



Review

Analyzing the impact of AI in healthcare: Trends, myths, challenges, and future directions

Anshika Jangra¹, Maidul Islam^{2,*}, Kuldeep Gurjar¹ and Swati Sharma³

¹ Department of Global Talent, The University of Suwon, Hwaseong-si, South Korea

² Department of International Business, Keimyung University, Daegu, South Korea

³ School of Pharmacy, Sungkyunkwan University, South Korea

* **Correspondence:** Email: maidul@kmu.ac.kr; Tel: +820535805967; Fax: +820535805967.

Abstract: The integration of artificial intelligence (AI) is significantly transforming healthcare, offering substantial improvements in efficiency, diagnostics, and patient care. This structured narrative review synthesizes peer-reviewed evidence retrieved from PubMed, Scopus, Web of Science, and Google Scholar (January 2017–October 2025) and applies an explicit eligibility protocol to map current applications, trends, and limitations of AI in healthcare. We explicitly distinguish three families of AI relevant to clinical practice—predictive AI (discriminative machine-learning models for risk stratification and diagnosis), generative AI (large language models and generative adversarial networks producing synthetic content), and agentic AI (goal-directed autonomous systems)—and analyze their distinct decision-making loci, adoption pathways, and complementarities with healthcare systems. Beyond technological advancements, this paper addresses prevalent misconceptions and myths surrounding AI in healthcare. These common misunderstandings often involve concerns about AI replacing human professionals, its impact on costs, and the absolute reliability of AI-driven decisions. Dispelling these myths is crucial for fostering informed understanding and professional acceptance. Furthermore, we provide an in-depth analysis of the multifaceted challenges that hinder AI's effective integration. These encompass technical hurdles like data quality and interpretability, alongside critical ethical, social, economic, and legal considerations. By synthesizing current developments and critically evaluating both potential benefits and obstacles, this review offers a balanced perspective on the present state and future direction of AI in healthcare. It emphasizes the need for strategic planning, ethical governance, and collaborative efforts to navigate these challenges responsibly and harness AI's transformative capacity for improved global health outcomes.

Keywords: artificial intelligence; digital healthcare; AI in healthcare trends; challenges of AI in healthcare; societal impact of AI in healthcare; predictive AI; generative AI; agentic AI; structured narrative review

1. Introduction

AI is rapidly becoming part of the everyday language of healthcare. For consistency, throughout this review, we use the following terminology: Artificial intelligence (AI) refers to the broad field of computational systems that perform tasks normally requiring human intelligence; machine learning (ML) is a subfield of AI in which algorithms learn patterns from data; deep learning (DL) is a subset of ML based on multi-layer artificial neural networks; generative AI (GenAI) denotes models that produce new content—text, images, code, or molecular structures—and includes large language models (LLMs) such as the GPT and Med-PaLM families and generative adversarial networks (GANs); natural language processing (NLP) is the set of techniques that enable computers to understand and generate human language; and agentic AI refers to goal-directed systems that plan, act, and adapt across multi-step tasks with limited human input. These terms are used in this restricted sense throughout the manuscript. Conversations about AI's role range from optimism about its ability to transform medical practice to hesitation about its limitations and risks. In this review, we look at AI in healthcare from two complementary perspectives. The first is what we call the AI perspective: a view that highlights the promise of AI tools such as machine learning and predictive modeling. From this angle, AI offers opportunities to improve patient outcomes, reduce costs, and make clinical and administrative work more efficient. At the same time, this perspective also reveals a set of challenges: questions about reliability, bias, data privacy, and the broader ethical concerns that come with automating aspects of care. We aim to unpack these trends and misconceptions, while also pointing to the ongoing debates about what role AI should realistically play in health systems. The second lens is the healthcare perspective, which grounds the discussion in the lived realities of clinical practice and patient care. Here, AI is not an abstract technology but a tool applied to concrete areas—diagnostic imaging, electronic health records, genetic data analysis, and even virtual assistants that support patient communication. These examples show both the exciting potential and the practical constraints of AI adoption in healthcare. While AI models trained on large datasets can detect diseases earlier or assist with rare conditions, they also raise questions about responsibility and the balance between automated decision-making and human expertise.

Taken together, these two perspectives allow us to review the promise and pitfalls of AI in healthcare, while emphasizing that its impact will depend as much on thoughtful implementation and ethical oversight as on technical innovation itself.

1.1. The AI perspective

Healthcare today is at a turning point, with AI—including ML, DL, and generative AI (Gen AI)—emerging as a force that could reshape the way medical care is delivered. In this section, we treat AI as an umbrella term but explicitly note where claims pertain specifically to predictive ML/DL systems versus generative or agentic systems, because these subtypes differ both technically and in their pathway to adoption (see Section 1.4). For the purpose of this article, we treat AI as an umbrella term,

without separating Gen AI from the broader discussion. The promise of machine learning and related technologies lies in their ability to improve patient outcomes, reduce costs, and streamline both clinical and administrative work. At the same time, this promise comes with a set of challenges: questions about how effective these systems really are, the misconceptions surrounding their use, and the ethical concerns raised by their deployment in sensitive medical contexts. Our aim in this section is to review the current trends, myths, and challenges that frame the conversation around AI in healthcare.

Recognizing AI's growing role in healthcare also means acknowledging the importance of weighing its benefits against its limitations. With this understanding, healthcare professionals and policymakers are better positioned to make informed choices about how, when, and where AI should be applied. Addressing misconceptions and untangling the complexities of implementation are not only essential for innovation but also for ensuring ethical and responsible use. This balance is best achieved through collaboration—bringing together healthcare providers, computer scientists, and ethicists to shape AI's role in ways that are practical and accountable. Seen in this light, the discussion of AI in healthcare is more than a technical review; it is a step toward understanding how its potential and its pitfalls converge.

The research base also highlights how far AI's influence has already spread. Studies such as Jiang [1], Yu [2], Kothinti [3], and Karunanayake [4] emphasized its transformative role in diagnosis, treatment, patient engagement, and administrative tasks. Breakthroughs range from the use of machine learning on structured datasets to natural language processing applied to unstructured clinical records. Faiyazuddin [5] mapped this landscape in detail, pointing to AI's contributions to diagnostic accuracy, treatment planning, patient interaction, and administrative efficiency. Similarly, Gupta [6] turned to emerging domains such as drug discovery, clinical trials, and patient care—areas where AI is accelerating discovery, improving target identification, and enabling navigation through massive datasets. Yet even alongside these advances, concerns remain. Ethical dilemmas, uneven implementation, and questions about the preservation of professional roles serve as reminders that adoption cannot be driven by technical capability alone.

1.2. The healthcare perspective

AI and its related technologies are now deeply embedded across multiple sectors, and healthcare is no exception. Within this context, the use of advanced algorithms and digital tools to support patient care, research, administration, and service delivery is broadly understood as AI in healthcare. These technologies carry the potential to reshape how care is delivered—enhancing clinical decision-making, streamlining administrative work, and supporting innovation across providers, payers, and pharmaceutical organizations [7].

AI's ability to process diverse types of patient data—ranging from medical records and imaging scans to genetic information—has opened new possibilities for diagnosis and treatment. Using machine learning algorithms, AI systems can detect subtle patterns, identify anomalies, and generate outcome predictions that contribute to more accurate diagnostics and personalized care [8]. In medical imaging, for instance, radiologists increasingly rely on AI-assisted analysis of X-rays, MRIs, and CT scans to identify diseases such as cancer with greater precision and speed [9]. AI can also help develop individualized treatment plans by considering a patient's unique characteristics, predicting medication responses, and adjusting dosages accordingly. Its capacity to learn from large datasets allows AI models to recognize rare or complex disorders, supporting physicians in cases that might otherwise be difficult to diagnose [10].

Beyond diagnostics, AI technologies are also transforming how patients interact with the healthcare system. AI-driven chatbots and virtual assistants now provide around-the-clock assistance—scheduling appointments, answering routine questions, and offering general medical information [11]. These tools not only improve accessibility but also allow healthcare professionals to dedicate more time to critical and complex cases. Similarly, in drug discovery, AI is accelerating research by analyzing massive datasets and predicting molecular properties. This capability holds promise for faster development of new treatments and therapies, offering better outcomes for a range of conditions [12]. On the administrative side, AI is helping to reduce human error and improve efficiency in areas such as medical coding, billing, and patient scheduling, ultimately giving healthcare workers more time to focus on direct patient care [13].

In summary, AI in healthcare presents an opportunity to transform patient care, personalize treatments, and advance medical research. Yet, alongside its benefits, it remains essential to ensure that AI is used ethically and responsibly. Striking the right balance between human judgment and technological capability will determine how successfully AI becomes an integrated, trusted partner in modern healthcare.

1.3. Tracing the evolution of AI in healthcare

The integration of artificial intelligence (AI) into healthcare has led to remarkable progress, shaping how medical care is delivered and understood. Early developments—such as the emergence of expert systems in the 1970s—laid the groundwork for AI’s diagnostic decision-making capabilities [14]. Although these systems provided a starting point, the real transformation began with the rise of machine learning and natural language processing, both of which expanded AI’s reach across clinical, research, and administrative domains, as mentioned in Figure 1.




Figure 1. Analytical timeline in healthcare.

