



International Journal of Innovative Technologies in Social Science

e-ISSN: 2544-9435

Operating Publisher
SciFormat Publishing Inc.
ISNI: 0000 0005 1449 8214

2734 17 Avenue SW,
Calgary, Alberta, T3E0A7,
Canada
+15878858911
editorial-office@sciformat.ca

ARTICLE TITLE	ARTIFICIAL INTELLIGENCE IN MEDICAL EDUCATION AND CLINICAL PRACTICE: A NARRATIVE REVIEW
----------------------	--

DOI	https://doi.org/10.31435/ijitss.4(52).2026.6261
------------	---

RECEIVED	02 August 2026
-----------------	----------------

ACCEPTED	26 September 2026
-----------------	-------------------

PUBLISHED	05 October 2026
------------------	-----------------

LICENSE	
----------------	--



The article is licensed under a **Creative Commons Attribution 4.0 International License**.

© The author(s) 2026.

This article is published as open access under the Creative Commons Attribution 4.0 International License (CC BY 4.0), allowing the author to retain copyright. The CC BY 4.0 License permits the content to be copied, adapted, displayed, distributed, republished, or reused for any purpose, including adaptation and commercial use, as long as proper attribution is provided.

ARTIFICIAL INTELLIGENCE IN MEDICAL EDUCATION AND CLINICAL PRACTICE: A NARRATIVE REVIEW

Anna Bociąg (Corresponding Author, Email: aniabociag4@gmail.com)

Specialist Hospital in Wejherowo, dr Alojzego Jagalskiego 10, 84-200 Wejherowo, Poland

ORCID ID: 0009-0006-9906-9906

Łukasz Ptach

Specialist Hospital in Wejherowo, dr Alojzego Jagalskiego 10, 84-200 Wejherowo, Poland

ORCID ID: 0009-0009-9906-8075

Rafał Marciniak

Independent Public Multidisciplinary Hospital in Nowa Sól, Chałubińskiego 7, 67-100 Nowa Sól, Poland

ORCID ID: 0000-0003-3123-3985

Konrad Kowalski

University Clinical Centre Gdańsk, Dębinki 7, 80-952 Gdańsk, Poland

ORCID ID: 0009-0006-6490-9653

Aleksandra Milewska

Copernicus Medical Entity Ltd., Nicolaus Copernicus Hospital, Nowe Ogrody 1-6, 80-803 Gdańsk, Poland

ORCID ID: 0009-0004-2301-0751

ABSTRACT

Background: Artificial intelligence is rapidly reshaping medicine, influencing clinicians' training and patient care. Despite growing enthusiasm for AI-enabled diagnostics and digital learning tools, integration of these technologies into medical education and clinical practice remains uneven.

Objective: This narrative review synthesises recent literature on the role of AI in medical education and clinical practice, mapping current applications, stakeholder perceptions, and the conditions that determine successful adoption.

Methods: Relevant articles were identified through structured searches of PubMed, Scopus, and EMBASE, prioritising literature from the last decade. Empirical studies, methodological reviews, and authoritative commentaries addressing AI, machine learning, or digital health technologies in clinical or educational settings were included and organised thematically.

Findings: AI-powered tools are embedded across medical training, spanning augmented and virtual reality simulations, bedside mobile learning, and a range of clinical specialties, with performance matching expert levels on narrowly defined diagnostic tasks. Stakeholder attitudes are positive but conditional: medical students expect digital health technology to feature in their curricula yet many consider their own competencies inadequate, and very few physicians-in-training report familiarity with AI. Patients express willingness to engage with AI for self-management while preferring human contact in sensitive clinical domains. Persistent barriers remain, including limited availability of curated data, algorithmic bias, difficulties reproducing published models, and underdeveloped regulatory and workflow structures.

Conclusion: AI delivers its greatest value when deployed as a collaborative instrument that augments rather than replaces clinical judgment. Realising this potential depends on sustained investment in digital literacy, explainable systems, clinician involvement in design, and equitable institutional infrastructure.

KEYWORDS

Artificial Intelligence, Medical Education, Clinical Practice, Healthcare Technology, Machine Learning

CITATION

Bociąg, A., Ptach, Ł., Marciniak, R., Kowalski, K., & Milewska, A. (2026). Artificial intelligence in medical education and clinical practice: A narrative review. *International Journal of Innovative Technologies in Social Science*, 4(52). [https://doi.org/10.31435/ijitss.4\(52\).2026.6261](https://doi.org/10.31435/ijitss.4(52).2026.6261)

COPYRIGHT

© The author(s) 2026. This article is published as open access under the Creative Commons Attribution 4.0 International License (CC BY 4.0), allowing the author to retain copyright. The CC BY 4.0 License permits the content to be copied, adapted, displayed, distributed, republished, or reused for any purpose, including adaptation and commercial use, as long as proper attribution is provided.

