecy, technical complexity and the slow pace of law-making. The overview is accurate and concise: the EU AI Act is mainly named, and a detailed analysis of it is an opportunity for future expansion; for policy-oriented readers, the module outlines the general contours of the field, which provides a good basis for further, more detailed study. As a constructive counterweight, the module presents successful examples of human–AI collaboration in medicine, law, education and programming, which supports the book's consistent thesis: the strongest model is the combination of a competent human and a powerful machine, in which the machine is the assistant and the human remains accountable. The book's identity is clearest in its vision for Georgian healthcare. It lists opportunities (remote diagnostics, data analysis, medical education, second opinions) and challenges (data availability, the language barrier for models trained predominantly in English, infrastructure and regulation), and proposes a phased vision for 2027–2037. This section is the book's calling card. No English-language text offers a comparable roadmap for a small post-Soviet health system, and it shows the reader how a global technology can be turned into a local plan. A natural next step, again left to the reader, is regional cooperation across the Caucasus with neighbouring health systems facing similar constraints.

In the Georgian edition, this vision – subsection 6.7 with its phased roadmap – is not presented separately: the opportunities, challenges and recommendations for Georgian healthcare are moved to the end of the seventh module, where they are read after the discussion of specific clinical systems. This decision is logical – the reader meets the Georgian perspective once they already know which tools are being discussed – but by omitting the phased vision, the Georgian edition loses one of the most original elements of the English one. The companion article partly restores this element: it proposes a three-phase roadmap for Georgia (2026–2028, 2028–2032 and 2032–2037), explicitly presented as the authors' framework rather than an evidence-based forecast [8].

Module Seven: A Practical Review of Medical AI Systems

The seventh module of the Georgian edition is the book's most

important innovation and differs from the preceding modules in genre as well. Whereas the first six modules answer the question "What is AI and how does it work?", the seventh module asks the practising physician's question: "Which tool should I use today, for this patient, and where does its reliability end?" The author presents the module as the integrative part of the course, connecting the foundations of machine learning (Module 3), generative AI (Module 4), ethics (Module 5) and regulation (Module 6) to clinical tools. This stated aim is consistently fulfilled in the text: almost every subsection illustrates concepts from earlier modules – the black box, hallucination, bias, automation bias – through the example of a specific product. Two methodological decisions of the module deserve particular mention. First, the text is accompanied by "⚠ Warnings" placed in red boxes which, in the author's words, are the most important part of the module and mark the boundary "where the system's capability ends and the physician's professional responsibility begins". Second, the module states from the outset that the described features, regulatory status and study results reflect the situation in 2026, that no figure should be used for a clinical decision without checking the manufacturer's current documentation and the regulator's official database, and that mentioning a system is not a recommendation of it. This disclaimer is not a formality but part of the module's message: it teaches the reader that information about medical AI, just like an AI's answer, requires verification.

Structure: Assistive and Autonomous Systems

The module is organized around one simple but decisive question: who makes the final decision? On this basis, the textbook arranges systems into four levels of assistance (Table 7).

Table 7: Levels of assistance of medical AI systems [7], Module 7.

Level	What the System Does	Who Makes the Decision	Examples in the Module
Triage	Flags suspicious cases and moves them up the worklist	The physician – in every case	Aidoc, Nanox.AI
Second reader	Points to a specific suspicious area on the image	The physician, taking the AI's cue into account	Paige Prostate, mammography AI
Decision support	Assesses risk or suggests treatment options	The physician and the multidisciplinary team	Epic's sepsis model, IBM Watson
Autonomous	Itself issues the screening result	The system – within strictly defined limits	IDx-DR (LumineticsCore)

This typology is the module's most valuable analytical tool. It translates the rational-agent framework of the first module into the language of accountability and shows the reader that "AI system" is not a homogeneous category: a triage algorithm that merely reorders a queue and an autonomous screening system that issues a conclusion without a physician involve different levels of risk, regulation and responsibility. The typology is accompanied by the warning "A 'second opinion' does not mean a 'second doctor'", which links automation bias to the "AI slave syndrome" of Module 6 – a good example of the organic binding of earlier and new material. This phenomenon has an extensive empirical basis in the literature [18], a brief reference to which would make the text more robust. Methodologically no less important is the insert "How to read the numbers: sensitivity and specificity". It explains both measures in everyday language and, crucially, warns that in a rare disease false alarms may outnumber true cases even when specificity is high. In this way the Georgian edition partly fills the gap noted with regard to the common modules (see 4.3): the reader now encounters the basic concepts of validation in the context of specific systems. The description of each system follows a single template: brief information on availability (institutional or individual access, pricing model, integration and training requirements), followed by "What it does", "How it works in practice", "Real-world data" and "Significance", an analogy and a warning. The template makes systems easy to compare, and the availability information is a rare and practically valuable detail: the reader learns that most systems are deployed at the institutional level and are not the choice of an individual physician. At the same time, in many places "What it does" and "How it works in practice" repeat the preceding introductory paragraph almost word for word, and the analogy is often duplicated in the "Significance" item as well. Editorial tightening would reduce the module's length without any loss of content.