Today, AI plays a critical role in medical imaging, enabling faster and more precise disease detection, while also analyzing electronic health record (EHR) data to identify patterns, predict outcomes, and assess patient risk [15]. In parallel, AI-driven methods have become integral to drug

discovery, accelerating the identification of new therapeutic candidates and reducing costs. Virtual assistants and chatbots now help manage appointments, guide patients, and provide accessible health information, easing the workload on medical professionals [16].

Earlier research laid the foundation for these developments. For instance, Esteva et al. developed a deep learning algorithm capable of diagnosing skin cancer from image data [17], and Song et al. used machine learning to predict the onset of sepsis [18]. Rajkomar et al. demonstrated AI's potential to forecast patient mortality, readmission rates, and hospital stays [19]. While these studies were groundbreaking, more recent work has refined AI's precision and broadened its clinical applications.

Contemporary studies such as Miha et al. highlight AI's capacity to predict cardiovascular risk using diverse, bias-mitigated datasets [20]. Natarajan et al. expanded this progress with models that anticipate adverse drug reactions by integrating genomic and patient data [21]. Ruamviboonsuk et al., in their work published in *The Lancet Digital Health*, developed a deep learning-based interpretation system capable of diagnosing diabetic retinopathy with high accuracy [22], while Jiang et al. introduced a system to predict acute kidney injury in critically ill patients [23]. These examples illustrate how AI is advancing from reactive diagnostics toward proactive, preventative healthcare.

However, the rapid pace of innovation also brings ongoing challenges. Gu et al. raised concerns about AI bias in mortality prediction, particularly among underrepresented racial groups, emphasizing the importance of fairness and transparency [24]. Likewise, data security, privacy, and potential misuse continue to present ethical and operational risks [25]. Recent cross-sectoral analyses have also begun to map the emerging occupational, safety, and societal risks created by predictive, generative, and agentic AI families—risks that healthcare must explicitly anticipate as adoption accelerates [26]. These issues underscore the need for continuous oversight and collaboration between healthcare professionals, technologists, and ethicists to ensure responsible AI integration.

Together, this growing body of research illustrates AI's evolving story in healthcare, marked by significant achievements, but equally by the need for vigilance, inclusivity, and ethical balance.

1.4. Taxonomy of AI families in healthcare: predictive, generative, and agentic

Discussions of AI in healthcare frequently treat technology as a single block, yet the systems currently being deployed differ substantially in their underlying methods, decision-making locus, regulatory pathways, and complementarity with health-system structures. Following recent cross-sector analyses of AI risk and adoption [26], we distinguish three families of AI that are most relevant to current clinical practice.

Predictive AI comprises discriminative ML and DL models—logistic regression, gradient-boosted trees, support-vector machines, convolutional, and recurrent neural networks—that map structured or imaging data to a label, score, or risk estimate. Examples include CNN-based diabetic retinopathy detectors, ML sepsis-onset predictors, and EHR-based mortality and readmission risk models [17–23]. Decisions about deployment are typically made at the institutional or health-system level: a hospital, payer, or public-health body integrates a predictive model into a workflow as Software as a Medical Device (SaMD), and the model's outputs feed standardized clinical pathways. Adoption, therefore, depends on procurement, validation, and regulatory clearance, and the model complements existing institutional decision-making by improving the accuracy and consistency of population-level triage and risk stratification.

Generative AI comprises GenAI/LLMs and GANs/VAEs that produce new content—free text, structured documentation, images, or candidate molecules. In healthcare, this family powers ambient documentation, discharge summary drafting, patient-facing chatbots, and synthetic-data augmentation [27–30]. Importantly, the locus of decision-making shifts from the institution to the individual professional: a clinician chooses, often through a generic interface, whether and how to use a GenAI assistant for a given note, query, or patient explanation. Adoption is therefore closer to the diffusion patterns of consumer software than to traditional medical-device procurement, raising distinct governance issues around reproducibility, hallucination, and accountability.

Agentic AI refers to goal-directed systems that plan and execute multi-step actions, orchestrating tools, querying records, drafting orders, and monitoring outcomes, with limited human prompting between steps [4,31]. These systems can be built on top of predictive and generative components, but are functionally distinct in that they make a sequence of operational decisions on behalf of users. Their adoption sits between institutional and individual modes: agents act inside enterprise workflows yet exercise discretion that has historically belonged to human professionals, which makes Human-in-the-Loop (HITL) safeguards, audit trails, and clear accountability frameworks essential.

Recognizing these three families clarifies several debates in the literature. Concerns about algorithmic bias and external validation [20,24,32,33] apply most strongly to predictive AI; concerns about hallucination, content provenance, and misuse [27,29,30] are characteristic of generative AI; and concerns about emergent behavior, automation bias, and accountability [31,34,35] are most acute for agentic systems. The complementarities also differ: predictive AI complements institutional risk management, generative AI complements individual professional cognition, and agentic AI complements operational workflow execution. These distinctions inform the structure of the analysis that follows and are summarized in Table 1 and Figure 2.

Table 1. Comparative taxonomy of predictive, generative, and agentic AI families in healthcare.

Dimension	Predictive AI	Generative AI	Agentic AI
Core method	Discriminative ML/DL (CNNs, RNNs, gradient boosting).	LLMs, GANs, VAEs producing synthetic content.	Planner–executor systems orchestrating tools and models.
Locus of decision-making	Institutional/health-system (procurement and pathway integration).	Individual professional (clinician-level adoption).	Hybrid: enterprise workflow with delegated discretion.
Adoption pathway	SaMD clearance, formal validation, integration in EHR.	Subscription/API access, governance via clinical policy.	Pilot deployments under HITL constraints; emerging frameworks.
Dominant risk	Distributional shift, demographic bias, miscalibration.	Hallucination, provenance loss, IP, and privacy leakage.	Emergent behavior, accountability gaps, automation bias.
Complementarity with healthcare system	Strengthens population-level risk management and triage.	Augments individual cognitive workload, communication, and documentation.	Coordinates operational workflows and longitudinal monitoring.
Representative references	[17–24]	[27–30,36–38]	[4,31,39]




Figure 2. Taxonomy of AI families in healthcare.

2. Materials and methods

2.1. Type of review

This work is a structured narrative review with a transparent search-and-selection protocol. We adopted a structured narrative format, rather than a full systematic review or meta-analysis, because the question we address (how AI is shaping clinical practice and what myths, challenges, and trajectories surround it) is intentionally broad, draws on heterogeneous study designs (technical evaluations, clinical validations, ethical analyses, and policy reports), and is not amenable to quantitative pooling. To compensate for the inherent limits of a narrative format and respond to reviewer feedback on methodological rigor, we explicitly pre-specified the search strategy, eligibility criteria, and study-selection workflow as described below and summarized in the PRISMA-style flow diagram (Figure 3).

Stage	Description	Count
 Identification	Records identified through database searching and hand-searching	5,891
 Duplicates Removed	Records after duplicates removed	3,124
 Screening	Records screened on title and abstract	3,124
 Excluded (Screening)	Records excluded with reasons	2,756
 Eligibility	Full-text articles assessed	368
 Excluded (Eligibility)	Full-text articles excluded with reasons	277
 Included	Studies included in the qualitative synthesis	91

Figure 3. PRISMA-style study-selection flow.

2.2. Search strategy

Literature was retrieved from four databases—PubMed/MEDLINE, Scopus, Web of Science, and Google Scholar—covering January 2017 to October 2025, a window selected to capture the period in which deep-learning, generative, and agentic AI matured in clinical contexts. Search terms combined three concept blocks using Boolean operators:

- Block A (technology): “artificial intelligence” OR “machine learning” OR “deep learning” OR “large language model” OR “generative AI” OR “foundation model” OR “agentic AI”.
- Block B (domain): healthcare OR medicine OR clinical OR “public health” OR diagnosis OR “drug discovery” OR “electronic health record*”.
- Block C (analytic angle): bias OR ethics OR governance OR adoption OR trust OR explainable OR “clinical decision support” OR myth.

Searches combined Block A AND Block B, then iteratively refined with Block C. Reference lists of seminal reviews [1,2,5,6,36,38,40] were hand-searched for additional eligible items (“snowballing”). Grey literature was permitted only when authored by recognized regulatory bodies, government agencies, or the corporate technical literature of vendors with peer-reviewed publication footprints.

2.3. Eligibility criteria

Studies were included if they (i) were published in English between January 2017 and October 2025; (ii) addressed AI/ML/DL/GenAI/agentic systems in human healthcare contexts (clinical, public-health, administrative, or pharmaceutical R&D); (iii) reported either empirical results, a clearly defined methodological framework, or an analytical/critical perspective relevant to the review's themes (trends, myths, challenges, or future directions); and (iv) were peer-reviewed journal articles, conference proceedings indexed in the above databases, or authoritative grey-literature reports. Studies were excluded if they (i) addressed AI in non-health domains without translational implications; (ii) were editorials, commentaries, or letters without original analysis; (iii) reported only patent or product-marketing material; or (iv) focused exclusively on animal or in vitro studies without clinical relevance. Where multiple publications reported overlapping content from the same research group, the most comprehensive or most recent version was retained.