Introduction

Artificial intelligence is rapidly reshaping industries across the globe, and its influence on healthcare represents one of the most significant shifts in modern medicine. At its core, artificial intelligence comprises computer systems designed to perform tasks that typically require human cognition—such as learning, reasoning, and processing complex information—enabling the analysis of clinical data to uncover patterns that might otherwise elude human observation. Within healthcare, these systems streamline a broad spectrum of processes, ranging from accurate disease diagnosis to the creation of personalized treatment plans. Specifically, machine learning algorithms are increasingly utilized to interpret medical images with greater precision, while AI-powered scribing tools help reduce the administrative burden on clinicians by automating documentation. These technologies support diagnostic accuracy, assist in predicting patient risks, and allow for the tailoring of treatment strategies to individual patient needs.

This article examines how these technological advancements are being integrated into medical training and daily clinical practice to improve patient outcomes and quality of life. By synthesizing recent research and practitioner perspectives, this review highlights the essential knowledge and tools necessary for healthcare providers to successfully incorporate digital solutions into patient care workflows while maintaining the human-centered empathy essential for clinical practice. Furthermore, integrating artificial intelligence into curricula is crucial for preparing future physicians to competently utilize these tools. By fostering a foundational understanding of AI technologies early in training, schools can better equip students to navigate the complexities of modern, data-driven healthcare. While the integration of these systems is technically complex, recent studies indicate that stakeholders generally hold positive attitudes toward the potential for AI-driven progress in the field.

Despite this optimism, successful implementation relies on balancing high-tech data processing with the human element of care. The future of healthcare depends on a model where AI's pattern-recognition strengths amplify clinical expertise, while doctors preserve the contextual wisdom and relational presence that machines cannot replicate. Consequently, bridging the gap between innovative algorithmic capabilities and the bedside experience remains a central challenge for medical institutions. To address this, we must critically evaluate both the pedagogical frameworks used to teach these digital competencies and the clinical workflows that currently govern their use. Through this exploration, we look at practical examples of these tools in action and consider how training programs can evolve to keep pace with rapid digital change. Educational institutions must ensure that future professionals possess the digital literacy required to work effectively alongside these advanced systems. This digital literacy is essential for students to effectively integrate technologies into their clinical practice. This shift toward a technologically integrated curriculum is vital, as many current medical students report feeling underprepared to manage these tools in real-world settings. However, formal digital health training in medical schools remains limited, highlighting the need to expand these offerings so that clinical workflows can accommodate emerging resources. By incorporating hands-on experience with AI-supported clinical aids, institutions can empower students to critically interpret algorithm-generated insights while maintaining the essential human judgment required at the bedside. This practical approach not only addresses the current feeling of underpreparedness among students but also fosters the adaptability necessary for physicians to serve as proficient partners to the evolving technologies they will inevitably encounter in practice. Ultimately, the successful adoption of these systems hinges on a collaborative environment where clinicians and developers work together to ensure AI serves as a transparent and reliable support for patient care.

Methodology

This narrative review was undertaken to provide an integrated overview of the current state of artificial intelligence in medical education and clinical practice. Rather than following the rigid protocol of a systematic review, we adopted a narrative approach that is appropriate to a field undergoing rapid technological change, characterised by methodological diversity across studies, and marked by uneven adoption across institutions and specialties.

Search strategy and article selection. Relevant articles were identified through structured searches of major biomedical databases, including PubMed, Scopus, and EMBASE. These database searches were complemented by manual review of citation lists in key narrative and scoping reviews identified during the initial search process. Priority was given to literature published over the last decade to ensure the technological context remained relevant to contemporary digital health standards, although seminal earlier work was retrieved through citation tracking where appropriate. Both empirical studies (including quantitative surveys, cohort analyses, and randomised evaluations) and authoritative expert commentaries were eligible for inclusion.

Themes and synthesis. Articles were organised around four intersecting themes that emerged consistently across the literature: the transformation of medical education and training through digital and AI-enabled tools; the deployment of AI in clinical diagnostics, treatment planning, and dose optimisation; stakeholder attitudes among patients, clinicians, and students toward AI and the evolving doctor–patient relationship; and implementation challenges alongside the institutional, educational, and regulatory enablers that determine successful adoption.

Inclusion criteria. Articles were eligible if they addressed artificial intelligence, machine learning, or digital health technologies in a clearly clinical or educational setting within medicine or allied health professions. Eligible study designs comprised empirical quantitative research, including cross-sectional surveys, cohort analyses, retrospective validation studies, and randomised evaluations of digital learning tools; qualitative investigations such as interview studies, focus groups, and ethnographic work; methodological, scoping, and systematic reviews; and authoritative expert commentaries published in peer-reviewed journals. Articles confined to engineering or computer science audiences without an explicit medical or educational framing were excluded, as were conference abstracts lacking full-text follow-up, non-peer-reviewed white papers, and industry marketing materials. Publications were limited to those in English or accompanied by an English-language abstract, a pragmatic choice reflecting the databases indexed in the primary search while acknowledging that relevant work published in other languages may have been overlooked. Both undergraduate and postgraduate medical education contexts were eligible, alongside hospital-based, primary care, and community clinical settings, in order to capture the educational continuum.