Diagnostic Systems: Accuracy, Intended use and Real-World Practice

The diagnostic section is the module's most extensive component and covers thirteen systems: in radiology – Aidoc, Nanox.AI (formerly Zebra Medical Vision) and Google's mammography AI; among multi-disease platforms – a2z-Unified-Triage and Annalise Enterprise CXR; in image reconstruction – SwiftMR and Subtle-HD-CT; in pathology – Paige Prostate and PathAI; in ophthalmology – IDx-DR (LumineticsCore); and in cardiology – Cardiologs, AliveCor Kardia 12L and 4TAVR. The selection is representative: it reflects the fact that the great majority of authorized AI devices concern medical imaging, while also illustrating the different logics of triage, second reading, reconstruction and autonomous screening. The subsection ends with a summary table matching each system to its field, task and level of autonomy. The main merit of this section is that it separates results declared by manufacturers from those observed in real-world practice. Using the example of Aidoc,

the reader sees that for a system that showed an average sensitivity of 97% and specificity of 98% in the study submitted to the FDA, only 61% of positive alerts turned out to be true when 14,707 examinations at three hospitals were analysed. In the case of IDx-DR [1,10], the reader learns that despite high sensitivity on images of sufficient quality, images could not be obtained for 10.5% of patients, and the system judged the images of a further 26.1% unsuitable for analysis. These examples teach the logic of validation better than any abstract definition: accuracy is a joint property not of the model alone but of the model, the conditions of data collection and the population.

The discussion of image reconstruction is conceptually original. The warning "A 'restored' detail may not be real" carries the concept of hallucination from Module 4 from text into the visual domain: a reconstruction algorithm can "erase" a small lesion or "draw" a non-existent structure. This is a risk rarely mentioned in the introductory literature, and knowing about it is of direct practical relevance to a radiology student. Similarly, the concept of a system's intended use – an age restriction in the case of a2z, a specific organ in the case of Paige, a single disease in the case of IDx-DR – recurs consistently throughout the module and builds in the reader the habit of not treating a "negative" AI result as ruling out all other pathology.

The main area for development in this section is the traceability of evidence. Many quantitative data (e.g. "80% faster", "more than 100 million cases", "more than 250 medical practices") rely on manufacturer information and, although the text often correctly notes "according to the manufacturer", the specific source and date are not given. Several links need technical correction (e.g. "cariologs" instead of "cardiologs"), and citing another company's website for Google's mammography AI requires a brief explanation. Since the module itself teaches the reader to distinguish manufacturer evidence from independent evidence, consistently applying this principle to every figure – for example, with a note of the form "manufacturer / independent study, year" – would further increase the module's credibility. The companion article [8] has since done much of this work for the key systems – and in doing so it shows that some figures need not only a source but also clarification or correction (see 9.2).

Clinical Decision Support: Failure as a Learning Resource

The subsection on Clinical Decision Support Systems (CDSS) deliberately focuses on failures and limitations through three examples: IBM Watson for Oncology, Epic's sepsis model, and the predictive models of Google and DeepMind. The textbook notes that the history of CDSS is filled more with lessons than with successes, and this position is fully justified pedagogically. The Watson case shows what happens to a system trained largely on hypothetical cases and the practice of a single centre; the independent validation of Epic's model [34] revealed that the widely deployed model missed a large proportion of sepsis cases while generating many false alarms; and DeepMind's acute kidney injury model [27], developed predominantly on data from men, performed less well in women. This subsection naturally continues the MYCIN lesson of Module 2 and thereby reinforces one of the book's strongest theses: clinical value depends on integration, validation and trust, not only on technical capability. In addition, the module introduces the reader to two new, practically important concepts – alert fatigue and dataset shift – and links the latter directly to the Georgian clinic: before a foreign model is used on Georgian patients, it must be established how far its training population resembles the local one in age, sex, disease prevalence and the characteristics of the health system. This is precisely the analytical skill whose development was only implied in general terms in the common modules.

LLM-Based Tools and the Regulatory "Grey Zone"

The discussion of four tools – OpenEvidence, DoxGPT, ChatGPT for Clinicians and iatroX – is the module's most topical and, at the same time, fastest-ageing part. Its analytical core is the subsection "There is no LLM on the FDA list": according to the textbook, of the 1,451 AI medical devices authorized by the end of 2025, about three quarters were in radiology, and as of February 2026 none of them used generative AI. The conclusion drawn from this – that physicians most often use precisely those AI tools that lack regulatory authorization – is a short, memorable and practically highly important formulation that shows the reader where responsibility lies: if an LLM's advice harms a patient, the argument "the system was wrong" does not exist. The warnings here are also well targeted: citing a source does not mean the answer is correct; funding through pharmaceutical advertising is a potential conflict of interest; a comparative study conducted by the company is not independent validation; and a tool built on US guidelines does not take into account Georgian protocols or the local registration of medicines. The iatroX example is particularly valuable: it establishes the principle that a clinical AI tool must conform to the medical system of the country in which it is used, and presents the absence of a tool based on Georgian guidelines as both a challenge and an opportunity for Georgian physicians and developers. The empirical claims in this section (usage percentages, monthly query volumes, results of companies' internal studies) particularly need sources and dates, because they are what changes fastest.

