

# AI in Healthcare: Mini-Review of AI Transforming Healthcare Globally & Ethically

Dr. S. S. Begum<sup>1,\*</sup>, Dr. S. Manham<sup>2</sup>

<sup>1</sup>Department of Chemistry, Gargaon College, Simaluguri, Assam: 785686, India

<sup>2</sup>Department of History, Gargaon College, Simaluguri, Assam: 785686, India

DOI: <https://doi.org/10.51244/IJRSI.2025.120800254>

Received: 23 Sep 2025; Accepted: 29 Sep 2025; Published: 03 October 2025

## ABSTRACT

Artificial Intelligence (AI) is no longer just a futuristic concept—it has become a trusted partner in transforming healthcare around the world. Today, AI quietly works alongside doctors, nurses, and healthcare teams, streamlining everything from diagnosing illnesses to managing hospital operations. In clinics and hospitals, AI-powered tools analyze medical images, genetic data, and patient histories with remarkable speed and accuracy. This means diseases like cancer, heart conditions, and neurological disorders can often be detected far earlier than before, giving patients a much better chance at successful treatment. For example, AI-assisted radiology can flag unusual patterns in X-rays, MRIs, or CT scans in just seconds, helping doctors make faster, more confident decisions. In everyday primary care, AI acts like a digital co-pilot—suggesting tests, offering evidence-based treatment options, and even pulling in data from wearable devices or electronic health records to personalize care. One of the biggest breakthroughs in recent times is AI's role in personalized medicine. By combining genetic information with lifestyle and medical history, AI helps design treatment plans tailored to each patient's unique needs. AI has become a game changer for drug discovery and clinical trials. It can simulate how molecules interact, identify promising treatments, and even suggest new uses for existing medicines—speeding up the process of getting life-saving drugs to market, thus enabling a more proactive, precise, and compassionate healthcare system.

## INTRODUCTION

Artificial Intelligence (AI) is no longer just a futuristic concept—it has become a trusted partner in transforming healthcare around the world. Today, AI quietly works alongside doctors, nurses, and healthcare teams, streamlining everything from diagnosing illnesses to managing hospital operations (Topol, 2019). In clinics and hospitals, AI-powered tools analyze medical images, genetic data, and patient histories with remarkable speed and accuracy (Esteva et al., 2019). This means diseases like cancer, heart conditions, and neurological disorders can often be detected far earlier than before, giving patients a much better chance at successful treatment. For example, AI-assisted radiology can flag unusual patterns in X-rays, MRIs, or CT scans in just seconds, helping doctors make faster, more confident decisions (McKinney et al., 2020). In everyday primary care, AI acts like a digital co-pilot—suggesting tests, offering evidence-based treatment options, and even pulling in data from wearable devices or electronic health records to personalize care. One of the biggest breakthroughs in recent times is AI's role in personalized medicine. By combining genetic information with lifestyle and medical history, AI helps design treatment plans tailored to each patient's unique needs. AI has become a game changer for drug discovery and clinical trials. It can simulate how molecules interact, identify promising treatments, and even suggest new uses for existing medicines—speeding up the process of getting life-saving drugs to market, thus enabling a more proactive, precise, and compassionate healthcare system (Stokes et al., 2020). The integration of Artificial Intelligence (AI) across various sectors has profoundly reshaped methodologies and outcomes, with its application in healthcare emerging as a particularly transformative domain (Jiang et al., 2017). AI has rapidly evolved into one of the most influential technologies driving innovation in healthcare. Broadly defined, AI refers to computer systems designed to perform tasks that typically require human intelligence, such as learning, reasoning, problem-solving, and decision-making. In the healthcare context, AI systems can process large volumes of structured and unstructured medical data,

identify complex patterns, and provide evidence-based insights to support clinical and administrative decision-making (Rajkomar et al., 2019). This capacity has opened new possibilities for improving diagnosis accuracy, predicting disease progression, personalizing treatment plans, and enhancing operational efficiency within healthcare institutions (He et al., 2019). The integration of AI into healthcare spans multiple domains. In clinical diagnostics, AI-powered algorithms are being used for image recognition in radiology, pathology, and dermatology, enabling early detection of diseases such as cancer and diabetic retinopathy with accuracy comparable to, or in some cases exceeding, human experts (Litjens et al., 2017). In predictive analytics, machine learning models are applied to electronic health records (EHRs) to forecast patient outcomes, hospital readmissions, and potential complications, thus allowing proactive intervention (Shickel et al., 2018). AI also supports precision medicine by analyzing genomic data and tailoring therapies to individual patient profiles (Kononenko, 2001). Beyond clinical care, AI is streamlining administrative tasks, including patient triage, appointment scheduling, and medical coding, thereby reducing the workload on healthcare staff and optimizing resource allocation (Bohr & Memarzadeh, 2020).