Search terms and screening. Search strings combined controlled vocabulary terms with free-text keywords across three intersecting conceptual domains: artificial intelligence and machine learning; medical education, clinical practice, and clinical decision-making; and stakeholder attitudes, encompassing both patient experience and clinician perception. Although date filters were not strictly imposed, priority was given to literature published from 2015 onward to ensure the technological context remained relevant to contemporary digital health standards. Title and abstract screening were conducted by two reviewers independently, with disagreements resolved through discussion. Full-text review followed for all articles that passed initial screening or for which relevance could not be determined from the abstract alone. Although this narrative review did not register a formal systematic-review protocol, the structured approach produced an evidence base adequate for the thematic synthesis that follows.

Scope and limitations. The review spans undergraduate medical training, postgraduate clinical settings, and emerging point-of-care technologies to capture the educational continuum. Because the rapid and uneven pace of technological deployment makes a deep quantitative meta-analysis impractical, the review highlights general trends and recurring themes rather than pooled effect sizes. Substantial variation in research design, instrumentation, and institutional context further complicates direct comparisons; findings are therefore interpreted as a thematic overview rather than an exhaustive audit of any single emerging platform. Reproducibility challenges that have been consistently raised in the machine learning literature also limit how confidently performance metrics — particularly those from image-based diagnostic studies — can be generalised across populations and institutions. Finally, conclusions drawn from stakeholder surveys must be read with caution, because respondent attitudes can be shaped by how researchers introduce or frame AI applications.

This methodology acknowledges that AI in medicine functions as a complex sociotechnical system, in which outcomes depend as much on human factors and institutional readiness as on the underlying algorithms.

Results

The integration of artificial intelligence and digital health tools is reshaping both how medical students are trained and how patient care is delivered. Our synthesis of recent literature highlights a transition from early experimental applications toward comprehensive, data-driven frameworks that emphasize interactive learning, diagnostic precision, and personalized treatment. The findings are organized around four intersecting themes: the transformation of medical education and training, the deployment of AI tools in clinical diagnostics and treatment, the evolving attitudes of patients and clinicians toward these technologies, and the persistent implementation challenges that determine whether AI can fulfill its promise at scale.

AI in Medical Education and Training

Recent advances in medical pedagogy have increasingly centred on extending the classroom into immersive, technology-rich environments. Augmented and virtual reality (AR/VR) platforms now allow students to rehearse complex clinical scenarios in safe, interactive settings, using immersive platforms to simulate realistic patient interactions without the risks associated with live practice. In one evaluation of AR-based training for operating-room emergencies, two-thirds of clinicians (67%) expressed a preference for AR-based simulation over textbook-based preparation. Tablet-based AR programs have shown learning outcomes comparable to head-mounted VR while causing fewer adverse effects such as nausea, headache, and disorientation. These platforms also align with Mayer's cognitive theory of multimedia learning, enhancing delivery, presentation, and sensory engagement in ways that traditional lecture formats rarely achieve.

Beyond immersive simulations, AI is reshaping the day-to-day fabric of medical training through intelligent tutoring systems, adaptive learning platforms, and the systematic use of digital anatomy resources. Digital anatomy has revolutionized undergraduate education through 3D reconstruction of human structures and is increasingly shifting the paradigm of pre-clinical and vocational training, providing safe, repeatable clinical experiences without the logistical constraints of cadaver-based dissection. Mobile and just-in-time learning have proven especially valuable during clinical placements. Students using personal digital assistants and mobile reference tools reported that bedside access to clinical information reinforced learning in context, with 76% of students in one UK cohort using such devices during placements. The principle of situated learning is reinforced when students access information on their own devices during real clinical encounters, a practice that bridges the gap between classroom theory and on-the-ward decision-making.

Despite these advances, a striking curricular gap remains. Recent surveys indicate that 80% of medical students expect digital health technology to be integrated into their curricula, yet more than 50% rate their own competencies in this area as poor or very poor. Among practising physicians and physicians-in-training, only 6% report being familiar with AI. These findings are echoed in specialty-specific data: among radiology residents, 86% agreed that AI will change and improve their future practice, while 71% felt that AI should be formally taught in medical schools. Schools are therefore under mounting pressure to embed AI literacy, digital fluency, and hands-on experience with algorithm-supported clinical tools into mainstream training, lest the next generation of clinicians graduate into a workplace they have not been adequately prepared to navigate.

AI Tools in Clinical Practice — Diagnostics and Treatment

Outside the classroom, machine learning and deep learning systems are quietly becoming routine partners in clinical decision-making. Across specialties, the evidence converges on a consistent pattern: AI performs at or above human expert levels on narrowly defined diagnostic tasks, while its most enduring value lies in accelerating routine workflows and surfacing patterns that might otherwise be missed.

In radiology, a mammography-based study showed that AI interpretation produced an absolute reduction of 5.7% in false positives and 9.4% in false negatives compared with standard double-reading. Separate work found AI to be 90% sensitive for breast cancer versus 78% for radiologists overall, and 91% versus 74% for early-stage disease. For pneumonia detection on chest radiographs, deep learning algorithms achieved 96% sensitivity and 64% specificity against radiologist performance of 50% and 73%, respectively. Similar momentum is visible in dermatology and pathology, where convolutional neural networks applied to skin lesion images and whole-slide pathology scans match specialist-level accuracy in detecting melanoma and other malignancies. Ophthalmology has likewise benefited, with automated diabetic retinopathy screening tools now deployed in several health systems.