Personalized Medicine, Documentation and Research

The discussion of AlphaFold and Foundation Medicine/OncoKB establishes two important conceptual distinctions. The first is between a scientific and a clinical tool: AlphaFold's prediction of a protein structure [13] is a hypothesis accompanied by its own confidence estimate and only the first link in the drug-development chain. The second is between the "AI" label and the actual technology: the textbook rightly notes that genomic profiling tools combine laboratory technology, algorithmic analysis and an expert-curated knowledge base, and are not always an "AI system" in the narrow sense. This observation brings the central idea of the first module –

that intelligence is a graded rather than a binary concept – back into a practical context. The subsection on documentation and research (AI scribes – Abridge, Nuance DAX, Heidi, Tortus, Tandem; general language models for patient instructions; generative AI in research) turns the principles of Module 4 into clinical routine. The analogy according to which the physician signs the record and, by signing, answers for every word in it accurately captures the main risk of AI scribes – the temptation to approve an "almost finished" note without reading it. It is also notable that the module reads the empirical results critically: time spent on documentation decreases, but most physicians spend the time saved on other work, meaning that the technology does not automatically reduce workload. This section is also the first place where something identified as a natural direction for expansion of the common modules (see 5.3) appears: the Georgian legal framework for data protection – the Law on Personal Data Protection, under which health data constitute a special category of data. It is complemented by the recommendations of the International Committee of Medical Journal Editors [11] on AI authorship and disclosure of its use, and by a warning that models are more likely to err in Georgian-language texts and that the accuracy of AI scribes in Georgian has not been independently tested.

Selection Criteria and Everyday Rules

The concluding, normative part of the module proposes eight criteria for selecting a system: clinical validation; data quality and bias; transparency; integration; regulation; price and availability; training and post-deployment monitoring; and accountability. The criteria emphasize that a decision should not rest on a single measure such as accuracy. The regulatory details are accurate and well explained: the 510(k) pathway, CE marking under the EU Medical Device Regulation (MDR), the high-risk status of medical AI under the EU AI Act [5], and – as a guard against a common error – a note that the EMA mainly regulates medicinal products rather than medical devices. The warning "Authorization does not mean proven clinical benefit" and the recommendation to test a system on local data in a "silent" pilot mode address exactly those aspects of implementation that were considered advisable additions to the common modules. The eight criteria are complemented by the physician's five everyday rules – start with one, always verify, document, remember regulatory status, protect data – and "the rule that stands above all the others": if an AI recommendation contradicts the clinical picture, a guideline or professional judgement, the physician must stop and must not "fit" the patient to the AI's answer. The book thus builds the practice of responsible use on two levels: the DOCTOR checklist of Module 5 works at the level of the individual clinical decision, and the criteria of Module 7 at the level of institutional selection and deployment. Such a two-level structure is rare in introductory educational literature on medical AI and can be regarded as one of the main achievements of the Georgian edition (Table 8).

Table 8: Typology of the warnings in Module 7 (systematized by the authors of this review).

Type of Warning	Examples in Module 7	What it Teaches the Reader
Limits of intended use	Age restriction (a2z), one organ (Paige), one disease (IDx-DR)	Use the system only for what it is authorized for
Interpreting a "negative" result	Triage only orders the queue; "no finding" is not normal	The AI's silence does not remove clinical suspicion
Manufacturer vs independent evidence	Aidoc, Epic, DoxGPT; 510(k) and randomized trials	Declared accuracy and clinical benefit are different things
Local adaptation	Dataset shift; US guidelines; the Georgian language	A model needs local testing
Accountability	Signature on an AI-scribe record; the "grey zone" of LLMs	Final responsibility rests with the physician
Data protection	Patient data in a general chatbot; storage of audio recordings	Health data are a special category

Medical AI in Georgia

The module ends with an overview of the opportunities (remote diagnostics, analysis of clinical data, medical education, second-opinion services) and challenges (data availability, the language barrier, infrastructure, the regulatory framework) of Georgian healthcare and four recommendations: creating an anonymized Georgian medical database, introducing AI courses in medical faculties, public–private partnership for infrastructure development, and developing an ethical and legal framework. The recommendations are sensible and consistent with the spirit of the common modules, but they are written in a more general register than the rest of the module. The subsection would gain analytical depth if it applied the tools of Module 7 directly to Georgian reality – for example, by applying the eight criteria to a Georgian clinic's purchase of a system, by discussing which national institution evaluates medical devices in Georgia and by what procedure, or by describing existing cases of AI use in Georgian clinics. The companion article [8] takes exactly this step: it draws on the State Audit Office's digital-health audit [26] and on regional health-system indicators, and applies the logic of Module 7 to a hypothetical screening programme for diabetic retinopathy in Kakheti (see 9.1).

Assessment of Module Seven

The seventh module adds a qualitatively new dimension to the

Georgian edition: it enriches an introductory textbook with elements of a clinical reference and, importantly, already responds to several of the development directions identified in this review with regard to the common modules. This relationship is summarized below (Table 9).

Table 9: Module 7 and the recommendations concerning the common modules.

Recommendation for the common modules	How Module 7 responds	Remaining scope
Extending validation metrics	Sensitivity, specificity and the role of disease prevalence are explained; manufacturer and real-world results are compared (Aidoc, IDx-DR, Epic)	Calibration, AUROC and predictive values are not discussed systematically
Discussion of implementation	Integration requirements for each system; "silent" pilot; post-deployment monitoring	Examples of change management and workflow redesign
Error scenarios and accountability	System of warnings; the "grey zone" of LLMs; responsibility attached to the signature	Liability for harm under Georgian law
Deepening regulation	510(k), CE/MDR, the high-risk category of the AI Act, clarification of the EMA's role	The national regulatory procedure
Georgian data protection framework	Law on Personal Data Protection; health data as a special category	A clinical scenario on consent and storage rules
Bias mitigation	Dataset shift; comparison of training and local populations; the DeepMind example	Subgroup reporting and fairness metrics