2.4. Study selection and synthesis

Two authors independently screened titles and abstracts, then full texts, against the eligibility criteria; disagreements were resolved by consensus discussion with a third author. Selected studies were charted into a structured matrix capturing AI family (predictive/generative/agentic), clinical domain, study design, data source, reported strengths, and reported limitations. Synthesis was performed thematically along the four organizing axes used in this review—trends and future directions, myths and misconceptions, challenges, and discussion/synthesis—which were defined a priori. The PRISMA-style study selection process is summarized in Figure 3, and the application-domain mapping derived from this synthesis is presented in Table 2; a comparative critical evaluation of representative studies appears in Table 3.

Table 2. Application areas of AI in healthcare with representative recent studies.

AI-based application area	Representative recent study	References
Medical imaging and radiology	Deep Learning Approaches for Medical Image Analysis and Diagnosis	[41]
Diagnostic decision support and clinical decision-making	Artificial Intelligence and Decision-Making in Healthcare: A Thematic Review	[42]
Electronic health records (EHRs) and predictive analytics	A Review of the Role of Artificial Intelligence in Healthcare	[43]
Drug discovery and personalized medicine	A Review on Revolutionizing Healthcare Technologies with AI and ML	[44]
Telehealth and virtual patient services	AI in Healthcare: Transforming the Practice of Medicine	[40]
Risk stratification, early warning, and clinical safety	Artificial Intelligence in Healthcare Delivery: Prospects and Pitfalls	[45]
Administration and workflow optimization	Implementing AI in Healthcare—The Relevance of Trust	[46]

Table 3. Comparative critical evaluation of representative studies on AI in healthcare across predictive, generative, and agentic families.

Study	AI family	Clinical task	Design	Strengths	Weaknesses/open questions
Esteva et al. [17]	Predictive (DL/CNN)	Skin-cancer classification from clinical and dermoscopic images.	Retrospective single-center benchmark.	Dermatologist-level accuracy on benchmark; demonstrated feasibility of CNNs in dermatology.	Limited skin-tone diversity; no prospective external validation; outcomes not measured.
Song et al. [18]	Predictive (ML)	Early prediction of late-onset neonatal sepsis.	Retrospective EHR cohort.	Practical lead time for intervention; interpretable feature set.	Single-site cohort; limited generalizability; alarm-burden not assessed.
Rajkomar et al. [19]	Predictive (DL on EHR)	Mortality, readmission, and length of stay.	Retrospective multi-site EHR.	Scales to raw EHR; broad clinical utility; reproducible methodology.	Limited interpretability; outcomes did not include downstream clinical action; deployment effects unknown.
Ruamviboonsuk et al. [22]	Predictive (DL imaging)	Diabetic retinopathy screening.	Prospective multi-site interventional cohort.	Real-world prospective deployment; documented operational performance and workflow integration.	Performance shifts vs. lab benchmarks; image-quality dependence; long-term outcome data still maturing.
Mihan et al. [20]	Predictive (ML)	Cardiovascular risk under bias-mitigated training.	Methodological/policy review with empirical examples.	Explicit attention to demographic fairness; pragmatic mitigation strategies.	Recommendations not yet translated into standardized reporting; lacks prospective trial evidence.
Hacking [29] / Chen & Esmailzadeh [30]	Generative (LLM, GAN/VAE)	Clinical documentation; synthetic data; patient communication.	Narrative/analytic perspective and case examples.	Maps real productivity gains and concrete privacy/security risks; identifies governance gaps.	No controlled clinical trials; hallucination rates not severity-stratified; IP and provenance unresolved.
Karunanayake [4] / Bonis [31]	Agentic	Autonomous clinical assistants; closed-loop monitoring.	Conceptual/industry analysis with early prototypes.	Articulates plausible operational gains; proposes HITL safeguards.	Limited peer-reviewed evidence; emergent-behavior and accountability frameworks still embryonic.

2.5. Classification framework

The methodology involves mapping the most prevalent areas of AI application through a review of representative scholarly studies. These application domains were identified and organized according to their prominence in recent research, clinical integration, and overall impact on healthcare delivery and innovation. This structured mapping enables a systematic understanding of how AI technologies are reshaping diverse aspects of healthcare—from diagnostic precision to patient monitoring and administrative optimization. The resulting classification, shown in Table 1, outlines six major domains of AI application, arranged in descending order from the most widely implemented to the least. Each domain includes a representative study that exemplifies AI's role within that specific context.

3. Trends and future directions in AI in healthcare

3.1. Overview

Artificial intelligence (AI) is rapidly transforming the healthcare landscape, with a growing number of companies, institutions, and healthcare providers adopting AI-based systems to enhance medical delivery and research. Its applications span diagnosis, treatment planning, data management, and patient monitoring, fundamentally altering how healthcare professionals engage with technology. AI's integration has led to a more precise understanding of patient needs, improved data handling, and better management of complex medical conditions, ultimately contributing to superior patient outcomes. Its capabilities in processing large datasets, improving surgical precision, and analyzing medical images continue to redefine clinical workflows [27].

3.2. Evolution of multimodal and generative AI architectures

The healthcare AI ecosystem is experiencing a decisive paradigm shift from traditional unimodal systems—such as image-only models in radiology or text-based natural language processing (NLP)—to multimodal AI architectures capable of integrating and contextualizing diverse information sources. These systems synthesize heterogeneous data types, including electronic health records (EHRs), Digital Imaging and Communications in Medicine (DICOM) files, structured genomic datasets, and continuous biosensor time-series data, resulting in richer, more comprehensive patient profiles [27,28]. In parallel, the emergence of generative AI (GenAI), underpinned by large language models (LLMs) and advanced generative adversarial networks (GANs), is transforming the way medical data is utilized. Beyond analytical functions, LLMs are increasingly applied to automate clinical documentation, streamline complex administrative processes, and enhance patient communication platforms [27,29]. Similarly, GANs and variational autoencoders (VAEs) have shown significant promise in generating high-fidelity synthetic patient datasets, which are vital for training AI models in privacy-sensitive settings and for overcoming the persistent challenge of limited annotated clinical data [30].

Collectively, these innovations mark a transition from data analysis to data synthesis and contextual intelligence, signifying the foundation for more adaptive, patient-centered AI systems. This evolution also provides the groundwork for the next critical trend—the integration of AI as a cognitive and workflow assistant, where generative and multimodal systems begin to directly augment clinical decision-making and operational efficiency within healthcare environments.

3.3. *AI as a clinical workflow and cognitive assistant*

Generative AI applications are now functioning as “cognitive assistants” that alleviate clinician burnout by automating repetitive documentation and communication tasks. Such systems can:

- Summarize patient histories for handovers [47].
- Draft clinical reports and progress notes [29].
- Manage communication within hospital systems [29].

These systems aim to restore clinician focus on direct patient care and decision-making, optimizing workflow efficiency [27].

3.4. *Advanced predictive analytics for early and precision disease detection*

One of the most transformative and widely recognized applications of artificial intelligence (AI) in healthcare lies in predictive analytics, particularly for early disease detection and personalized risk stratification. Unlike traditional statistical models, modern AI frameworks can process and interpret vast amounts of longitudinal patient data from heterogeneous sources—ranging from genetic and epigenetic markers, continuous monitoring data from wearable and implantable devices, and lifestyle and behavioral indicators to high-resolution medical imaging and laboratory results [6,48]. By integrating these multifaceted datasets, AI systems are capable of uncovering subtle, nonlinear correlations that might otherwise go undetected by human clinicians. This capability enables the prediction of disease onset, progression trajectories, and treatment responses for complex and chronic conditions such as cardiovascular disease, cancer, and neurological disorders [6,48]. For instance, predictive algorithms can identify early biomarkers of Alzheimer’s disease or forecast the likelihood of cancer recurrence based on patient-specific molecular profiles and imaging trends.

These advances collectively support a proactive and preventative model of medicine, in which care shifts from reactive treatment to early intervention and continuous management. Evidence from recent studies suggests that such AI-powered predictive models not only enhance diagnostic accuracy but also improve clinical decision-making by enabling personalized treatment pathways, reducing hospital readmissions, and optimizing healthcare resource allocation [6]. In essence, predictive analytics serves as a cornerstone in the ongoing transition toward precision healthcare, where prevention and individualized care supersede traditional one-size-fits-all approaches.

3.5. *Acceleration of the drug discovery and development pipeline*

Artificial intelligence, particularly generative AI, is also reshaping the landscape of pharmaceutical research and development (R&D) by dramatically shortening the drug discovery timeline—from initial target identification to lead optimization and clinical validation [40]. This transformation stems from AI’s ability to model complex biological systems, simulate molecular interactions, and generate novel drug candidates at a scale and speed previously unattainable through conventional experimental techniques. Advanced deep learning architectures, including variational autoencoders (VAEs), generative adversarial networks (GANs), and transformer-based models, are increasingly used to design and evaluate new molecular entities *in silico*—predicting their pharmacokinetic and pharmacodynamic characteristics with remarkable precision [30,49]. Simultaneously, the integration of multi-omics data (spanning genomics, proteomics, and

metabolomics) allows AI systems to identify novel druggable targets, elucidate complex biomarker signatures, and predict therapeutic responses across diverse patient populations [49]. Crucially, these AI-driven methodologies are not limited to hypothesis generation; they are actively expediting tangible progress in the drug development pipeline. Emerging evidence indicates that AI-designed molecules are demonstrating higher success rates in early-phase clinical trials compared to traditionally derived compounds, thereby reducing both time and cost associated with preclinical candidate identification [49,50]. Moreover, by leveraging synthetic biology and virtual screening, AI enables researchers to rapidly test millions of compounds and optimize candidates before laboratory synthesis, making the entire process more efficient and sustainable. Ultimately, AI's contribution to drug discovery represents one of the most promising intersections of computational innovation and biomedical science, with the potential to accelerate the delivery of safer, more effective therapies to patients worldwide.