Predictive analytics extends beyond imaging. Cardiology-focused deep learning tools detect EKG abnormalities and stratify cardiovascular risk. Oncology models trained on genomic profiles have predicted

chemotherapy response with over 80% accuracy across multiple standard-of-care regimens in a study of 175 cancer patients . Electronic health record data from 17,556 patients have been used to predict antidepressant class response, opening the door to AI-informed prescribing . In dose optimization, an AI-derived platform dynamically adjusts chemotherapy dosing based on individual patient response and has shown reduced doses alongside improved response rates compared with standard regimens . A separate warfarin maintenance study analysed data from 19,719 inpatients across three institutions and produced individualized dosing recommendations that outperformed expert physicians . AI is also supporting individualized treatment beyond pharmacology, including gene and cell therapies that require precise patient stratification to be delivered safely . Across these applications, the recurring finding is that AI is most effective when treated as a complement to clinician judgment — flagging anomalies, prioritizing urgent cases, and refining rather than replacing diagnostic decisions.

Stakeholder Attitudes and the Doctor–Patient Relationship

Across patients, clinicians, and students, attitudes toward AI are broadly positive but consistently conditional. A systematic synthesis of 27 stakeholder studies found that attitudes were generally favourable, particularly among those with direct experience of AI, but heavily dependent on safeguards: trustworthy outputs, preserved clinician oversight, and protection of the clinician–patient relationship . The same analysis noted that AI applications which automate data interpretation and synthesis are received more favourably than those that directly influence clinical decisions or appear to displace the human element of care .

Patient data reinforce this picture. In a US-based survey, 60% of participants reported discomfort with providers relying on AI for their care, while 80% simultaneously indicated willingness to use AI-powered tools to help manage their health . Comfort varied by clinical context: most patients considered that AI would improve healthcare overall, but expressed pronounced reservations about AI in sensitive domains such as mental health, chronic disease management, and end-of-life decisions, where human contact remains the expected standard . Higher educational attainment and greater familiarity with technology consistently predict more positive attitudes .

Clinician attitudes largely mirror those of patients but with sharper concerns around professional identity. Doctors predominantly support AI as a complement to their practice rather than a substitute, yet many express unease about the erosion of clinical skills, the potential displacement of professional epistemic authority, and loss of decisional control . Practical concerns surface in studies of emerging tools: in qualitative work on AI scribes, doctors reported insufficient knowledge about how the technology works, including what happens to consultation data once captured, and worried about over-reliance . Crucially, trust in AI-generated advice is context-dependent: when AI output and clinician opinion disagree, or when decisions involve serious illness, clinician judgment is expected to take precedence . For medical students, attitudes trend positive but remain tightly bound to the perceived centrality of the doctor–patient relationship, with empathy, trust, and ethical awareness cited repeatedly as the values AI must safeguard rather than replace .

Implementation Challenges and Enablers

The persistence of barriers explains why a technology with so much demonstrated capability has yet to be evenly embedded in everyday practice. A common thread across the reviewed literature is concern about input data: machine learning accuracy depends on the quantity and quality of available data, and many clinical settings face shortages of curated, representative datasets . This is compounded by reproducibility challenges that have been consistently raised in the AI and machine learning literature, making some published findings difficult to validate across institutions and populations . Difficulty reproducing models is especially pronounced in fields such as neurology and psychiatry, where predictive tools are increasingly being explored but patient-level heterogeneity remains high .

Ethical and regulatory gaps create additional friction. Privacy, confidentiality, informed consent, and patient autonomy remain unresolved at the policy level, while algorithmic bias — including the risk that human biases embedded in training data are amplified by AI outputs — threatens to entrench existing health disparities if left unchecked . Clinicians themselves identify the absence of practical, trustworthy regulatory guidance as a barrier to adoption, calling for clear recommendations from medical colleges, government agencies, and insurers to support safe integration into practice .

Workflow integration presents a parallel set of constraints. The promise of digital tools will only be realized if apps and AI tools can be prescribed, monitored, and documented through the same electronic health record systems clinicians already use — a step that requires digital formularies, clinical decision support, and

reimbursement models that have not yet matured. Without these structures, adoption is likely to remain limited to early champions while broader clinician engagement stalls.

The literature also points to a consistent set of enablers. A human-in-the-loop design philosophy, in which AI handles routine interpretation and prediction while clinicians retain final authority, emerges as the dominant recommended model. Meaningful involvement of clinicians in the design, validation, and ongoing oversight of AI tools is associated with stronger acceptance and safer deployment. Explainable AI is increasingly recognized not just as a technical goal but as an ethical requirement: opaque, "black box" models undermine shared decision-making and erode trust. Finally, sustained investment in training and digital literacy — for clinicians, students, and patients alike — is repeatedly cited as the foundation on which the entire integration effort rests. Together, these enablers suggest that the path to routine clinical AI depends less on the technology itself and more on the institutional, regulatory, and educational scaffolding around it.