At the same time, the module has areas of development of its own. First, its factual fabric – authorization dates, usage statistics, company mergers and product names – is the most rapidly changing material in the entire book. The author's disclaimer rightly acknowledges this risk; in future editions, however, every quantitative claim should be accompanied by a source, a date and the type of evidence (manufacturer, independent retrospective or prospective study), and the module should be updated regularly between editions, perhaps as an online supplement. Second, the introduction to the module announces three real clinical cases (Aidoc, DeepFAN, IDx-DR), a situational case, exercises and discussion topics, and the text refers to Case 7.3 (subsection 7.8.3) in connection with IDx-DR; in the version reviewed, this material is not present, and DeepFAN appears only in the FAQ. Since it is cases that turn knowledge into skill, adding them would fully unlock the module's pedagogical potential. Third, editorial tightening is desirable (repeated paragraphs, spelling inaccuracies in links and terms), as is a short quiz at the end of the module, by analogy with the English edition of the common modules. The companion article already develops the DeepFAN study as the "human + AI" paradox [8,37], together with real-world data on Aidoc [24] and IDx-DR [10], and thus provides ready-made material for these cases.

The Companion Article: Testing the Textbook Against Independent Evidence

In October 2026, the textbook's author and the first author of this review published a narrative review, Artificial Intelligence in Medicine: Opportunities, Risks, and Prospects for Georgia [8], in this journal. The article uses the textbook as its conceptual framework but, by its own account, bases empirical claims wherever possible on independent peer-reviewed sources. Its methods are reported transparently: a targeted search of PubMed/MEDLINE and Google Scholar (April–September 2026) yielded about 50 sources after three-stage screening; regulatory and institutional documents – the FDA list of AI-enabled devices, EU regulations, and reports of the State Audit Office of Georgia, the European Observatory on Health Systems and Policies and WHO – were added purposively; and specific systems were selected according to their regulatory status, the availability of clinical data and their illustrative value. For the most important clinical claims, the article thus performs exactly the operation that this review recommends for the textbook as a whole: it separates the conceptual framework, which the textbook supplies, from the evidence, which must be traceable to primary sources. Because the first author of this review co-authored the article, the comparison below is confined to verifiable differences in content and sourcing; it does not attempt to assess the article's overall quality.

Directions for Development the Article has Already Addressed

Many of the directions for development identified in Sections 3–8 are addressed in the article (Table 10). Its contribution is greatest precisely where the textbook was weakest: the traceability of figures, the completeness of Module 7 and the application of its tools to Georgian reality.

The Kakheti scenario deserves particular mention because it answers the main criticism made of subsection 8.7. Instead of general recommendations, the article takes one autonomous system (IDx-DR), one region in which hospital-bed density and outpatient contacts are several times lower than in Tbilisi, and the real-world data of an independent study, and works through the consequences: if the proportions observed in Karlsburg [10] were reproduced, roughly one in three screened patients would still need referral to an ophthalmologist. From this the article derives the preconditions for implementation – equipment and a trained operator, internet

access and EHR integration, a "silent" pilot, a referral and treatment pathway, legal clarity on liability, and public financing. This is the selection logic of Module 7 applied to a concrete Georgian decision, and it would make an excellent situational case for the Georgian edition.

Table 10: How the companion article [8] responds to the directions for development identified in this review.

Direction for Development	Response in the Companion Article	Still To Be Done in the Textbook
Linking figures to sources	Key figures linked to FDA 510(k) records, peer-reviewed studies and regulatory documents: Aidoc [24,30], IDx-DR [1,10], mammography AI [15,32], Epic's sepsis model [34], AI scribes [14]	Carry sources, dates and the type of evidence over into Module 7
Validation metrics	Positive and negative predictive values; AUC with the caution that it measures discrimination only, while calibration and thresholds must also be reported [4]	Worked examples of calibration and predictive values in the technical modules
Missing cases in Module 7	DeepFAN "human + AI" paradox [37]; Aidoc in a real clinical workflow [24]; IDx-DR in real-world screening [10]	Turn them into teaching cases with discussion questions
Georgian depth of Module 7 (8.7)	State Audit Office data on EHR coverage and telemedicine [26]; regional indicators; a diabetic-retinopathy screening scenario for Kakheti	The national procedure for evaluating medical devices
Phased vision for Georgia (7.3)	Three-phase roadmap for 2026–2037, labelled as the authors' framework rather than a forecast	Restore in the Georgian edition; reconcile with the textbook's 2027–2037 vision
Implementation	Three stages: "silent" mode → visible support → broad deployment with monitoring; eight procurement criteria including energy efficiency	A change-management example; reconcile the two sets of eight criteria
Environmental claims	GPT-3 training emissions [20] and energy per ChatGPT query [12] sourced, with a caveat on variability over time	Correct and source the textbook's figures (see 9.2)

Discrepancies to be Harmonized

The comparison also reveals places where the textbook and the article give different figures or explanations (Table 11). Most of these differences are not errors in the strict sense: they arise because the textbook does not identify the version of the system, the study or the date on which a figure rests. They are, however, precisely the kind of untraceable variation against which the textbook warns its readers. The two sets of Aidoc figures are not necessarily contradictory – they may refer to different versions of the algorithm and different real-world studies – but without a version, a source and a date the reader cannot tell. The GPT-3 comparison, by contrast, appears to need correction: at roughly 4.6 tonnes of CO_2 per passenger car per year, 550 tonnes corresponds to somewhat more than one hundred cars, not five hundred. The two explanations of MYCIN's non-adoption, already noted in 4.1, should likewise be reconciled with the article's account. In every case the recommended action is the same: each figure in the textbook should state its source, the type of evidence and the date, so that the reader can do what the DOCTOR checklist asks – find the original source (Table 11).

Table 11: Differences between the textbook [6,7] and the companion article [8].