Despite these advancements, the adoption of AI in healthcare is not without challenges. Data-related issues, such as incomplete records, poor interoperability, and limited access to high-quality annotated datasets, can reduce the reliability of AI models (Kelly et al., 2019). Ethical concerns, including algorithmic bias, data privacy, and the lack of transparency in decision-making (“black box” problem), also present significant barriers to trust and acceptance (Price & Cohen, 2019). Furthermore, regulatory frameworks for AI in healthcare are still evolving, leading to uncertainty in clinical deployment (FDA, 2021). Building clinician confidence through explainable AI, rigorous validation, and continuous monitoring is essential to ensure patient safety and ethical compliance (Amann et al., 2020). Artificial intelligence is beginning to touch almost every part of healthcare, from diagnosis to day-to-day hospital management. In clinical care, AI tools are now being used to read X-rays, CT scans, pathology slides, and skin images. For example, Google Health’s AI model for diabetic retinopathy screening has demonstrated performance comparable to ophthalmologists in detecting early signs of the disease (Gulshan et al., 2016), while IBM Watson for Oncology has been used to suggest personalized cancer treatment options based on a patient’s genetic profile and medical history. In many cases, these systems can spot diseases at an early stage with accuracy similar to, or sometimes better than, experienced doctors. Beyond diagnosis, machine learning applied to electronic health records helps predict who is at risk of complications, hospital readmission, or poor treatment response, allowing doctors to intervene earlier and tailor therapies more precisely to individual patients (Rajkomar et al., 2018). AI is also speeding up drug discovery and medical research, helping researchers identify promising compounds faster than traditional methods (Zhavoronkov et al., 2019). Virtual consultations and rehabilitation tools powered by AI are making care more efficient and accessible—even in remote or underserved areas (Wahl et al., 2018). Yet the road to full integration of AI in healthcare is not without problems. Medical data is far more complicated and messier than data in other fields. Patients differ in their genetics, lifestyles, and health conditions, and even doctors may treat the same disease differently. This makes it very hard for AI systems to generalize reliably across different hospitals or patient groups (Zech et al., 2018). Ethical concerns add another layer of difficulty. Many AI systems are “black boxes,” meaning they provide results without clearly showing how those results were reached. This lack of transparency raises fears of bias, unfair treatment, and errors that could harm patients, while also making it difficult for regulators to approve and monitor these systems (Rudin, 2019; Obermeyer et al., 2019). Trust is therefore central. Doctors and patients alike need to feel confident that AI systems are accurate, fair, and accountable. Recent initiatives, such as the FUTURE-AI guidelines, suggest principles like fairness, robustness, and usability to make AI safer and easier to trust (Lekadir et al., 2022). But trust alone is not enough. There are also social and economic barriers. Not every hospital has access to advanced AI systems, and underserved areas risk being left further behind due to lack of digital infrastructure (The Lancet Global Health, 2020). If AI is to deliver on its promise, it will require not just better technology but also careful attention to ethics, regulation, and equity, ensuring that its benefits reach everyone—not just the privileged few (Panch et al., 2018).

## Aims of This Review

This review seeks to explore how AI contributes to healthcare by examining:

### 1. Input and Processing: The nature and quality of medical data leveraged by AI systems.

2. Applications: The roles AI plays across diagnostics, prognostics, administration, and research.
3. Challenges & Ethical Dimensions: Technical limitations, data biases, ethical dilemmas, trust barriers, and policy frameworks.
4. Future Directions: Strategies for developing transparent, equitable, and effective AI systems, including interdisciplinary collaboration and improved governance.