These four themes set up the Discussion that follows, in which we consider how medical curricula, clinical workflows, and policy frameworks must co-evolve to translate the demonstrated potential of AI into reliable improvements in patient care.

Discussion

AI as a Collaborative Instrument, Not a Substitute

The findings reported in the Results section point to a clear and consistent conclusion: artificial intelligence is best understood not as a replacement for clinicians but as a sophisticated instrument that augments specific dimensions of clinical capability. The reviewed literature identifies three domains in which AI demonstrates clear superiority — speed and scale of data processing, pattern recognition across large imaging datasets, and predictive analytics — alongside three corresponding domains of clinical work that remain irreducibly human, including the therapeutic relationship, ethical judgment, and contextual reasoning in complex and ambiguous cases.

When viewed through this lens, the central question ceases to be whether AI will replace physicians. Instead, the more productive question is how AI can be integrated to amplify the dimensions of care that physicians already provide. The literature reviewed here frames AI as a tool analogous to advanced imaging technology — a powerful instrument that extends the clinician's perceptual and analytical reach without displacing the interpretive and relational role of the human practitioner. This framing aligns with the broader finding, from a systematic synthesis of 27 stakeholder studies, that AI applications which automate data interpretation and synthesis are received more favourably than those that directly influence clinical decisions or appear to displace the human element of care.

A particularly illustrative example is the deployment of AI scribes, which automate clinical documentation that has long contributed to professional burnout. Focus group studies indicate that doctors perceive these tools as a meaningful improvement in efficiency, quality of notes, and even doctor–patient interaction, provided that safeguards around transparency, oversight, and data governance are in place. However, the same studies highlight an underlying tension: both AI users and non-users expressed concern about over-reliance, particularly under time pressure, and worried that automation could erode clinical skills if outputs are accepted uncritically. In one telling example, a user interpreted unfamiliar content in a note as their own omission rather than recognizing it as a potential AI hallucination — illustrating the precise cognitive risk that human-in-the-loop design is meant to mitigate.

Bridging the Education–Practice Divide

The curricular signals reported in the Results are best read as a symptom of a deeper misalignment rather than as a shortfall of curricular content: the speed of clinical digitisation has outpaced the speed at which institutions can reconceptualise what it means to train a physician. Because there is no significant difference in AI knowledge between practising physicians and medical students, exposure during medical school is not, on its own, closing the gap. Postgraduate learning is therefore happening on the job—through experimentation, vendor training, and informal peer exchange.

The radiologist-resident data reinforce this view in a more specialised way. Even within a discipline at the forefront of AI integration, residents simultaneously expected the technology to change their practice and reported insufficient formal training to engage with it. That convergence across general medical education and a heavily image-oriented specialty indicates the deficit is structural rather than domain-bound. The deeper

implication is that AI literacy is unlikely to be built through classroom instruction alone, because the deficit lies in how clinical learning is *organized* rather than in what is *covered*.

The situated-learning evidence is most informative when read against this backdrop. When students access point-of-care information on their own devices during authentic clinical encounters, learning is anchored in the specific problem the patient presents and carries forward into practice because it was acquired in the context where it will be used. Read in light of the curricular shortfall, this means mobile and bedside learning tools are valuable vehicles through which AI fluency can be built at scale before graduation. Medical schools that take the 80% student expectation seriously will need to align assessment, faculty development, and clinical placement design around this principle—treating digital fluency as a workplace-acquired competence scaffolded by clinical exposure, not as a discrete lecture topic.

The Doctor–Patient Relationship in an AI Era

The convergence reported across stakeholder groups in the Results is itself a substantive finding. Patients, clinicians, and students approach AI from different positions—with different stakes, expertise, and professional identities—yet they arrive at substantially overlapping conclusions about the role AI should occupy in clinical encounters. That convergence is intellectually significant because it suggests resistance to AI substitution is not the product of any single group's interests: it appears to be a structural feature of clinical practice, in which the doctor–patient relationship carries functions that all three groups independently recognise as irreducible.

For clinicians, the concern is more specific than job displacement and centres on professional epistemic authority: the capacity to make and justify clinical decisions under uncertainty. AI-supported workflows can erode this authority in subtle ways: when algorithms recommend a course of action that diverges from clinical judgment, the path of least resistance for a time-pressured clinician is often to accept the recommendation. The qualitative evidence summarised in the Results about loss of decisional control and unease around professional judgement suggests clinicians are alert to this dynamic. The deeper question is whether institutions are designing workflows that actively protect against it.

Medical students framed their concern as a *curricular imperative* rather than a defensive reaction: they described empathy, trust, and ethical awareness as values that AI must actively safeguard rather than simply preserve. The word "actively" matters here - it implies that protecting these values is a design goal for AI systems, not an automatic by-product of integration. This framing connects directly to the question of explainability. If explainability is treated as a technical requirement (providing interpretable outputs for technical auditing), it underperforms its ethical role. Treated instead as a design principle aimed at preserving shared decision-making and protecting clinicians' interpretive authority, it becomes one of the practical mechanisms through which professional epistemic authority can be defended against the subtle erosion mechanisms described above.