Topic	Textbook	Companion Article
Aidoc: data submitted to the FDA	Sensitivity 97%, specificity 98%	93.6% and 92.3% (198 cases) [30]; 96.2% and 94.8% for a later version
Aidoc: real-world practice	Only 61% of positive alerts true (14,707 examinations, three hospitals)	Basel, 4,450 patients: sensitivity 87.2%, PPV 70.5%; notification time 73 → 58 min [24]
MYCIN	At or above the level of junior physicians; Georgian timeline attributes non-adoption to the cost of computers	At the level of infectious-disease specialists in a blinded evaluation [36]; never implemented
GPT-3 training emissions	Annual emissions of about five hundred cars	About 550 t CO_2e, comparable to more than one hundred cars [20]
Energy per ChatGPT query	About ten times a web search; no source given	About 2.9 Wh versus 0.3 Wh; estimates vary and are declining [12]
DeepMind kidney injury model	Developed predominantly on data from men	Trained on data from US veterans, most of them men [27]

What the Article Cannot Replace

The article complements the textbook but does not replace it. It is a narrative review rather than a teaching text: it contains no analogies, quizzes, cases with questions or graded learning objectives, and it covers the conceptual modules only briefly. It also states its own limitations openly – screening was not carried out by two independent reviewers, the Georgian data rely partly on international reports, the Kakheti scenario is constructed rather than observed, and the roadmap is a proposal rather than a forecast. The most productive relationship between the two texts is therefore a division of labour: the textbook teaches concepts and habits of critical use, while the article – and its future updates – serves as the citable evidence layer and as a model for the regularly updated online supplement to Module 7 proposed in 8.8. Used together, they also offer students a rare learning opportunity: to compare how the same claim appears in a textbook and in a peer-reviewed article, and to understand why the difference matters.

Comparative Positioning in the Field

Gagnidze's textbook occupies a distinctive niche among the works likely to appear beside it on a reading list. Compared with Topol's Deep Medicine [28], it is more explicitly educational and more accessible: Topol writes for a broad audience and interweaves clinical cases, interviews and personal reflection, whereas Gagnidze writes for students and organizes the content around learning objectives, quizzes and cases. Topol discusses specific clinical applications in greater depth in ophthalmology, cardiology and oncology, while Gagnidze covers foundational concepts more broadly: the four faces of AI, the rational-agent model, the three types of machine learning. Both authors devote considerable attention to ethics, but Gagnidze gives considerably more space to algorithmic bias and digital poverty and does so from inside a health system that English-language texts rarely address. Against Russell and Norvig's Artificial Intelligence: A Modern Approach [22], the difference lies in depth rather than direction: both textbooks share a common organizing terminology, but Russell and Norvig's book is a rigorous, mathematically demanding reference work for computer science, whereas Gagnidze's textbook is deliberately pitched at a lower technical level for readers who will probably never write algorithms themselves but must learn when to trust AI and when to question its output. Compared with the policy register of the WHO's Ethics and Governance of Artificial Intelligence for Health [35], Gagnidze treats governance mechanisms more concisely but is far more useful at the bedside, since he turns high-level principles into the practical DOCTOR routine. The comparison that matters most to a curriculum committee is summarized below (Table 12).

Table 12: Comparative positioning of the textbook in the field.

Dimension	Gagnidze [6,7]	Topol [28]	Russell, Norvig [22]	WHO [35]
Main audience	Medical students, clinicians	General audience	Computer science students, engineers	Policymakers, regulators
Technical depth	Introductory	Intermediate, clinical	Advanced, mathematical	Low; governance-oriented
Medical grounding	High – medical examples in every module	High – on applied aspects	Low – general AI	High – at policy level
Specific clinical systems	High – Module 7, more than 20 systems	High – narrative cases	None	Selective
Coverage of ethics	Integrated throughout the text	Substantial	Selective	Comprehensive
Regional context	Explicit (Georgia)	US-centred	None	Global
Pedagogical apparatus	Strong (quizzes, cases, FAQ, warnings)	None	Strong (exercises)	None

The table presents a consistent picture: no other existing text combines an introductory technical level, strong medical grounding, integrated ethics, a well-developed pedagogical apparatus and an explicit non-Western regional context. It is precisely this combination that constitutes Gagnidze's contribution. The seventh module of the Georgian edition adds one more element – a critical review of clinical systems actually in use – and brings the book closer to the applied depth of Topol's Deep Medicine, but in a systematic, comparative and warning-based form rather than a narrative one. At the same time, quantitative validation methods, the organizational side of implementation and national regulatory analysis still need further elaboration in all the books discussed, including Gagnidze's – a shared area of development on which work would benefit the field as a whole.

Review and Prospects for Development

The aim of a review is not only description but also reasoned evaluation. In this respect, the textbook, in both editions, can be regarded as a valuable, deliberately introductory resource: its main strengths lie in pedagogical accessibility and, in the Georgian edition, also in its direct connection to clinical practice, while its opportunities for development stem largely from the deliberate

choice made in favour of accessibility and from the asymmetry between the editions. Where the companion article has already addressed a direction for development, this is noted below; the corresponding changes, however, still need to be carried over into the textbook itself.

Strengths

Table 13: Main strengths of the textbook.