By weaving together technical insights, ethical considerations, and real-world challenges, this review aims to provide a cohesive, realistic overview of AI's contributions and limitations in healthcare—informing both present practice and future innovations. Of course, with this rapid growth come challenges. Protecting patient privacy, avoiding algorithmic bias, and ensuring that AI systems are transparent and trustworthy remain top priorities (GDPR, 2016; FDA, 2021). Governments, healthcare organizations, and tech companies are working together to set ethical standards and create safeguards that keep patients safe while embracing innovation. Ultimately, AI in 2025 is not here to replace healthcare professionals—it's here to empower them. By handling repetitive tasks, analyzing complex data, and delivering insights instantly (Topol, 2019).

This review comprehensively examines the current state and future prospects of AI in healthcare, focusing on its applications, challenges, and ethical considerations. Specifically, AI's potential to revolutionize healthcare stems from its capacity to enhance diagnostics, personalize treatment regimens, and significantly improve operational efficiencies within clinical environments (Jiang et al., 2017). The ability of AI to analyze vast and intricate datasets, including medical imaging and electronic health records, enables more accurate and timely diagnoses, thereby improving the quality and efficiency of healthcare decision-making (Rajkomar et al., 2019). The healthcare industry has witnessed a significant surge in AI adoption due to its potential to enhance service delivery and operational efficiency, although uncertainties persist regarding its practical effectiveness and value (He et al., 2019). Despite the employment of several useful technologies in healthcare, AI is not yet widely deployed, and its algorithms often remain opaque, presenting challenges in understanding their decision-making processes (Rudin, 2019). This complexity necessitates the development of more transparent and interpretable AI models, allowing healthcare professionals to comprehend the rationale behind AI-driven recommendations and foster greater confidence in their clinical utility (Amann et al., 2020). This opacity also necessitates careful consideration of the ethical implications associated with AI deployment in sensitive healthcare contexts, particularly regarding patient safety, data privacy, and accountability (Price & Cohen, 2019). It investigates the 'why, how, and when' of XAI model usage and their implications, aiming to formalize the XAI field and detail how trustworthy AI can be developed for healthcare (Adadi & Berrada, 2018). A systematic review, conducted in accordance with PRISMA guidelines, examined studies published between 2012 and 2022 that applied explainable artificial intelligence (XAI) for patient screening and diagnosis. By restricting the scope to English-language records that aligned with the PICO framework, the review provided a structured synthesis of research comparing XAI-based models with conventional diagnostic approaches (Angraal et al., 2020). The same paper identified several methodological approaches—such as dimension reduction, feature selection, attention mechanisms, knowledge distillation, and surrogate representations—as central to the development of explainable models in medicine. The study emphasized that explainability is not merely a technical preference but a prerequisite for clinical adoption, since interpretable models are perceived as more trustworthy by medical practitioners. This argument resonates with the broader consensus in the field: Ghassemi et al. (2021), for instance, similarly contend that the clinical value of AI depends less on predictive accuracy alone and more on the ability of systems to provide justifiable explanations that clinicians can act upon. At the same time, the paper's findings are consistent with other critiques of the field regarding persistent technical and operational challenges. Issues such as system performance, security, evaluation of explanations, and generalization are widely acknowledged limitations of current XAI frameworks. For example, Tonekaboni et al. (2019) observed that while XAI techniques such as feature attribution provide some interpretability, they often lack consistency across models and datasets, thereby undermining their reliability in clinical decision-making. The review adds weight to this concern by noting that the evaluation of explanations remains underdeveloped, particularly in healthcare contexts where the stakes are high. In addition to technical barriers, the paper also highlights practical concerns such as high false-positive rates, bias, privacy risks, and limited transparency. These concerns parallel the arguments of Amann et al. (2020), who stress that XAI applications risk perpetuating structural inequities in healthcare if bias in training data is not adequately addressed.

However, the review is more critical in noting that despite widespread enthusiasm, most current applications remain “strong on promise but rather lacking in evidence and demonstration.” This skepticism aligns with recent systematic reviews, which caution against the premature clinical deployment of XAI without rigorous validation and external benchmarking (Angraal et al., 2020). Where the analysis makes a distinctive contribution is in highlighting operational and organizational obstacles, such as legal and socio-relational issues and communication barriers between AI developers and healthcare practitioners. While these themes are less emphasized in the predominantly technical literature, other scholars have begun to echo similar concerns. For instance, Cabitza et al. (2017) argue that institutional readiness, professional trust, and medico-legal accountability are as critical as algorithmic performance in determining whether XAI can be integrated effectively into healthcare systems.