3.6. Emergence of agentic AI and autonomous systems

Agentic AI marks a new phase in healthcare innovation, introducing systems capable of self-directed learning, task execution, and adaptive decision-making [31]. In clinical environments, these models continuously monitor patient vitals, interpret real-time data, and propose context-specific interventions, reducing response times in critical care. Beyond patient management, they streamline administrative operations by automating complex workflows such as prior authorizations and clinical documentation. Although Human-in-the-Loop (HITL) oversight remains essential for validation and accountability, ongoing developments indicate a steady shift toward higher autonomy and greater operational independence [31]. As discussed in Section 1.4, agentic systems blur the institutional/individual boundary that has historically governed clinical AI adoption, and their oversight requirements differ qualitatively from those that apply to predictive or generative components used in isolation.

3.7. Ethical, regulatory, and interoperability challenges

Despite significant advancements, several enduring barriers continue to impede the seamless integration of AI in healthcare:

- Algorithmic bias: The risk of perpetuating inequities due to biased or non-representative training data [51].
- Transparency and explainability (XAI): The need for interpretable models to preserve clinician trust and accountability [52].
- Data governance and interoperability: Challenges in securely integrating and standardizing diverse datasets across healthcare systems [53].
- Regulatory lag: The lack of adaptive regulatory frameworks for dynamically evolving AI models [54].
- Cross-sectoral risk spillovers: Recent analyses of occupational and societal AI risk emphasize that the categories of harm created by predictive, generative, and agentic systems differ in kind, requiring family-specific governance rather than a single “AI ethics” envelope.

3.8. Future trends

3.8.1. Large medical models (LMMs) and foundational AI

AI research in healthcare is rapidly advancing toward large medical models (LMMs), multimodal foundational systems trained on diverse biomedical data, including clinical notes, imaging, genomics, and the full corpus of biomedical literature [36–38]. These models are expected to provide generalizable intelligence with emergent reasoning and deep contextual understanding of complex medical concepts. By enabling efficient fine-tuning or prompting for a wide range of downstream applications, such as high-level diagnostic reasoning, report summarization, and initial drug target hypothesis generation, LMMs have the potential to replace many narrow, task-specific AI systems, thereby standardizing core AI infrastructure across healthcare [36,38].

3.8.2. Closed-loop and agentic systems

Future AI systems will evolve toward agentic and closed-loop models capable of autonomously planning, monitoring, and executing multi-step clinical tasks within safety-constrained environments [2,39]. Applications may include continuous management of chronic diseases, where AI agents process real-time data from wearables and medical devices (e.g., continuous glucose monitors) to adjust therapeutic interventions or alert clinicians based on predictive risk assessments. Human oversight will remain essential, but the paradigm will shift from AI-assisted decision support to autonomous, safety-governed action [2,39].

3.8.3. AI-enabled digital twins

The concept of digital twins—dynamic virtual replicas of individual patients—represents the next frontier in personalized medicine [55]. By integrating multi-scale biological data, from single-cell genomics and transcriptomics to organ-level imaging and whole-body scans, AI-driven digital twins will allow *in silico* simulation of disease progression, prediction of adverse drug interactions, and pre-testing of therapeutic strategies. This capability moves precision medicine from reactive, trial-and-error approaches toward predictive, highly individualized treatment planning [56].

3.8.4. Ambient intelligence and continuous data capture

Clinical AI tools, such as AI scribes, are expected to evolve into pervasive ambient intelligence [57]. Passive sensors in clinical environments will capture interactions, examinations, and procedural details in real time, automatically generating fully coded EHR documentation. Extending beyond clinical settings, home-based and wearable devices will provide continuous passive monitoring, enable the ultra-early detection of physiological deviations, and support predictive, proactive healthcare [57].

3.8.5. Ethical AI and robust governance

The widespread integration of AI necessitates mature ethical and governance frameworks [58]. Key priorities include rigorous bias mitigation across diverse populations, implementation of privacy-preserving techniques such as federated learning, and validated transparency through external audits rather than relying solely on explainable AI (XAI) methods. This approach emphasizes robust clinical performance and safety over superficial interpretability, ensuring that AI adoption maintains trust, fairness, and patient-centered care [59].

To synthesize these emerging trends within a unified conceptual framework, Figure 4 presents a maturity–impact lifecycle model of AI in healthcare. The figure illustrates a continuous, iterative progression spanning AI development, maturity assessment, clinical implementation, impact prediction, performance monitoring, and refinement. This cyclical representation is particularly relevant in the context of the trends discussed above, as it highlights that advancements such as multimodal learning, agentic systems, and digital twins do not evolve linearly but rather through repeated feedback-driven optimization loops. For instance, predictive analytics and generative AI models discussed in Sections 3.2 and 3.4 correspond to the “Predict Impact” and “Assess Maturity” phases, while agentic and closed-loop systems (Section 3.8.2) align with the “Implement AI” and “Monitor Performance” stages. Consequently, the figure reinforces the notion that future AI systems in healthcare will operate within adaptive, continuously learning ecosystems rather than static deployment pipelines.




Figure 4. Maturity vs. impact matrix of future trends.

4. Myths and misconceptions surrounding AI in healthcare

The emergence of sophisticated AI, particularly large language models (LLMs) and generative AI, has ignited a transformative debate within medicine. While enthusiasm for AI in healthcare continues to grow, misconceptions persist that can distort expectations and hinder ethical adoption by conflating the very different capabilities and risks of predictive, generative, and agentic AI families. These misconceptions are systematically contrasted with evidence-based realities in Table 4, which provides a structured comparison between commonly held beliefs and their corresponding factual clarifications. This tabular synthesis supports the critical evaluation of AI narratives and complements the lifecycle perspective illustrated in Figure 4 by emphasizing that both technological maturity and societal perception must evolve in parallel for responsible AI adoption.

4.1. *AI will replace medical professionals*

Reality: AI is an augmentation tool, not a substitute for clinicians.

The belief that AI will fully replace doctors, nurses, or other healthcare workers remains one of the most persistent and least substantiated fears [60,61]. While AI systems outperform humans in narrow, data-intensive domains, such as image recognition, triage, and anomaly detection, they lack the contextual intelligence and emotional depth intrinsic to medical care (Table 1).

i. Contextual judgment and empathy: Clinical decision-making extends beyond data interpretation. It integrates structured information (EHRs, lab values) with unstructured cues such as patient history, emotional state, and social context [2]. Current AI systems cannot replicate empathy, moral reasoning, or the nuanced understanding necessary to build trust and deliver compassionate care [60,61].

ii. Evolving role of experts: Historically, technological innovations have enhanced rather than eliminated clinical roles. In radiology and pathology, for instance, AI automates routine pattern recognition, allowing specialists to focus on quality assurance, consultation, and complex diagnostics [62]. The future is thus collaborative: AI processes data, while clinicians provide insight and empathy.

iii. Job transformation, not elimination: Automation may reduce repetitive administrative tasks but simultaneously creates new opportunities in AI development, validation, oversight, and patient engagement [61].

4.2. *AI makes unbiased and infallible decisions*

Reality: AI inherits and amplifies systemic biases from its training data.

The notion that AI operates as a neutral, purely objective decision-maker is dangerously misleading. An AI model's accuracy and fairness are only as strong as the data it is trained on—a principle often summarized as “garbage in, garbage out” [32,33].

i. Data bias and disparities: Models trained on homogeneous or non-representative datasets tend to perform poorly across diverse populations. This has been observed, for example, in AI systems diagnosing skin lesions, where model accuracy varies across skin tones and demographic groups [60,32,33].

ii. The human factor in algorithm design: Bias can emerge from human decisions during feature

selection, data labeling, and algorithmic modeling. Implicit biases of developers or clinicians may inadvertently propagate inequities [32].

iii. Need for validation, not just accuracy: High accuracy in training environments does not guarantee fairness or generalizability. For high-stakes applications, continuous external validation, fairness audits, and real-world performance monitoring are essential [63].

Table 4. Representation of the myths and facts of artificial intelligence in today's world.

Myths	Facts
AI is replacing medical professionals in healthcare.	AI and clinicians are working collaboratively to improve overall patient care.
Expenditure on healthcare has increased with the introduction of AI.	Healthcare is becoming affordable with the introduction of new technologies, helping reduce manpower.
AI decisions are unbiased and precise.	AI decisions are influenced by the dataset sources used for training.
AI ensures no data security and privacy.	AI uses encryption for data and ensures the security of personal information.
AI is the solution to all healthcare problems.	AI can enhance problem-facing ability but cannot eliminate the underlying challenges.
AI is better at identifying diseases.	AI can support physicians in making more accurate, data-driven clinical decisions.
AI is complex and difficult to understand.	Most AI tools are user-friendly, though some require targeted training for effective use.

4.3. *AI will always increase healthcare costs*

Reality: AI drives long-term efficiency and cost reduction.