Strength	Evidence in the Text	Significance
Accessible language	Analogies (cook, child, librarian); no jargon	Suitable for readers without technical training
Medical focus	Medical examples and clinical cases in every module	Directly relevant to health professionals
Transition to clinical systems	Module 7: more than 20 systems, a typology of assistance levels (Georgian edition)	Turns knowledge into the skill of working with clinical tools
System of warnings	A "⚠ Warning" for every system: intended use, negative result, accountability	Teaches readers to recognize the limits of reliability
Balanced treatment	Opportunities and risks weighed; hype avoided	Produces critical rather than naive readers
Practical tools	DOCTOR checklist (individual level); 8 selection criteria (institutional level); prompt examples	Directly usable both at the bedside and in clinic-level decisions
Focus on equity	Digital poverty; bias; dataset shift; Georgian cases	Addresses the social determinants of the benefits of AI
Structured pedagogy	Quizzes, cases, tests, exercises (Eng.); FAQ (Geo.); discussion topics (both)	Supports active and case-based learning
Historical depth	AI winters; MYCIN; Moravec's paradox; lessons of Watson and Epic	Places current developments in historical context
Integrated ethics	Ethics embedded in every module	Reflects the real intertwining of ethics and practice

Opportunities for Development

Table 14: Opportunities for development of the textbook.

Opportunity For Development	Evidence in the Text	Significance
Linking empirical claims to sources	Empirical claims in the common modules and quantitative data in Module 7 need a source and a date	Will increase scientific credibility and traceability
Deepening technical issues	Validation, overfitting and generalization are discussed briefly in the common modules; Module 7 partly fills the gap	Adding calibration and AUROC will strengthen the ability to appraise models critically
Expanding the bibliography	A short bibliography with a disclaimer that it was compiled with AI assistance	An expanded bibliography will support guided further reading
Deepening regulatory analysis	The AI Act is named in the common modules; Module 7 explains 510(k), the MDR and the risk category; the national framework is not discussed	Useful for policy-oriented and Georgian readers
Harmonizing the pedagogical apparatus of the editions	The Georgian edition lacks quizzes, cases for Modules 1-5, tests and exercises; the English edition lacks Module 7 and the FAQ	Makes both editions complete and interchangeable resources
Completing Module 7	The announced cases, exercises and discussion topics are not in the text	Cases turn knowledge into clinical skill
Regular updating of Module 7	Data reflect the situation in 2026	An online supplement would keep it up to date
Empirical evaluation of learning outcomes	No data on learning outcomes are yet available	Empirical data would confirm pedagogical effectiveness

Recommendations for Future Editions

a) Sources should be added for every empirical claim and, in Module 7, the date and type of evidence (manufacturer or independent study) should accompany each source; all bibliographic entries should be verified independently, separately from the AI-assisted draft. The sources assembled in the companion article [8] provide a ready starting point.

b) Building on the line begun in Module 7, the discussion of validation should be expanded – sensitivity, specificity, calibration, AUROC, dataset shift and prospective evaluation – so that read-

ers can critically appraise claims about model performance.

c) Explainable AI techniques (SHAP, LIME, attention maps) should not only be mentioned but also explained in the technical modules.

d) It would be useful to link the implementation elements already proposed in Module 7 (the "silent" pilot, post-deployment monitoring) to the MYCIN lesson and to supplement them with a practical example of change management.

e) Error scenarios – what happens clinically and legally when a system makes a mistake – should also be discussed with reference to the Georgian legal and regulatory environment.

f) The editions should be harmonized: Module 7 and the FAQ should be added to the English edition, and the Georgian edition should gain the quizzes, module cases, pre- and post-tests, practical exercises, Georgian-context inserts (including the 2027–2037 vision) and "A Dialogue Across Time".

g) The announced cases (Aidoc, DeepFAN, IDx-DR), a situational case, exercises, discussion topics and a short quiz should be added to Module 7, and internal cross-references (e.g. "Case 7.3") should be aligned with the text.

h) A bilingual (Georgian–English) glossary of terms and an index should be added and, given the pace of development in the field, regular updating of the textbook, and of Module 7 in particular, should be ensured.

i) The textbook should be harmonized with the companion article [8]: discrepant figures (Aidoc's performance, the GPT-3 emissions comparison, the explanation of MYCIN's non-adoption) should be corrected or clarified; the Georgian health-system data and the Kakheti scenario should be added to Module 7 of the Georgian edition; and the article's eight procurement criteria should be reconciled with the eight selection criteria of Module 7.

Implications for Medical Education

The practical value of a review lies in showing how the textbook can be used. The recommendations below rest on a realistic premise: in an overloaded curriculum, AI literacy should find its own place and merge with existing clinical teaching rather than replace it.

AI literacy as a Core Competence

The book's thesis that AI literacy is no longer optional for future physicians should be reflected at the curriculum level [19,33]. Literacy here means not the ability to build models but understanding what AI is, critically evaluating its outputs, using it ethically and communicating about it with patients and colleagues. In keeping with the book's integrated-ethics approach, this literacy is best distributed across the whole curriculum and embedded in clinical rotations.

Balancing Technical and Clinical Skills

The textbook's main message – that AI augments rather than replaces clinical skill – has direct consequences for the curriculum. Physical examination, history-taking, communication and empathy remain irreplaceable and should not be pushed into the background. To these are added data interpretation, critical appraisal of machine outputs and disciplined prompting. The book's warning about cognitive atrophy should also be reflected in assessment design: it is important that students demonstrate clinical reasoning both without the tool and with its competent use.