Taken together, the literature suggests that while there is broad agreement on the necessity of explainability for trustworthy AI in medicine, there is less consensus on how to achieve it effectively. This review reflects the broader tension in the field: the technological advancements in XAI are significant, but their clinical utility remains under-validated. Compared with other studies, the assessment is more cautious, stressing the urgent need for regulatory frameworks, systematic validation, and improved data quality reporting. Without these, the promise of XAI in healthcare risks being overshadowed by its limitations (Rudin, 2019).

## DISCUSSION

For AI systems to function effectively and equitably in healthcare, they require access to large, diverse, high-quality, and standardized datasets. Limitations such as incompleteness, inconsistency, and bias in medical data directly undermine model performance and increase the risk of patient harm (Zech et al., 2018). Clinical information—including administrative records, diagnostic images, laboratory results, and patient demographics—is frequently dispersed across multiple platforms and institutions. As a result, extensive processes of cleaning, standardization, and normalization are essential before such data can be reliably used for AI training and deployment. A foundational limitation of medical AI—and one that conditions every downstream claim about explainability, fairness, and clinical utility—is the provenance and quality of the training data. Large, labelled image sets have powered early clinical breakthroughs: for example, CheXNet trained on ChestX-ray14 demonstrated radiologist-level performance for several chest pathologies (Rajpurkar et al., 2017), and deep CNNs trained on large dermoscopy collections have shown dermatologist-level performance for melanoma detection (Esteva et al., 2017). These successes illustrate how scale and curated labels can produce impressive predictive accuracy. However, multiple studies show that dataset artifacts and label choices frequently act as hidden shortcuts that undermine model validity in new settings. Models trained on hospital billing or cost proxies (rather than direct clinical need) can reproduce—and amplify—systemic disparities: Obermeyer et al. (2019) demonstrated that an algorithm used to allocate care systematically underestimated the clinical needs of Black patients because it used healthcare cost as a proxy for illness. This example shows how a data choice (proxy label) can produce seemingly high-performing models that nonetheless perpetuate inequity. “Datasheets for datasets” and “Model Cards” are widely recommended practices to disclose dataset composition, labeling processes, known limitations, and evaluation statistics—practices intended to make input provenance explicit and interpretable for clinicians and regulators (Gebru et al., 2021; Mitchell et al., 2019). Adoption remains partial, however, and many clinical datasets lack standardized metadata about selection bias, demographic coverage, and preprocessing steps—gaps that limit reproducibility and fair deployment.

AI applications in healthcare cluster into four practical domains: diagnostics, prognostics, administration and research. Landmark papers have shown strong performance in image-based diagnosis: Rajpurkar et al.’s (2017) CheXNet for chest X-rays and Esteva et al.’s (2017) skin-lesion classifier are canonical examples where deep learning matched or exceeded specialists on held-out test sets. These results catalyzed optimism about point-of-care decision support and automating screening tasks.

Prognostics (risk prediction and early warning): AI has been applied to sepsis prediction, readmission risk, and deterioration forecasting. Some institution-level deployments show improved triage speed, yet prognostic models often suffer from high false-positive rates and limited temporal generalizability, especially when surveillance practices or clinical workflows change (Shickel et al., 2018). Moreover, prognostic labels can



themselves be noisy (e.g., outcome definitions vary), complicating both model training and the interpretability of model outputs. Predictive models are used to optimize scheduling, resource allocation, and billing workflows. These applications can yield operational efficiencies but also raise distinct fairness concerns—when administrative proxies correlate with patient disadvantage, optimization can worsen inequities, as seen in the Obermeyer et al. (2019) example where cost proxies misallocated care. In research and drug discovery, AI accelerates candidate screening, target identification, and retrospective pattern mining in large observational datasets; successes are notable in accelerating hypothesis generation (Stokes et al., 2020). Yet translational gaps persist: many in-silico leads fail in biological validation, and black-box models make it harder for bench scientists and clinicians to interpret why a candidate was prioritized. This strengthens the argument for XAI methods that provide mechanistic or feature-level rationales (Jiménez-Luna et al., 2020).