Although implementing AI infrastructure and training personnel requires substantial initial investment, economic evidence increasingly supports AI as a driver of long-term cost savings and system-wide efficiency [64,65].

i. Efficiency and optimization: AI reduces administrative overhead through automated billing, scheduling, and documentation—activities that can account for nearly 25% of healthcare expenditure [64]. It also enhances resource allocation, surgical scheduling, and patient throughput.

ii. Improved outcomes, lower costs: By improving diagnostic precision and early detection, AI reduces unnecessary procedures and complications, resulting in fewer hospitalizations and lower treatment costs [65].

iii. Accessibility and scale: AI-driven telemedicine, remote diagnostics, and wearable monitoring solutions expand access to quality care in underserved regions, at a fraction of conventional costs [62].

4.4. *AI data collection is inherently less secure*

Reality: AI presents new risks but also strengthens data protection.

The training of AI systems on vast health datasets raises legitimate privacy concerns, yet many of the same technologies powering AI can also be used to enhance data security [66,67].

- i. The security paradox: Aggregating data from EHRs, wearables, and trials increases exposure risks. However, strong regulatory adherence to HIPAA and GDPR, combined with technical safeguards, mitigates these threats [66].
- ii. Privacy-enhancing technologies (PETs): a) Federated learning enables collaborative model training without transferring raw data [67]. B) Differential privacy injects controlled noise to protect individual identities. C) AI-based threat detection proactively identifies cyber risks faster than traditional systems [66].
- iii. Governance and trust: Security depends on strong data governance frameworks, encryption, and privacy-by-design principles that sustain public confidence [67].

4.5. AI is a universal healthcare solution

Reality: AI is a specialized tool, not a panacea. AI has transformative potential, but it is not a universal remedy for the systemic challenges in global healthcare.

- i. Narrow vs. general intelligence: Today's systems operate within artificial narrow intelligence (ANI), excelling only at specific tasks (e.g., classifying tumors, reading X-rays). Artificial general intelligence (AGI), the kind that can think and reason like a human, remains hypothetical and far from realization [61,68].
- ii. Limits of systemic impact: While AI can improve workflow efficiency, data-driven policy, and patient management, it cannot resolve structural issues such as inequitable access, workforce shortages, or the complexity of disease itself [2].
- iii. Empowering the patient: The next evolution of AI in healthcare emphasizes patient-facing tools, systems that help individuals interpret symptoms, track chronic conditions, and make informed health decisions (Figure 2). This model redefines clinicians as guides and collaborators rather than sole decision-makers [69].

5. Challenges associated with AI in healthcare

The integration of AI into healthcare is transforming clinical practice, enabling unprecedented advances in diagnosis, treatment, and patient monitoring. However, this transformation is not without its complexities. The deployment of AI in healthcare brings forth a host of technical, ethical, and socio-economic challenges that must be addressed to ensure that these technologies are secure, equitable, and clinically reliable. The following subsections outline the key barriers to responsible and sustainable AI adoption.

5.1. Technical and data infrastructure challenges

The success of AI implementation in healthcare relies on the availability of high-quality data and a robust, interoperable technical infrastructure. Healthcare data, by nature, presents unique challenges that can hinder AI performance and generalizability.

5.1.1. Data availability, quality, and fragmentation

AI models require large, diverse, and well-labeled datasets for optimal training and generalization. However, real-world healthcare data is often fragmented and inconsistent [70]:

- i. Heterogeneity and complexity: Electronic health records (EHRs) include a mix of structured data (lab results, vitals), semi-structured data (imaging, waveform), and unstructured narratives (clinical notes, pathology reports). These heterogeneous sources complicate data harmonization and model training [15,71].
- ii. Inadequate representation: Both rare diseases and healthy populations tend to be underrepresented in datasets, leading to poor calibration in differential diagnosis and population-level screening tasks [72].
- iii. Lack of standardization: Information is stored across multiple legacy systems and incompatible formats, impeding data aggregation and large-scale AI deployment. This lack of semantic interoperability remains a critical bottleneck [73].

5.1.2. Algorithmic robustness and trustworthiness (explainability)

Trust in AI systems depends on transparency, validation, and reliability. The “black box” problem—AI’s opacity in decision-making—poses a major barrier to clinical trust and adoption.

- i. Model validation and rigor: Algorithms must be rigorously validated across diverse, real-world clinical settings rather than just research environments, ensuring safety and generalizability [74].
- ii. Explainable AI (XAI): Clinicians must understand why a model produces a given recommendation. Explainability supports accountability, enhances user trust, and allows clinician oversight in critical care contexts [75].
- iii. Workflow integration: Integrating AI-driven decision support into existing clinical processes often leads to friction due to usability issues, workflow misalignment, or limited infrastructure, especially in resource-constrained environments [76].

Beyond these established issues, three weaknesses of current AI approaches deserve sharper emphasis. First, most published predictive models are validated retrospectively on convenience samples and rarely undergo prospective, multi-site external validation, leading to documented performance drops when deployed (“distributional shift”) [22,74]. Second, generative models in clinical contexts continue to produce confidently stated but factually incorrect outputs (“hallucination”), and current evaluation pipelines for clinical LLMs do not yet stratify failures by clinical risk severity, which limits the actionability of reported accuracy figures [27,29]. Third, agentic systems amplify both of these failure modes by chaining model calls without principled uncertainty propagation, so a small upstream error can cascade into a clinically significant downstream action [4,31]. These observations are revisited and quantified in the comparative-evaluation table presented in Section 6 (Table 3).

5.2. Ethical and legal challenges

AI’s role in healthcare introduces ethical dilemmas surrounding fairness, transparency, accountability, and patient autonomy. These challenges underscore the need for robust ethical frameworks and adaptive legal systems.

- i. Bias, fairness, and equity: Algorithmic bias poses one of the most pressing ethical challenges. When AI systems reflect the biases present in their training data, they risk amplifying health disparities, particularly for underrepresented populations [77]. Ensuring fairness and inclusivity requires balanced datasets, continuous auditing, and clear ethical standards [78].
- ii. Patient autonomy and informed consent: The complexity of AI systems makes it difficult to achieve truly informed consent. Patients must understand how algorithms use their data and influence their care decisions [79]. Achieving meaningful transparency may require new patient education frameworks and regulatory communication models.
- iii. Liability and accountability: Establishing clear accountability for AI-related errors is a significant legal challenge. In the case of a misdiagnosis or malfunction, responsibility can be diffuse, falling on the clinician, hospital, developer, or manufacturer [80]. Defining liability frameworks is critical to protect both patients and practitioners while promoting innovation.
- iv. Data privacy and cybersecurity: The vast amount of personal health information (PHI) used in AI development increases the risk of cyberattacks and data misuse. Traditional privacy laws, such as HIPAA, must evolve to match the advanced data aggregation and analytics capabilities of modern AI systems [43,81].

5.3. *Socio-economic and human factors challenges*

AI's success depends not only on algorithms but also on human acceptance and economic sustainability. Technology adoption in healthcare is deeply shaped by behavioral, organizational, and financial factors.

5.3.1. Human factors and adoption

Clinician engagement is critical for successful AI deployment [34]. However, several behavioral factors complicate this process:

- i. Automation bias: Over-reliance on AI recommendations can lead clinicians to overlook contextual nuances or patient-specific details [35].
- ii. Clinician–AI communication: Poorly designed interfaces and opaque AI outputs can disrupt workflow, increase cognitive load, and lead to clinician fatigue [82].
- iii. Reskilling and education: Continuous professional training is essential for clinicians to interpret AI outputs, recognize limitations, and collaborate effectively with AI systems [82,83].

5.3.2. Economic sustainability and return on investment (ROI)

Although AI can reduce costs long-term, short-term economic hurdles remain significant:

- i. High initial costs: The expenses associated with development, validation, and implementation, including infrastructure upgrades, can be prohibitive [84].
- ii. Demonstrating ROI: Health systems must see measurable improvements in patient outcomes and operational efficiency to justify investments [85].
- iii. Regulatory burden: The cost of obtaining approvals (e.g., FDA clearance for Software as a Medical Device) adds complexity and delays implementation [86].

6. Discussion and future directions

The review so far has traced how AI is reshaping healthcare—across imaging, predictive analytics, drug discovery, and autonomous systems—while also spotlighting myths, ethical tensions, and infrastructural constraints. In this section, we synthesize these threads and suggest paths forward.

6.1. *Synthesis of key themes*

Across all domains, a few patterns emerge:

- i. Human–AI collaboration, not substitution: The myth that AI will replace clinicians is unfounded. The evidence suggests AI works best when paired with human judgment, empathy, and oversight.
- ii. Data quality as a limiting factor: Even the best algorithms fail if fed poor, biased, fragmented, or unrepresentative data. Many challenges trace back to data infrastructure.
- iii. Ethics, governance, and trust: No matter how advanced the model, adoption will lag if clinicians and patients do not trust the system. That requires transparency, accountability, and clear regulatory frameworks.
- iv. Uneven implementation and infrastructure gaps: AI’s potential is often limited by resource constraints, lack of standardization, or user resistance, particularly in less-resourced settings (e.g., rural, lower-income regions).
- v. Evolving role of evidence generation: AI is pushing beyond retrospective studies into prospective trials, synthetic controls, and adaptive learning systems—forcing new norms for validation and regulation.