What emerges is a constructive reframing in which AI's role is to free clinicians for the higher-order functions of medicine that depend on human presence. The stakeholder evidence on familiarity reinforces this reframing in a more subtle way: since those with direct experience of AI in practice tend to hold more positive attitudes than those encountering it hypothetically, the relationship between trust and familiarity is bidirectional and experience-dependent. Trust in AI is, accordingly, not a property that can be declared but one that emerges from repeated, transparent, demonstrably useful experience.

Patient Experience, Trust, and Acceptance

The patient statistics reported in the Results sit comfortably with what technology-acceptance theory predicts but reveal something more specific about how patients conceptualise AI. The same respondents who expressed discomfort with provider-driven AI also expressed willingness to use AI in self-management, a pattern that suggests the crucial variable is not AI itself but the *locus of agency* in the clinical decision. Patients appear willing to draw on AI as a tool that extends their own capacity to manage health while resisting AI when it appears to displace the clinician whose judgment they trust in unfamiliar territory. This is consistent with the broader moderation in the literature on familiarity and educational attainment, with acceptance falling sharply for applications in contexts where relational presence is the expected standard.

The asymmetry that runs through these findings is that stakeholders expect AI to meet a higher standard than equivalent human practice. The evidence that stakeholders are less forgiving of AI errors than of human errors, particularly at high stakes, indicates AI is being held to a stricter evidentiary and procedural bar than clinical practice has historically imposed on any other tool. This is not an unreasonable standard but it does

mean AI systems cannot be deployed with the same tolerance for occasional failure that clinicians and patients routinely extend to human practitioners. Institutional design (quality assurance, continuous monitoring, transparent reporting of error) must consequently be more rigorous than what is currently routine.

These dynamics recast patient education and digital-literacy initiatives. They are best understood not as communication campaigns but as structural prerequisites for sustainable adoption: a population that understands what these tools can and cannot do is better prepared for clinical integration. The case for trusted intermediaries - government and medical bodies - to lead public-facing communication reflects the same logic: trust transfers more reliably from institutions whose role is to protect public welfare than from the commercial vendors of the tools themselves. Without these investments, even well-designed AI tools risk adoption resistance driven less by their performance than by the absence of the framing within which their performance can be evaluated.

Implementation Pathway—Enablers and Persistent Barriers

The enablers identified in the Results converge on a single underlying principle: AI adoption succeeds when the design of the system reflects the workflow, identity, and authority of the clinicians who will use it. Human-in-the-loop design is the dominant recommended model precisely because it accommodates the trust asymmetry and stakes sensitivity discussed above, allowing AI to contribute its computational strengths without forcing clinicians to surrender interpretive authority. Meaningful clinician involvement in design, validation, and oversight goes beyond procedural inclusion; it is necessary for safer and more effective integration, as professional participation fosters critical reflection on AI use. Explainable AI sits at the intersection of the technical and the ethical because it is necessary to maintain clinical confidence and shared decision-making. Sustained investment in training and digital literacy across the workforce is, finally, the social infrastructure holding these enablers together. Each technical design choice depends on a population of users capable of exercising the judgment those designs are intended to support.

The barriers described in the Results—data quality, algorithmic bias, reproducibility, hallucination, and workflow integration—are not independent technical problems but different facets of a single underlying question: how does a healthcare system extend its evidence base, governance arrangements, and clinical workflows in step with the rapid pace of algorithmic development? Each barrier reflects a domain in which institutional infrastructure for evaluating and integrating new tools has not matured at the same pace as the tools themselves.

Input data shortages, for instance, are not only a technical challenge but an equity challenge: under-represented populations are precisely those whose data is scarce, so any model trained on the available corpus will be more accurate for over-served groups. Algorithmic bias compounds this dynamic in ways that cannot be resolved by technical fixes alone, because bias detection requires governance arrangements for ongoing monitoring in deployment and not only for initial validation. Reproducibility problems, particularly pronounced in high-variance fields such as neurology and psychiatry, mean that AI performance validated in one population cannot be assumed to transfer to another, raising the bar for the kind of external validation that evidence-based medicine has historically required of pharmacological and device interventions. Hallucination interacts directly with the cognitive trust calibration discussed above: the more fluent and authoritative AI output appears, the more vulnerable clinicians are to accepting it uncritically. Taken together, these barriers carry an implication that is not merely clinical but epistemic: AI forces clinical practice to take a more rigorous stance toward evidence than it has historically taken, because the systems involved demand a kind of validation they are uniquely susceptible to failing.