Pedagogy, Assessment and Faculty Development

The book's case-based structure is consistent with established principles of adult and case-based learning and offers ready-made teaching material. It can be extended through AI-enhanced simulation, structured debates on the ethical dilemmas raised in the book, and practical exercises in prompting and fact-checking [9]. Assessment should follow the same logic: rather than testing the memorization of definitions, examinations should require students to evaluate AI-generated clinical content, detect a deliberately embedded hallucination and reason through an ethical case – precisely the skills developed by the book's practical exercises. Faculty development is especially important for the success of these measures, since most current teachers received their professional training before these tools appeared; interdisciplinary collaboration between medicine, computer science and ethics is an important condition for credible teaching of the subject.

A concrete Model of Integration

One workable model of integration treats the textbook as the core reading for a short compulsory AI-literacy module in the preclinical years, deliberately supplemented with a few peer-reviewed sources – the companion article [8], which links the textbook's clinical claims to primary sources, being a natural first choice – both to provide additional sourcing and to model scholarly practice. The module is then linked to the clinical years through exercises embedded in rotations – for example, evaluating a real decision-support tool in radiology or auditing an AI-generated discharge summary – with assessment by practical demonstration rather than written reproduction of knowledge. Used in this way, even the book's areas for development become learning opportunities: the need to link sources becomes a source-verification exercise, and the disclaimer about AI assistance becomes a live example of exactly the critical appraisal the module exists to build.

Two Editions in the Curriculum

The existence of two editions creates additional opportunities for the curriculum. The seventh module of the Georgian edition is a natural bridge between the preclinical module and clinical rotations: its eight criteria can be used as a structured exercise in which a group of students evaluates a specific system – for example, a triage tool under consideration at a teaching clinic – and justifies its

evaluation in a written report, while the typology of assistance levels and the warnings can become a shared language for discussing AI outputs during rotations. The English edition, for its part, suits English-language programmes and international students, and its quizzes, cases and tests can also be used as supplementary material in Georgian-language groups. Until the editions are harmonized, the optimal solution is to use them in combination: the pedagogical apparatus of the English edition for the common modules, and the seventh module of the Georgian edition for the clinical stage.

Conclusion

A. Gagnidze's textbook, in both editions, fulfils its main aim: it offers readers without technical training an accessible introduction to AI grounded in the medical context. Across the six common modules it consistently covers definitions of and approaches to AI, the history of the field, the foundations of machine learning, the mechanisms and risks of generative AI, ethical and social issues, and the future of the health professions, while the seventh module of the Georgian edition connects this knowledge to specific systems already deployed in clinics. In evaluative terms, the strengths of the common modules and their room for improvement are two sides of the same editorial choice: readability, accessibility, a balanced approach and easy-to-use practical tools make it a strong introductory resource, while the depth of citation, the technical details of model validation and the analytical treatment of regulation are natural directions for future development. The seventh module already responds to some of these directions – through concrete discussion of the main validation measures, implementation, accountability and regulation – and creates an original teaching tool in its system of warnings. Its main challenge is the transience of its information and the traceability of its evidence, while the main challenge of the two editions is the asymmetry of their pedagogical apparatus.

The companion article shows that these challenges are manageable: once the key clinical claims were linked to independent and regulatory sources, the textbook's central arguments were confirmed, while a small number of figures proved to need clarification or correction – a result that itself illustrates the habit of verification the textbook teaches. For medical education, the conclusion is clear. AI literacy is already essential for future health professionals, and both editions provide a solid foundation that, with future development – peer-reviewed literature, up-to-date regulatory analysis and integration with practical assessment – will further strengthen medical education. Their integration should be planned in advance, draw on the experience of clinical rotations, and be supported by faculty development and interdisciplinary collaboration. The book's concluding message is sound: AI will not replace the physician, but a physician who uses AI well will outperform, in many tasks, one who does not use artificial intelligence. The seventh module adds an important refinement to this message: using AI "well" means knowing what the system is authorized for, what its silence means and who answers for its error. Human oversight, empathy and critical judgement remain irreplaceable throughout. A. Gagnidze's textbook – in both its English and Georgian editions – will therefore make a genuine contribution to preparing the health professions for the age of artificial intelligence. In future, updated editions, the author would do well to align the two versions, enrich them with an empirical basis and complement them with additional accompanying resources, such as a regularly updated online supplement to Module 7, a case database and online modules. The companion article [8] offers a model for such a supplement.

Declarations

Conflict of interest. Mari Manjgaladze is a co-author, with the author of the reviewed textbooks, of the companion article discussed in Section 9 [8]. To limit the effect of this relationship, the comparison in Section 9 is confined to verifiable differences in content and sourcing. No other conflict of interest is declared. Statement on the use of AI (ICMJE). A generative AI tool (Claude, Anthropic) was used to assist with structuring and language editing of the manuscript. The authors checked the content, sources and conclusions and accept full responsibility for them [11].

Acknowledgment

This work was prepared with the support of East European University, Tbilisi, Georgia. The authors thank Professor Avtandil Gagnidze for his extensive and fruitful collaboration.

Conflict of Interest

None.

References

1. Abràmoff MD, Lavin PT, Birch M, Shah N, Folk JC, et al. (2018) Pivotal trial of an autonomous AI-based diagnostic system for detection of diabetic retinopathy in primary care offices. NPJ Digit Med 1: 39.
2. Beam AL, Kohane IS (2018) Big data and machine learning in health care. JAMA 319(13): 1317-1318.
3. Char DS, Shah NH, Magnus D (2018) Implementing machine learning in health care – addressing ethical challenges. N Engl J Med 378(11): 981-983.
4. Collins GS, Reitsma JB, Altman DG, Moons KGM (2015) Transparent reporting of a multivariable prediction model for individual prognosis or diagnosis (TRIPOD): the TRIPOD statement. BMJ 350: g7594.
5. European Union (2024) Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Off J Eur Union.
6. Gagnidze A (2026) Introduction to Artificial Intelligence for Health Professions. East-West University; East European University, Tbilisi. ISBN 978-9941-8-9256-1.
7. Gagnidze A (2026) Khelovnuri inteleqtis shesavali medikosebistvis [Introduction to Artificial Intelligence for Medical Professionals; in Georgian]. East-West University, Tbilisi. ISBN 978-9941-8-9385-8.
8. Gagnidze A, Manjgaladze M (2026) Artificial intelligence in medicine: opportunities, risks, and prospects for Georgia. Am J Biomed Sci & Res 32(5): 797-807.