6.2. *Gaps in current literature*

- i. Prospective, real-world validation: Many models report high accuracy retrospectively, but fewer have been tested prospectively, externally, or in operational settings.
- ii. Explainability and usability studies: There is limited empirical work on how clinicians interpret AI outputs, how interfaces affect trust, and how “black box” systems perform in live workflows.
- iii. Long-term impact on skills and practice: As AI becomes more common, how will it shape clinicians’ diagnostic acumen, reasoning, and continued learning?
- iv. Global and equity perspectives: Most studies come from high-income countries; fewer examine deployment in low- and middle-income settings, where infrastructure and data constraints differ significantly [87].
- v. Failure-mode stratification for generative clinical AI: Reported accuracy figures for clinical LLMs rarely separate hallucinations that are clinically inert from those that are clinically dangerous; standardized, severity-stratified failure-mode evaluations are still missing from the literature.

6.3. *Policy, organizational, and educational implications*

- i. Adaptive governance frameworks: Regulators must keep pace with evolving models (e.g., model update cycles, re-training, post-market evaluation) [88].

- ii. Data stewardship and ecosystem design: Concepts like data stewardship (balancing societal, institutional, and individual interests) may fill gaps in accountability and trust [89].
- iii. Embedding ethics from the start: The ethical framework must be integrated into every stage—data collection, model building, validation, deployment—not tacked on afterward [90].
- iv. Workforce education and cultural change: Clinicians, data scientists, ethicists, and administrators must learn to collaborate. AI literacy must be part of training curricula, not optional.
- v. Infrastructure investment and interoperability: Standards for semantic interoperability, data exchange, and annotation must be enforced [91].

6.4. Future directions

Looking ahead, we propose several directions for research and practice:

- i. Hybrid clinical trials with AI components: Incorporating AI arms, synthetic controls, and real-world validation.
- ii. Explainability tools co-designed with clinicians: Interfaces that translate AI reasoning into actionable, understandable recommendations.
- iii. Federated and privacy-preserving learning: To bridge data silos without compromising patient privacy.
- iv. Longitudinal, real-world outcome studies: Tracking how AI affects health outcomes, cost trajectories, and clinician behavior over time.
- v. Frameworks for autonomous AI deployment: As “agentic AI” becomes more feasible, frameworks for safe, human-supervised autonomy will be critical.

6.5. Critical comparison of representative studies

To strengthen the critical analysis of the studies cited above, Table 3 presents a head-to-head comparison of seven representative works that span the predictive, generative, and agentic AI families and the main clinical domains discussed in this review. For each study, we summarize the AI family, clinical task, study design, principal strengths, and principal weaknesses or unresolved questions. Three observations follow from this comparison. First, high-profile predictive models with reportedly excellent retrospective performances, such as the CNN-based dermatology classifier of Esteva et al. [17], remain limited by demographic skew in their training data and by the absence of multi-site prospective validation. Second, prospective real-world studies are rare, with the diabetic-retinopathy program of Ruamviboonsuk et al. [22] standing out as an example of how operational deployment uncovers failure modes that are invisible in benchmark accuracy figures. Third, generative and agentic systems [27,29,31] are still reported predominantly through narrative case studies and vendor-driven evaluations rather than through controlled clinical trials, which makes claims about their workflow benefits provisional. Together, these patterns motivate the family-specific governance recommendations in Section 5 and the future research agenda in Section 6.6.

7. Conclusion and outlook

Artificial intelligence is no longer a futuristic aspiration in healthcare—it is actively reshaping how medicine is practiced, researched, and delivered. From imaging and diagnostics through drug development

to autonomous systems, AI's influence has spread across virtually every domain of clinical care.

However, this review underscores that technological sophistication alone is insufficient. For AI to fulfill its promise, integration must be guided by ethical foresight, robust infrastructure, human-centered design, and adaptive governance. AI must remain a tool of augmentation, not replacement, blending computational power with human insight. Recognizing that predictive, generative, and agentic AI follow distinct adoption logics and create distinct risks (Sections 1.4, 5.2, and 6.1) is, in our view, a prerequisite for governance frameworks that are both effective and proportionate.

As the field strides into the era of large medical models, continuous learning systems, and semi-autonomous agents, we must remain vigilant. The next chapter of AI in healthcare depends not just on better models, but on trust, fairness, and equity.

If chosen wisely, AI can transform healthcare from reactive to proactive, from one-size-fits-all to personalized, and from fragmented systems to interconnected ecosystems—never replacing the human touch but elevating it.

Limitations

Several limitations should be noted. First, this is a structured narrative review rather than a systematic review or meta-analysis: although the search strategy and eligibility criteria were pre-specified (Section 2.2–2.4), no formal risk-of-bias appraisal was performed, and no quantitative synthesis was attempted. Second, the search was restricted to four databases and to English-language publications, which may have under-represented evidence from non-English health systems and from regional repositories that are particularly relevant to low- and middle-income settings. Third, the field is moving rapidly, particularly for generative and agentic AI, so some technical claims will inevitably be superseded between submission and publication; we have attempted to mitigate this by privileging peer-reviewed sources and by anchoring claims to study design rather than to vendor announcements. Fourth, although we have introduced a three-family taxonomy, real systems increasingly combine predictive, generative, and agentic components, and the boundaries between families are themselves a moving target.

Use of AI tools declaration

The authors declare they have used Artificial Intelligence (AI) tools to refine the diagram (Figure 1, 2, 3, and 4) in the article.

Authors' contribution

Conceptualization, A.J., K.G., M.I.; methodology, A.J., S.S; formal analysis, M.I., K.G, and A.J.; investigation, A.J., K.G.; writing—original draft preparation, A.J., K.G., M.I.; writing—review and editing, A.J., K.G., M.I. and S.S.; visualization, A.J., K.G., M.I.; All authors have read and agreed to the published version of the manuscript.

Conflict of interest

All the authors declare that there is no conflict of interest.

References

1. Jiang F, Jiang Y, Zhi H, et al. (2017) Artificial intelligence in healthcare: past, present and future. *Stroke Vasc Neurol* 2: 230–243. <https://doi.org/10.1136/svn-2017-000101>
2. Yu KH, Beam AL, Kohane IS (2018) Artificial intelligence in healthcare. *Nat Biomed Eng* 2: 719–731. <https://doi.org/10.1038/s41551-018-0305-z>
3. Kothinti RR (2024) Deep learning in healthcare: Transforming disease diagnosis, personalized treatment, and clinical decision-making through AI-driven innovations. *World J Adv Res Rev* 24: 2841–2856. <https://doi.org/10.30574/wjarr.2024.24.2.3435>
4. Karunanayake N (2025) Next-generation agentic AI for transforming healthcare. *Inform Health* 2: 73–83. <https://doi.org/10.1016/j.infoh.2025.03.001>
5. Faiyazuddin M, Rahman SJQ, Anand G, et al. (2025) The impact of artificial intelligence on healthcare: a comprehensive review of advancements in diagnostics, treatment, and operational efficiency. *Health Sci Rep* 8: e70312. <https://doi.org/10.1002/hsr2.70312>
6. Gupta U, Pranav A, Kohli A, et al. (2024) The contribution of artificial intelligence to drug discovery: Current progress and prospects for the future, *Microbial data intelligence and computational techniques for sustainable computing*, Springer Nature, 1–23. https://doi.org/10.1007/978-981-99-9621-6_1
7. Yusuf A, Olaniyan L, Ayobami HS, et al. (2025) The Role of Technology in Enhancing Healthcare Administration and Service Delivery. *CogNexus* 1: 37–50. <https://doi.org/10.63084/cognexus.v1i02.77>
8. Y Kumar, A Koul, R Singla, et al. (2023) Artificial intelligence in disease diagnosis: a systematic literature review, synthesizing framework and future research agenda. *J Amb Intel Hum Comp* 14: 8459–8486. <https://doi.org/10.1007/s12652-021-03612-z>
9. Tang X (2019) The role of artificial intelligence in medical imaging research. *BJR Open* 2: 20190031. <https://doi.org/10.1259/bjro.20190031>
10. Visibelli A, Roncaglia B, Spiga O, et al. (2023) The impact of artificial intelligence in the odyssey of rare diseases. *Biomedicines* 11: 887. <https://doi.org/10.3390/biomedicines11030887>
11. Amjad A, Kordel P, Fernandes G (2023) A review on innovation in healthcare sector (telehealth) through artificial intelligence. *Sustainability* 15: 6655. <https://doi.org/10.3390/su15086655>
12. Bohr A, Memarzadeh K (2020) The rise of artificial intelligence in healthcare applications, *Artificial Intelligence in healthcare*, Elsevier, 25–60. <https://doi.org/10.1016/B978-0-12-818438-7.00002-2>
13. Haddad F (2022) 10 Applications of Intelligent Automation in Healthcare. [cited 2026 February 28]. Available from: <https://www.healthforce.ai/blog/10-applications-of-intelligent-automation-in-healthcare>.
14. Buchanan BG, Smith RG (1988) Fundamentals of expert systems. *Annu Rev Comput Sci* 3: 23–58. <https://doi.org/10.1146/annurev.cs.03.060188.000323>
15. Cheng Y, Wang F, Zhang P, et al. (2016) Risk prediction with electronic health records: A deep learning approach, *Proceedings of the 2016 SIAM international conference on data mining*, SIAM, 432–440. <https://doi.org/10.1137/1.9781611974348.49>
16. Elahi Naraqi A (2024) Virtual Medical Assistant Model Using Natural Language Processing in Healthcare System. Polytechnique Montréal.