Workflow integration is, in this light, the most underappreciated barrier because it is the point at which every other barrier converges in clinical practice. The findings about digital formularies, decision support infrastructure, and reimbursement models in the Results reflect the broader lesson that healthcare adoption is rarely gated by what a tool can do but by how it fits into existing professional, financial, and administrative routines. The diffusion-of-innovation pattern visible in early AI scribe uptake—where early champions drive adoption while the broader workforce stalls—mirrors the trajectory of nearly every prior healthcare technology, from electronic health records to clinical decision support. The call from AI scribe focus group participants for clear, time-efficient guidance from medical colleges, government agencies, and insurers is, in this sense, a request for top-down governance scaffolding that time-pressured clinicians can rely on. Absent that scaffolding, adoption will continue to concentrate in well-resourced centres and among clinicians willing to absorb the learning cost.

Limitations and Future Directions

Several limitations bound the conclusions of this review. First, the underlying literature is methodologically heterogeneous: variation in study design, instrumentation, and institutional context constrains direct comparison and prevents pooled quantitative estimates. Survey-based studies of stakeholder attitudes are subject to priming effects, and many do not cite concrete examples of AI applications that could orient respondents, leaving participants to imagine the technology rather than respond to specified use cases. Second, longitudinal studies measuring the sustained impact of AI-augmented training on clinical performance, decision-making, or patient outcomes remain scarce—the evidence base is dominated by short-term feasibility and acceptance studies rather than downstream outcome evaluations. Third, much of the existing literature is concentrated in well-resourced healthcare systems and academic medical centres, which may limit generalisability to lower-resource settings.

Future research should therefore prioritise three directions. The first is prospective, multimodal evaluation of human–AI collaboration in real clinical environments, with attention to both technical performance and the lived experience of clinicians and patients. The second is the development of validated, standardised instruments for measuring AI knowledge, attitudes, and competencies among medical students and practising physicians—without such instruments, the literature's recurring finding that AI literacy is low will remain difficult to quantify, compare across settings, or track over time. The third is the explicit study of equity. Given that stakeholders consistently warn about the risk of AI-driven innovation concentrating benefits in already well-served populations, future research should examine how AI can be deployed to extend rather than narrow access to high-quality care.

Conclusions

This narrative review has synthesised the current state of artificial intelligence in medical education and clinical practice through the intersecting themes of medical training, the deployment of AI in diagnostics and treatment, the attitudes of patients and clinicians, and the implementation challenges that shape adoption. Across all of these, a single coherent direction emerges.

The most consistent message across the literature is that AI delivers its greatest value when positioned as a collaborative instrument that augments the work of clinicians rather than substituting for them. When treated this way, AI extends perceptual reach, accelerates routine workflows, and surfaces patterns that might otherwise be missed, while the relational presence, ethical judgment, and contextual reasoning that define clinical practice remain firmly in human hands. The convergence of stakeholder views is the strongest evidence for this framing: patients, clinicians, and students approach AI from very different positions yet arrive at substantially overlapping conclusions about the role it should occupy in clinical encounters, and that overlap suggests resistance to AI substitution is a structural feature of clinical practice rather than the product of any single group's interests.

In medical education, the evidence describes a clear mismatch between expectation and provision. Most medical students expect digital health technology to be part of their formal curricula but perceive their own competencies as inadequate. The most promising path forward is situated, clinically embedded learning. Mobile and bedside access to point-of-care information anchors knowledge in the specific problems students will later manage and turns digital fluency into a workplace-acquired competence rather than a lecture topic. Schools that align assessment, faculty development, and clinical placement design around this principle will be best placed to produce graduates who can work alongside algorithmic tools rather than around them.

In clinical practice, the evidence supports a consistent operating model. AI performs at or above expert level on narrowly defined diagnostic tasks across radiology, dermatology, pathology, and ophthalmology, with its most enduring contribution in accelerating routine workflows and refining rather than replacing clinical judgment. The qualitative experience of clinicians using AI scribes captures both the promise and the risk: meaningful gains in efficiency, note quality, and even the clinical interaction are accompanied by concerns about over-reliance, transparency, and what happens to consultation data once captured, particularly under time pressure. The model that best aligns with this sentiment, and that is most likely to earn the trust of clinicians who currently report insufficient understanding of how these systems function, is one in which AI flags and prioritises while clinicians retain interpretive authority for the final decision.

For patients, the central message is conditional optimism. Most are willing to engage with AI tools in managing their own health while simultaneously preferring human contact in sensitive domains such as mental health, chronic disease management, and end-of-life decisions. The asymmetry that runs through the

stakeholder data is consequential: stakeholders expect AI to meet a higher standard than equivalent human practice, and they are less forgiving of AI errors than of equivalent human ones, particularly in serious illness or when AI output diverges from clinician opinion. This imposes on AI a stricter evidentiary and procedural bar than has historically been applied to other clinical tools, with direct implications for quality assurance, continuous monitoring, and transparent reporting of error. Public-facing communication led by trusted government and medical bodies is therefore best understood as a structural prerequisite for sustainable adoption rather than a peripheral exercise in information dissemination.