**Challenges & ethical dimensions: technical limitations, bias, explainability, trust and policy:** The challenges of explainable AI (XAI) in healthcare can be understood across technical, ethical, and legal dimensions. Technically, many widely used XAI tools—such as saliency maps, feature attributions, and surrogate models—offer only post-hoc explanations that often appear convincing but can be unstable, inconsistent across models, or even misleading for individual patients. Ghassemi et al. (2021) caution that such methods risk creating a false sense of understanding, giving clinicians the impression that the AI is reasoning like them when it is not. Tonekaboni et al. (2019) similarly note that clinicians prefer explanations tied to medically meaningful features, yet most current approaches fail to provide this level of clarity, highlighting the need for evaluation frameworks that measure explanation fidelity and clinical usefulness rather than mere visual appeal. Ethical concerns center on bias and fairness, with Obermeyer et al. (2019) demonstrating how a widely used care allocation algorithm underestimated the needs of Black patients because it relied on healthcare costs as a proxy for illness. Scholars such as Amann et al. (2020) argue that bias can emerge from missing variables, flawed labels, or unrepresentative data, problems that explainability alone cannot address without deliberate strategies like relabeling, fairness constraints, or targeted data collection. Privacy and security present another major obstacle, as sensitive medical data carries high risks of re-identification and leakage. While methods such as federated learning and differential privacy aim to mitigate these risks, they often reduce accuracy and add engineering complexity, and how to generate explanations for such models remains an open question (Kaissis et al., 2020). Trust and adoption also pose challenges: empirical studies show that clinicians use explanations to justify, audit, and learn from AI, but when outputs do not align with their judgment, they may either disregard the model or over-rely on it, both of which can compromise patient safety (Cabitza et al., 2017). Finally, legal and regulatory issues remain unresolved, particularly around liability and accountability in cases of AI-related harm. Although frameworks like the EU AI Act are beginning to require transparency, documentation, and human oversight for high-risk AI systems in healthcare, practical standards for evaluating explanation quality are still lacking (European Commission, 2021). Overall, the literature agrees that explainability cannot replace rigorous data curation, bias mitigation, privacy safeguards, and strong governance; instead, XAI must evolve alongside robust regulation, trustworthy data practices, and organizational readiness to be genuinely effective and safe in clinical contexts (Price & Cohen, 2019).

**Future directions: pathways to transparent, equitable, and effective AI systems:** Future directions for explainable AI (XAI) in healthcare revolve around practical pathways that can enhance transparency, equity, and clinical effectiveness. A key lever lies in better dataset and model documentation. Proposals such as Model Cards and Datasheets for Datasets provide structured templates to disclose intended uses, data composition, evaluation protocols, and failure modes (Mitchell et al., 2019; Gebru et al., 2021). Their adoption in medical AI, where opacity and hidden assumptions are common, would enable external auditing and informed clinical use. Empirical studies suggest that access to such documentation improves practitioners' ability to judge applicability and risks, yet implementation in healthcare lags behind other domains. Equally pressing is the need for rigorous, clinically-oriented evaluation of explanations. Current practice too often settles for visually plausible heatmaps or saliency maps, but the literature calls for metrics that capture fidelity to the model's actual reasoning, robustness across perturbations, and measurable utility for clinical outcomes. Prospective user studies with clinicians, as recommended by Ghassemi et al. (2021), are especially critical to move beyond theoretical benchmarks toward evidence of real-world benefit. Addressing upstream data biases remains another priority. Obermeyer et al.'s (2019) widely cited work demonstrates how flawed label selection and narrow data sampling can systematically disadvantage vulnerable populations. Future systems must prioritize

clinically meaningful outcomes, diverse sampling strategies, and fairness audits, making bias mitigation a routine rather than exceptional practice. In addition, scholars emphasize interdisciplinary and human-centered development. Studies by Tonekaboni et al. (2019) and others show that clinician and patient involvement from problem framing through deployment not only clarifies what constitutes a “useful” explanation but also ensures that outputs are presented in workflow-compatible ways. On the governance front, emerging regulation such as the EU AI Act is steering the field toward mandatory documentation, human oversight, and continuous monitoring (European Commission, 2021). Yet literature cautions that compliance alone is insufficient without operational mechanisms like incident reporting, post-market surveillance, and ongoing validation, which are still underdeveloped in clinical AI practice. Finally, technical research priorities remain unresolved: how to design explanation methods with provable fidelity and uncertainty quantification, how to generate faithful explanations for federated or privacy-preserving models, and how to account for longitudinal challenges such as deskilling or distribution shifts across hospitals. Mixed-method reviews highlight that without sustained clinician training and role adaptation, XAI could inadvertently undermine rather than strengthen medical expertise (Cabitza et al., 2017). Taken together, the literature suggests that progress requires an integrated approach: transparent documentation, clinician-centered evaluation, data curation, interdisciplinary co-design, and strong governance frameworks. Only by aligning technical advances with regulatory standards and organizational readiness can healthcare AI bridge the persistent “promise versus evidence” gap that Rudin (2019) and others critique.