17. Esteva A, Kuprel B, Novoa RA, et al. (2027) Dermatologist-level classification of skin cancer with deep neural networks. *Nature* 542: 115–118. <https://doi.org/10.1038/nature21056>
18. Song W, Jung SY, Baek H, et al. (2020) A predictive model based on machine learning for the early detection of late-onset neonatal sepsis: development and observational study. *JMIR Med Inform* 8: e15965. <https://doi.org/10.2196/15965>
19. Rajkomar A, Oren E, Chen K, et al. (2018) Scalable and accurate deep learning with electronic health records. *NPJ Digit Med* 1: 18. <https://doi.org/10.1038/s41746-018-0029-1>
20. Mihan A, Pandey A, Van Spall HG (2024) Mitigating the risk of artificial intelligence bias in cardiovascular care. *Lancet Digit Health* 6: e749–e754. [https://doi.org/10.1016/S2589-7500\(24\)00155-9](https://doi.org/10.1016/S2589-7500(24)00155-9)
21. K Natarajan, P Natarajan, S Muthusamy, et al. (2025) A novel method for bioinformatics analysis in gene expression profiling framework for personalized healthcare applications. *Multimed Tools Appl* 84: 24805–24837. <https://doi.org/10.1007/s11042-024-20066-5>
22. Ruamviboonsuk P, Tiwari R, Sayres R, et al. (2022) “Real-time diabetic retinopathy screening by deep learning in a multisite national screening programme: A prospective interventional cohort study. *Lancet Digit Health* 4: e235–e244. [https://doi.org/10.1016/S2589-7500\(22\)00017-6](https://doi.org/10.1016/S2589-7500(22)00017-6)
23. Jiang YJ, Xi XM, Jia HM, et al. (2021) Risk factors, clinical features and outcome of new-onset acute kidney injury among critically ill patients: a database analysis based on prospective cohort study. *BMC Nephrol* 22: 289. <https://doi.org/10.1186/s12882-021-02503-x>
24. Gu T, Pan W, Yu J, et al. (2025) Mitigating bias in AI mortality predictions for minority populations: a transfer learning approach. *BMC Med Inform Decis Mak* 25: 30. <https://doi.org/10.1186/s12911-025-02862-7>
25. Goswami SS, Mondal S (2026) Data Security and Privacy Concerns in Digitized Medical Services: Implications for Malpractice Risk Management. *Spectr Decis Mak Appl* 3: 212–242. <https://doi.org/10.31181/sdmap31202643>
26. Baraza X, Torrent-Sellens J (2026) Artificial intelligence and emerging risks in occupational safety and health. *Encyclopedia* 6: 25. <https://doi.org/10.3390/encyclopedia6010025>
27. Ottawa (ON): Canadian Agency for Drugs and Technologies in Health (2025) 2025 Watch List: Artificial Intelligence in Health Care: Health Technologies.
28. Zhang P, Shi J, Kamel Boulos MN (2024) Generative AI in medicine and healthcare: Moving beyond the ‘peak of inflated expectations’. *Future Internet* 16: 462. <https://doi.org/10.3390/fi16120462>
29. Hacking S (2024) ChatGPT and medicine: Together we embrace the AI renaissance. *JMIR Bioinform Biotechnol* 5: e52700. <https://doi.org/10.2196/52700>
30. Chen Y, Esmailzadeh P (2024) Generative AI in medical practice: In-depth exploration of privacy and security challenges. *J Med Internet Res* 26: e53008. <https://doi.org/10.2196/53008>
31. Bonis J (2025) Agentic AI in Healthcare: The Rise of Autonomous Clinical Assistants. [cited 2026 February 28]. Available from: <https://www.johnsnowlabs.com/agentic-ai-in-healthcare-the-rise-of-autonomous-clinical-assistants/>.
32. JL Cross, MA Choma, JA Onofrey (2024) Bias in medical AI: Implications for clinical decision-making. *PLOS Digit Health* 3: e0000651. <https://doi.org/10.1371/journal.pdig.0000651>
33. M Ghassemi, L Oakden-Rayner, AL Beam (2021) The false hope of current approaches to explainable artificial intelligence in health care. *Lancet Digit Health* 3: e745–e750. [https://doi.org/10.1016/S2589-7500\(21\)00208-9](https://doi.org/10.1016/S2589-7500(21)00208-9)

34. Fihn S, Saria S, Mendonça E, et al. (2019) Deploying AI in clinical settings, *Artificial intelligence in health care: The hope, the hype, the promise, the peril*, 145.
35. Henry KE, Kornfield R, Sridharan A, et al. (2022) “Human–machine teaming is key to AI adoption: clinicians’ experiences with a deployed machine learning system. *NPJ Digit Med* 5: 97. <https://doi.org/10.1038/s41746-022-00597-7>
36. Moor M, Banerjee O, Abad ZSH, et al. (2023) Foundation models for generalist medical artificial intelligence. *Nature* 616: 259–265. <https://doi.org/10.1038/s41586-023-05881-4>
37. Su Y, Lu Z, Liu J, et al. (2025) “Applications of large models in medicine. *arXiv preprint arXiv* 2502.17132.
38. Fahrner LJ, Chen E, Topol E, et al. (2025) The generative era of medical AI. *Cell* 188: 3648–3660. <https://doi.org/10.1016/j.cell.2025.05.018>.
39. Kruszynska M (2025) Closed Loop AI Systems pave the way for personalised medicine. [cited 2026 February 28]. Available from: <https://spyro-soft.com/blog/healthcare/closed-loop-ai-systems-pave-the-way-for-personalised-medicine>
40. Bajwa J, Munir U, Nori A, et al. (2021) Artificial intelligence in healthcare: transforming the practice of medicine. *Future Healthc J* 8: e188–e194. <https://doi.org/10.7861/fhj.2021-0095>
41. Thakur GK, Thakur A, Kulkarni S, et al. (2024) Deep learning approaches for medical image analysis and diagnosis. *Cureus* 16: e59507. <https://doi.org/10.7759/cureus.59507>
42. Khosravi M, Zare Z, Mojtabaeian SM, et al. (2024) Artificial intelligence and decision-making in healthcare: A thematic analysis of a systematic review of reviews. *Health Serv Res Mana* 11: 23333928241234863. <https://doi.org/10.1177/23333928241234863>
43. Al Kuwaiti A, Nazer K, Al-Reedy A, et al. (2023) A review of the role of artificial intelligence in healthcare. *J Pers Med* 13: 951. <https://doi.org/10.3390/jpm13060951>
44. Kandhare P, Kurlekar M, Deshpande T, et al. (2025) A review on revolutionizing healthcare technologies with AI and ML applications in pharmaceutical sciences. *Drugs Drug Cand* 4: 9. <https://doi.org/10.3390/ddc4010009>
45. Olawade DB, David-Olawade AC, Wada OZ, et al. (2024) Artificial intelligence in healthcare delivery: Prospects and pitfalls. *J Med Surg Public Health* 3: 100108. <https://doi.org/10.1016/j.glmedi.2024.100108>
46. Steerling E, Siira E, Nilsen P, et al. (2023) Implementing AI in healthcare—the relevance of trust: A scoping review. *Front Health Serv* 3: 1211150. <https://doi.org/10.3389/frhs.2023.1211150>
47. Sáez C, Ferri P, García-Gómez JM (2024) Resilient Artificial Intelligence in Health: Synthesis and Research Agenda Toward Next-Generation Trustworthy Clinical Decision Support. *J Med Internet Res* 26: e50295. <https://doi.org/10.2196/50295>
48. Hui AT, Ahn SS, Lye CT, et al. (2021) Ethical challenges of artificial intelligence in health care: A narrative review. *Ethics Biol Eng Med* 12: 55–71. <https://doi.org/10.1615/EthicsBiologyEngMed.2022041580>
49. Kandregula N (2025) Accelerating drug discovery with generative AI: A paradigm shift in pharmaceutical innovation and development. *World* 25: 1645–1658. <https://doi.org/10.30574/wjarr.2025.25.3.0845>
50. Malhotra N (2025) Is AI hype in drug development about to turn into reality? [cited 2026 February 28]. Available from: <https://codeblue.galencentre.org/2025/09/is-ai-hype-in-drug-development-about-to-turn-into-reality/>.