9. Holmes W, Bialik M, Fadel C (2019) Artificial intelligence in education: promises and implications for teaching and learning. Center for Curriculum Redesign.
10. Hunfeld E, Tayar A, Paul S (2026) Real-world performance of the AI diagnostic system IDx-DR in the diagnosis of diabetic retinopathy and its main confounders. Sci Rep 16: 4349.
11. International Committee of Medical Journal Editors (2026) Recommendations for the conduct, reporting, editing, and publication of scholarly work in medical journals (updated January 2026). ICMJE.
12. International Energy Agency (2024) Electricity 2024: analysis and forecast to 2026. IEA, Paris.
13. Jumper J, Evans R, Pritzel A, Tim Green, Michael Figurnov, et al. (2021) Highly accurate protein structure prediction with AlphaFold. Nature 596(7873): 583-589.
14. Lukac PJ, Turner W, Vangala S, Chin AT, Khalili J, et al. (2025) Ambient AI scribes in clinical practice: a randomized trial. NEJM AI 2(12).
15. McKinney SM, Sieniek M, Godbole V, Godwin J, Antropova N, et al. (2020) International evaluation of an AI system for breast cancer screening. Nature 577(7788): 89-94.
16. Meskó B, Hetényi G, Győrffy Z (2018) Will artificial intelligence solve the human resource crisis in healthcare? BMC Health Serv Res 18(1): 545.
17. Obermeyer Z, Powers B, Vogeli C, Mullainathan S (2019) Dissecting racial bias in an algorithm used to manage the health of populations. Science 366(6464): 447-453.
18. Parasuraman R, Manzey DH (2010) Complacency and bias in human use of automation: an attentional integration. Hum Factors 52(3): 381-410.
19. Park SH, Do KH, Kim S, Park JH, Lim YS, et al. (2019) What should medical students know about artificial intelligence in medicine? J Educ Eval Health Prof 16: 18.
20. Patterson D, Gonzalez J, Le Q, Liang C, Munguia LM, et al. (2021) Carbon emissions and large neural network training. arXiv: 2104.10350.
21. Rajkomar A, Dean J, Kohane I (2019) Machine learning in medicine. N Engl J Med 380(14): 1347-1358.
22. Russell S, Norvig P (2021) Artificial intelligence: a modern approach. (4th edn), Pearson.
23. Searle JR (1980) Minds, brains, and programs. Behav Brain Sci 3(3): 417-424.
24. Seyam M, Weikert T, Sauter A, Brehm A, Psychogios MN, et al. (2022) Utilization of artificial intelligence–based intracranial hemorrhage detection on emergent noncontrast CT images in clinical workflow. Radiol Artif Intell 4(2): e210168.
25. Shortliffe EH, Cimino JJ (2021) Biomedical informatics: computer applications in health care and biomedicine. (5th edn), Springer.
26. State Audit Office of Georgia (2024) Digital Health Performance Audit [in Georgian]. State Audit Office of Georgia, Tbilisi.
27. Tomašev N, Glorot X, Rae JW, Michal Zielinski, Harry Askham, et al. (2019) A clinically applicable approach to continuous prediction of future acute kidney injury. Nature 572(7767): 116-119.
28. Topol EJ (2019) Deep medicine: how artificial intelligence can make healthcare human again. Basic Books.
29. Turing AM (1950) Computing machinery and intelligence. Mind 59(236): 433-460.
30. US Food and Drug Administration (2018) 510(k) premarket notification K180647: BriefCase (Aidoc Medical). FDA, Silver Spring, MD.
31. UNESCO (2023) Guidance for generative AI in education and research. United Nations Educational, Scientific and Cultural Organization.
32. Warren LM, Venton J, Young KC, Halling-Brown M, Kelly CJ, et al. (2026) Impact of using artificial intelligence as a second reader in breast screening including arbitration. Nat Cancer 7(3): 507-521.
33. Wartman SA, Combs CD (2018) Medical education must move from the information age to the age of artificial intelligence. Acad Med 93(8): 1107-1109.
34. Wong A, Otles E, Donnelly JP, Andrew Krumm, Jeffrey McCullough, et al. (2021) External validation of a widely implemented proprietary sepsis prediction model in hospitalized patients. JAMA Intern Med 181(8): 1065-1070.
35. World Health Organization (2021) Ethics and governance of artificial intelligence for health. WHO.
36. Yu VL, Fagan LM, Wraith SM, Clancey WJ, Scott AC, et al. (1979) Antimicrobial selection by a computer: a blinded evaluation by infectious diseases experts. JAMA 242(12): 1279-1282.
37. Zhu Z, Ge Hu, Weixiong Tan, Kai Gao, Chao Sun, et al. (2026) DeepFAN, a transformer-based deep learning model for human–artificial intelligence collaborative assessment of incidental pulmonary nodules in CT scans: a multi-reader, multi-case trial. Nat Cancer 7(5):773-789.