For policy and institutional design, several priorities emerge from the synthesis. First, regulators and professional colleges must provide practical, trustworthy guidance on AI integration that is specific enough to be actionable for time-constrained clinicians. Second, healthcare systems must build the digital formularies, clinical decision-support infrastructure, and reimbursement mechanisms that allow AI tools to be prescribed, monitored, and documented within existing electronic health record workflows. Third, sustained investment in digital literacy across clinicians, students, and patients must accompany the technical rollout, since every design choice around explainability, human-in-the-loop oversight, and clinician involvement depends on a population of users capable of exercising the judgment those designs are intended to support.

Equity runs through these priorities and deserves to be named explicitly. Stakeholders consistently warn that AI-driven innovation risks concentrating benefits in already well-served populations, while lower-resource settings, primary care clinics, rural hospitals, and disadvantaged communities risk being left behind. Designing for equity involves strategies such as improving dataset representation and algorithmic transparency, and ensuring reimbursement structures support access beyond tertiary centers.

The trajectory of AI in medicine is unlikely to follow a simple substitution narrative. The evidence reviewed here supports a future in which AI's computational strengths are aligned with the relational and ethical capacities of human clinicians, producing a model of care that is, if guided deliberately, more precise and more humane than either humans or machines could achieve alone. Several research directions are emerging: the need for improved AI competencies among physicians and strategies for extending access to care. Realising the promise of AI in medicine will depend less on the next technical breakthrough than on the education, governance, and equity scaffolding built around it.

Funding statement: The article did not received any funding.

Conflict of Interest Statement: No conflicts of interest to declare.

REFERENCES

1. Alowais, S. A., Alghamdi, S. S., Alsuhebany, N., Alqahtani, T., Alshaya, A., Almohareb, S. N., Aldairem, A., Alrashed, M., Saleh, K. B., Badreldin, H. A., Yami, M. S. A., Harbi, S. A., & Albekairy, A. (2023). Revolutionizing healthcare: The role of artificial intelligence in clinical practice. *BMC Medical Education*. <https://doi.org/10.1186/S12909-023-04698-Z>
2. Boillat, T., Otaki, F., Baghestani, A., Zarnegar, L., & Kellett, C. (2024). A landscape analysis of digital health technology in medical schools: Preparing students for the future of health care. *BMC Medical Education*. <https://doi.org/10.1186/S12909-024-06006-9>
3. Corridon, P. R., Wang, X., Shakeel, A., & Chan, V. (2022). Digital technologies: Advancing individualized treatments through gene and cell therapies, pharmacogenetics, and disease detection and diagnostics. *Biomedicines*. <https://doi.org/10.3390/BIOMEDICINES10102445>
4. Davies, B., Rafique, J., Vincent, T. R., Fairclough, J., Packer, M. H., Vincent, R., & Haq, I. (2012). Mobile medical education (MoMed) - how mobile information resources contribute to learning for undergraduate clinical students - a mixed methods study <https://doi.org/10.1186/1472-6920-12-1>
5. Dhar, P., Rocks, T., Samarasinghe, R. M., Stephenson, G., & Smith, C. M. (2021). Augmented reality in medical education: Students' experiences and learning outcomes <https://doi.org/10.1080/10872981.2021.1953953>

6. Gordon, W. J., Landman, A., Zhang, H., & Bates, D. W. (2020). Beyond validation: Getting health apps into clinical practice <https://doi.org/10.1038/S41746-019-0212-Z>
7. Irons, J., Duenser, A., Pickett, T., Haysom, G., & McGrath, M. J. (2026). AI in medical practice: Doctors' perspective on the benefits, challenges and facilitators of artificial intelligence scribe use. *Health and Technology*. <https://doi.org/10.1007/S12553-025-01042-X>
8. Merlo, E. M., Sparacino, G., Silvestro, O., Giacobello, M. L., Meduri, A., Casciaro, M., Gangemi, S., & Martino, G. D. (2026). The role of artificial intelligence in shaping the doctor–patient relationship: A narrative review. *Healthcare*. <https://doi.org/10.3390/HEALTHCARE14040481>
9. Nağıyeva, G., & Süleymanov, H. (2026). The irreplaceable role of human doctors in the age of artificial intelligence. *EuroGlobal Journal of Linguistics and Language Education*. <https://doi.org/10.69760/EGJLLE.2602028>
10. Scott, I., Carter, S. M., & Coiera, E. (2021). Exploring stakeholder attitudes towards AI in clinical practice. *BMJ Health & Care Informatics*. <https://doi.org/10.1136/BMJHCI-2021-100450>
11. Stroud, C., Onnela, J., & Manji, H. K. (2019). Harnessing digital technology to predict, diagnose, monitor, and develop treatments for brain disorders. *Npj Digital Medicine*. <https://doi.org/10.1038/S41746-019-0123-Z>
12. Ward, J., Gordon, J., Field, M. J., & Lehmann, H. P. (2001). Communication and information technology in medical education [https://doi.org/10.1016/S0140-6736\(00\)04173-8](https://doi.org/10.1016/S0140-6736(00)04173-8)
13. Wickramasinghe, N., Thompson, B., & Xiao, J. (2021). The opportunities and challenges of digital anatomy for medical sciences: Narrative review <https://doi.org/10.2196/34687>