51. Awasthi R, Ramachandran SP, Mishra S, et al. (2025) Artificial Intelligence in Healthcare: 2024 Year in Review. *MedRxiv* 2025: 2025.02. 26.25322978. <https://doi.org/10.1101/2025.02.26.25322978>
52. Vani MS, Sudhakar RV, Mahendar A, et al. (2025) Personalized health monitoring using explainable AI: Bridging trust in predictive healthcare. *Sci Rep* 15: 31892. <https://doi.org/10.1038/s41598-025-15867-z>
53. Gujar P (2025) Data standardization and interoperability, in *Data usability in the enterprise: How usability leads to optimal digital experiences*, ed: Springer, 89–110. https://doi.org/10.1007/979-8-8688-1183-8_4
54. Akpobome O (2024) The impact of emerging technologies on legal frameworks: A model for adaptive regulation. *Int J Res Publ Rev* 5: 5046–5060. <https://doi.org/10.55248/gengpi.5.1024.3012>
55. Lotter W, Hassett MJ, Schultz N, et al. (2024) Artificial intelligence in oncology: current landscape, challenges, and future directions. *Cancer Discovery* 14: 711–726. <https://doi.org/10.1158/2159-8290.CD-23-1199>
56. Cellina M, CèM, AliM, et al. (2023) Digital twins: The new frontier for personalized medicine? *Appl Sci* 13: 7940. <https://doi.org/10.3390/app13137940>
57. J Singh, C Prabha, G Singh, et al. (2022) Addressing role of ambient intelligence and pervasive computing in today's world, *The International Conference on Recent Innovations in Computing*, 257–269. https://doi.org/10.1007/978-981-99-0601-7_20
58. De Almeida PGR, Dos Santos CD, Farias JS (2021) Artificial intelligence regulation: A framework for governance. *Ethics Inf Technol* 23: 505–525. <https://doi.org/10.1007/s10676-021-09593-z>
59. Odumbo OR (2025) Explainable AI and Federated Learning in Healthcare Supply Chain Intelligence: Addressing Ethical Constraints, Bias Mitigation, and Regulatory Compliance for Global Pharmaceutical Distribution. *Int J Comput Appl T* 14: 16–29.
60. Trovato MG (2025) The Myths We Still Believe About AI in Healthcare. [cited 2026 February 28]. Available from: <https://ai.northeastern.edu/news/ai-healthcare-myths>.
61. Sharma M (2024) The Impact of AI on Healthcare Jobs: Will Automation Replace Doctors. *Am J Data Min Knowl Discov* 9: 32–35. <https://doi.org/10.11648/j.ajdmkd.20240902.11>
62. Khan SJ (2025) Myth and Fact: Brief Reality of AI in Healthcare. [cited 2026 February 28]. Available from: <https://ezovion.com/myth-and-fact-brief-reality-of-ai-in-healthcare/>.
63. NIH/National Library of Medicine (2024) Risks and benefits of integrating AI into medical decision-making. [cited 2026 February 28]. Available from: <https://www.sciencedaily.com/releases/2024/07/240723123334.htm>.
64. Sahni N, Stein G, McKinsey O, et al. (2023) The potential impact of artificial intelligence on healthcare spending, *The economics of artificial intelligence*, MA, USA: University of Chicago Press, 49–86. <https://doi.org/10.3386/w30857>
65. Jeong J, Kim S, Pan L, et al. (2025) Reducing the workload of medical diagnosis through artificial intelligence: A narrative review. *Med* 104: e41470. <https://doi.org/10.1097/MD.00000000000041470>

66. E Nwachukwu, G Obasi, D Omotayo (2025) Data Privacy and Security Concerns in AI-Driven Healthcare. [cited 2026 February 28]. Available from: https://www.researchgate.net/profile/Ige-Gabriel/publication/394275058_Data_Privacy_and_Security_Concerns_in_AI-Driven_Healthcare/links/6890adf041943575bcb64df1/Data-Privacy-and-Security-Concerns-in-AI-Driven-Healthcare.pdf?__cf_chl_tk=yZxjvaSMFHVIBHa2YjaFDuYNv3.oXGwThFAIleoZdPc-1786431849-1.0.1.1-jtnPhfdn8txHnOFELnwsFtPT725JPd47oNkiJGCvc9Q.
67. BHM (2024) Part 2: Identify & Address Data Privacy and Security Concerns with AI System Enhancements. [cited 2026 February 28]. Available from: <https://bhmpc.com/2024/09/data-privacy-and-security-concerns-with-ai-system-enhancements/>.
68. McLean S, Read GJ, Thompson J, et al. (2023) The risks associated with Artificial General Intelligence: A systematic review. *J Exp Theor Artif In* 35: 649–663. <https://doi.org/10.1080/0952813X.2021.1964003>
69. Son C (2025) Why the future of AI in medicine is patient-facing. [cited 2026 February 28]. Available from: <https://kevinmd.com/2025/09/why-the-future-of-ai-in-medicine-is-patient-facing.html>.
70. A Bechar, R Medjoudj, Y Elmir, et al. (2025) Federated and transfer learning for cancer detection based on image analysis. *Neural Comput Appl* 37: 2239–2284. <https://doi.org/10.1007/s00521-024-10956-y>
71. Kaswan KS, Gaur L, Dhatteerwal JS, et al. (2021) AI-based natural language processing for the generation of meaningful information electronic health record (EHR) data, *Advanced AI techniques and applications in bioinformatics*, CRC Press, 41–86. <https://doi.org/10.1201/9781003126164-3>
72. Baynam G, Bowman F, Lister K, et al. (2017) Improved diagnosis and care for rare diseases through implementation of precision public health framework. *Rare diseases epidemiology: update and overview*, Advances in Experimental Medicine and Biology, 55–94. https://doi.org/10.1007/978-3-319-67144-4_4
73. Arora A (2025) Challenges of Integrating Artificial Intelligence in Legacy Systems and Potential Solutions for Seamless Integration. *Res J* 6: 44–51. <https://doi.org/10.2139/ssrn.5268176>
74. Liu F, Panagiotakos D (2022) Real-world data: a brief review of the methods, applications, challenges and opportunities. *BMC Med Res Methodol* 22: 287. <https://doi.org/10.1186/s12874-022-01768-6>
75. Abgrall G, Holder AL, Chelly Dagdia Z, et al. (2024) Should AI models be explainable to clinicians? *Critical Care* 28: 301. <https://doi.org/10.1186/s13054-024-05005-y>
76. Almadani B, Kaisar H, Thoker IR, et al. (2025) A systematic survey of distributed decision support systems in healthcare. *Systems* 13: 157. <https://doi.org/10.3390/systems13030157>
77. Norori N, Hu Q, Aellen FM, et al. (2021) Addressing bias in big data and AI for health care: A call for open science. *Patterns* 2: 100347. <https://doi.org/10.1016/j.patter.2021.100347>
78. Ueda D, Kakinuma T, Fujita S, et al. (2024) Fairness of artificial intelligence in healthcare: review and recommendations. *Jpn J Radiol* 42: 3–15. <https://doi.org/10.1007/s11604-023-01474-3>
79. Hurley ME, Lang BH, Kostick-Quenet KM, et al. (2025) Patient consent and the right to notice and explanation of AI systems used in health care. *Am J Bioethics* 25: 102–114. <https://doi.org/10.1080/15265161.2024.2399828>

80. Nolan P (2022) Artificial intelligence in medicine: how do we determine legal liability when things go wrong? Macquarie University. [cited 2026 February 28]. Available from: <https://figshare.mq.edu.au/ndownloader/files/39357701>.
81. Khalid N, Qayyum A, Bilal M, et al. (2023) Privacy-preserving artificial intelligence in healthcare: Techniques and applications. *Comput Biol Med* 158: 106848. <https://doi.org/10.1016/j.combiomed.2023.106848>
82. TK Gandhi, D Classen, CA Sinsky, et al. (2023) How can artificial intelligence decrease cognitive and work burden for front line practitioners? *JAMIA Open* 6: ooad079. <https://doi.org/10.1093/jamiaopen/oad079>
83. Asgari E, Kaur J, Nuredini G, et al. (2024) Impact of electronic health record use on cognitive load and burnout among clinicians: narrative review. *JMIR Med Inform* 12: e55499. <https://doi.org/10.2196/55499>
84. Briceño-Garmendia C, Estache A (2004) Infrastructure services in developing countries: Access, quality, costs, and policy reform, World Bank Publications, 3468. <https://doi.org/10.1596/1813-9450-3468>
85. Arah OA, Klazinga NS, Delnoij DM, et al. (2003) Conceptual frameworks for health systems performance: A quest for effectiveness, quality, and improvement. *Int J Qual Health Care* 15: 377–398. <https://doi.org/10.1093/intqhc/mzg049>
86. Fargen KM, Frei D, Fiorella D, et al. (2013) The FDA approval process for medical devices: an inherently flawed system or a valuable pathway for innovation? *J Neurointerv Surg* 5: 269–275. <https://doi.org/10.1136/neurintsurg-2012-010400>
87. Yousefi F, Dehnavieh R, Laberge M, et al. (2025) Opportunities, challenges, and requirements for Artificial Intelligence (AI) implementation in Primary Health Care (PHC): a systematic review. *BMC Primary Care* 26: 196. <https://doi.org/10.1186/s12875-025-02785-2>
88. Stroud AM, Anzabi MD, Wise JL, et al. (2025) Toward safe and ethical implementation of health care artificial intelligence: Insights from an academic medical center. *Mayo Clin Proc Digit Health* 3: 100189. <https://doi.org/10.1016/j.mcpdig.2024.100189>
89. Bartlett B (2024) Towards accountable, legitimate and trustworthy AI in healthcare: Enhancing AI ethics with effective data stewardship. *New Bioethics* 30: 285–309. <https://doi.org/10.1080/20502877.2025.2482282>
90. Mennella C, Maniscalco U, De Pietro G, et al. (2024) Ethical and regulatory challenges of AI technologies in healthcare: A narrative review. *Heliyon* 10: e26297. <https://doi.org/10.1016/j.heliyon.2024.e26297>
91. Li YH, Li YL, Wei MY, et al. (2024) Innovation and challenges of artificial intelligence technology in personalized healthcare. *Sci Rep* 14: 18994. <https://doi.org/10.1038/s41598-024-70073-7>



AIMS Press

© 2026 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0>)