## CONCLUSION

Artificial intelligence is no longer a futuristic add-on in medicine but a transformative force reshaping how care is delivered, managed, and experienced. Its impact is visible from precision oncology—where therapies are fine-tuned to target specific genetic mutations, improving effectiveness while minimizing side effects—to the use of natural language processing that helps clinicians navigate the overwhelming volume of medical literature and patient data. Beyond the exam room, AI drives operational efficiency: intelligent scheduling reduces wait times, predictive analytics anticipate staffing and supply needs, and chatbots handle routine queries so that healthcare professionals can focus on urgent, human-centered care. Crucially, AI is also narrowing access gaps. In rural or underserved regions, AI-enabled remote monitoring keeps doctors connected to patients in real time, while in mental health, conversational agents and emotion-recognition tools provide timely support between therapy sessions. These developments show that what was once a high-tech luxury has become an essential part of global health. Yet, as the literature highlights, the promise of AI must be balanced with ethical safeguards, rigorous validation, and governance frameworks that ensure safety, equity, and trust. Moving forward, the challenge is not simply building more powerful models but embedding them responsibly within health systems so that patients and providers alike benefit from AI’s full potential.

## REFERENCES

1. Adadi, A., & Berrada, M. (2018). Peeking inside the black-box: A survey on explainable artificial intelligence (XAI). *IEEE Access*, 6, 52138-52160.
2. Amann, J., Blasimme, A., Vayena, E., Frey, D., & Madai, V. I. (2020). Explainability for artificial intelligence in healthcare: A multidisciplinary perspective. *BMC Medical Informatics and Decision Making*, 20(1), 310.
3. Angraal, S., Krumholz, H. M., & Schulz, W. L. (2020). Blockchain technology: Applications in healthcare. *Circulation: Cardiovascular Quality and Outcomes*, 13(7), e006025.
4. Bohr, A., & Memarzadeh, K. (2020). The rise of artificial intelligence in healthcare applications. In *Artificial Intelligence in Healthcare* (pp. 25-60). Academic Press.
5. Cabitza, F., Rasoini, R., & Gensini, G. F. (2017). Unintended consequences of machine learning in medicine. *JAMA*, 318(6), 517-518.
6. European Commission. (2021). Proposal for a regulation of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). COM(2021) 206 final.

7. Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639), 115-118.
8. Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., ... & Dean, J. (2019). A guide to deep learning in healthcare. *Nature Medicine*, 25(1), 24-29.
9. Food and Drug Administration (FDA). (2021). Artificial intelligence and machine learning (AI/ML) software as a medical device. U.S. Department of Health and Human Services.
10. Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumée III, H., & Crawford, K. (2021). Datasheets for datasets. *Communications of the ACM*, 64(12), 86-92.
11. General Data Protection Regulation (GDPR). (2016). Regulation (EU) 2016/679 of the European Parliament and of the Council. *Official Journal of the European Union*, L119, 1-88.
12. Ghassemi, M., Oakden-Rayner, L., & Beam, A. L. (2021). The false hope of current approaches to explainable artificial intelligence in health care. *The Lancet Digital Health*, 3(11), e745-e750.
13. Gulshan, V., Peng, L., Coram, M., Stumpe, M. C., Wu, D., Narayanaswamy, A., ... & Webster, D. R. (2016). Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA*, 316(22), 2402-2410.
14. He, J., Baxter, S. L., Xu, J., Xu, J., Zhou, X., & Zhang, K. (2019). The practical implementation of artificial intelligence technologies in medicine. *Nature Medicine*, 25(1), 30-36.
15. Jiménez-Luna, J., Grisoni, F., & Schneider, G. (2020). Drug discovery with explainable artificial intelligence. *\*Nature Machine Intelligence*, 2(10), 573-584.
16. Jiang, F., Jiang, Y., Zhi, H., Dong, Y., Li, H., Ma, S., ... & Wang, Y. (2017). Artificial intelligence in healthcare: Past, present and future. *Stroke and Vascular Neurology*, 2(4), 230-243.
17. Kaissis, G. A., Makowski, M. R., Rückert, D., & Braren, R. F. (2020). Secure, privacy-preserving and federated machine learning in medical imaging. *Nature Machine Intelligence*, 2(6), 305-311.
18. Kelly, C. J., Karthikesalingam, A., Suleyman, M., Corrado, G., & King, D. (2019). Key challenges for delivering clinical impact with artificial intelligence. *BMC Medicine*, 17(1), 195.
19. Kononenko, I. (2001). Machine learning for medical diagnosis: History, state of the art and perspective. *Artificial Intelligence in Medicine*, 23(1), 89-109.
20. The Lancet Global Health. (2020). Artificial intelligence for global health. *The Lancet Global Health*, 8(7), e875.
21. Lekadir, K., Feragen, A., Fofanah, A. J., Frangi, A. F., Buyx, A., Emelie, A., ... & Sanz, J. (2022). FUTURE-AI: Guiding principles and consensus recommendations for trustworthy artificial intelligence in medical imaging. *arXiv preprint arXiv:2209.02435*.
22. Litjens, G., Kooi, T., Bejnordi, B. E., Setio, A. A. A., Ciompi, F., Ghafoorian, M., ... & Sánchez, C. I. (2017). A survey on deep learning in medical image analysis. *Medical Image Analysis*, 42, 60-88.
23. McKinney, S. M., Sieniek, M., Godbole, V., Godwin, J., Antropova, N., Ashrafi, H., ... & Shetty, S. (2020). International evaluation of an AI system for breast cancer screening. *Nature*, 577(7788), 89-94.
24. Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., ... & Gebru, T. (2019). Model cards for model reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 220-229).
25. Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447-453.
26. Panch, T., Mattie, H., & Celi, L. A. (2019). The “inconvenient truth” about AI in healthcare. *NPJ Digital Medicine*, 2\*(1), 77.
27. Price, W. N., & Cohen, I. G. (2019). Privacy in the age of medical big data. *Nature Medicine*, 25(1), 37-43.
28. Rajkomar, A., Dean, J., & Kohane, I. (2019). Machine learning in medicine. *New England Journal of Medicine*, 380(14), 1347-1358.
29. Rajkomar, A., Oren, E., Chen, K., Dai, A. M., Hajaj, N., Hardt, M., ... & Sundberg, P. (2018). Scalable and accurate deep learning with electronic health records. *NPJ Digital Medicine*, 1\*(1), 18.
30. Rajpurkar, P., Irvin, J., Zhu, K., Yang, B., Mehta, H., Duan, T., ... & Lungren, M. P. (2017). CheXNet: Radiologist-level pneumonia detection on chest X-rays with deep learning. *\*arXiv preprint arXiv:1711.05225*.

31. Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206-215.
32. Shickel, B., Tighe, P. J., Bihorac, A., & Rashidi, P. (2018). Deep EHR: A survey of recent advances in deep learning techniques for electronic health record (EHR) analysis. *IEEE Journal of Biomedical and Health Informatics*, 22(5), 1589-1604.
33. Stokes, J. M., Yang, K., Swanson, K., Jin, W., Cubillos-Ruiz, A., Donghia, N. M., ... & Collins, J. J. (2020). A deep learning approach to antibiotic discovery. *Cell*, 180(4), 688-702.
34. Tonekaboni, S., Joshi, S., McCradden, M. D., & Goldenberg, A. (2019). What clinicians want: Contextualizing explainable machine learning for clinical end use. *Proceedings of Machine Learning Research*, 106, 359-380.
35. Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine*, 25(1), 44-56.
36. Wahl, B., Cossy-Gantner, A., Germann, S., & Schwalbe, N. R. (2018). Artificial intelligence (AI) and global health: How can AI contribute to health in resource-poor settings? *BMJ Global Health*, 3(4), e000798.
37. Zech, J. R., Badgeley, M. A., Liu, M., Costa, A. B., Titano, J. J., & Oermann, E. K. (2018). Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study. *PLOS Medicine*, 15(11), e1002683.
38. Zhavoronkov, A., Ivanenkov, Y. A., Aliper, A., Veselov, M. S., Aladinskiy, V. A., Aladinskaya, A. V., ... & Aspuru-Guzik, A. (2019). Deep learning enables rapid identification of potent DDR1 kinase inhibitors. *Nature Biotechnology*, 37(9), 1038-1040.