International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 5, Sep-Oct-2026, ISSN (Online): 2395-566X

Artificial Intelligence in Healthcare: A Review of Diagnostic Imaging, Drug Discovery, Personalized Medicine and Responsible Adoption in India

Alfiya Mehboob Kazi[1] Prof. S. K. Said[2] Prof. M. C. Kolse[3]
Department Of Artificial Intelligence And Data Science Jaihind College Of Engineering Pune, Maharashtra, India

Abstract—Artificial Intelligence (AI) is reshaping healthcare by accelerating diagnosis, shortening parts of the drug-development cycle and making treatment decisions more individual. This review consolidates evidence from four connected areas: medical imaging, drug discovery, personalized medicine, and the ethical and regulatory conditions under which these tools can be trusted. Imaging is the most mature domain, where deep convolutional networks now detect abnormalities, triage urgent studies and improve reconstruction, while pathology, cardiology and neurology are rapidly following. In drug discovery, learning-based models help identify targets, generate candidate molecules and select trial participants, although gains are concentrated in early pipeline stages. Personalized medicine uses electronic health records and multi-omics data to estimate individual treatment response. Alongside these benefits, the paper analyses risks concerning patient privacy, dataset bias, opacity and accountability, and maps them to the Indian context of the Ayushman Bharat Digital Mission, the Digital Personal Data Protection Act 2023 and national AI strategy documents. A layered reference architecture for responsible clinical AI is proposed, together with a research gap analysis and objectives. The review concludes that healthcare AI should be judged by measurable improvement in patient outcomes, supported by prospective validation, secure data practice, fairness auditing and human oversight.

Index Terms—Artificial intelligence, healthcare, medical imaging, drug discovery, personalized medicine, clinical decision support, data privacy, health ethics

I. INTRODUCTION

Health systems around the world are under simultaneous pressure from ageing populations, the growing burden of chronic disease and a shortage of specialist clinicians. India illustrates this tension sharply: a large share of specialists practise in cities, while most of the population lives in smaller towns and villages where diagnostic delays are common. At the same time, hospitals now generate digital data at an unprecedented rate in the form of radiology studies, laboratory values, prescriptions and clinical notes. Artificial Intelligence (AI), and in particular deep learning [1], offers a way to turn this data into timely clinical insight.

Landmark studies have shown that carefully trained models can match specialists on narrowly defined tasks such as grading diabetic retinopathy [2] and classifying skin lesions [3]. However, laboratory accuracy alone does not guarantee benefit at the bedside. Topol [4] argues that the real value of AI lies in a partnership in which machines handle pattern recognition at scale and clinicians retain judgement, empathy and accountability.

This paper reviews the application of AI across the healthcare value chain and makes three contributions: (i) a structured survey of imaging, drug discovery and personalized medicine literature; (ii) an analysis of ethical, legal and regulatory requirements with emphasis on India; and (iii) a layered reference architecture that links data governance, model development and clinical deployment. The following sections develop each of these contributions in turn.

II. BACKGROUND AND FUNDAMENTAL CONCEPTS

2.1 Learning Paradigms In Clinical Ai

Supervised learning maps labelled inputs, such as a chest radiograph and its diagnosis, to outputs. Convolutional Neural Networks (CNNs) and residual networks [5] dominate image tasks, while encoder–decoder designs like U-Net [6] are standard for segmenting organs and lesions. Sequential and tabular records are handled with gradient-boosted trees or recurrent and transformer models [7]. More recently, large language models have shown that clinical knowledge can be encoded in general-purpose text models [8].

2.2 Healthcare Data Sources

Clinical AI draws on imaging archives (PACS), Electronic Health Records (EHR), genomics, wearables and claims data. Public resources such as MIMIC-III [9] and large chest X-ray collections [10] have been critical for reproducible research, but they mostly originate from high-income settings, which affects how models generalize to Indian populations. Fig. 1 summarizes the major application domains discussed in this review.

© 2026 IJSRET

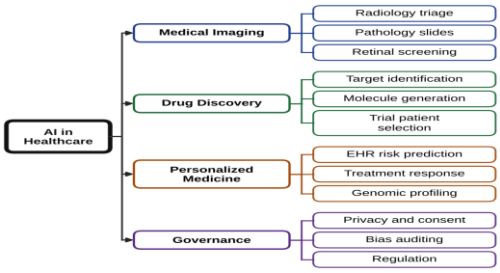


Fig. 1. Taxonomy of AI applications in healthcare covered in this review.

III. LITERATURE SURVEY

The survey of Litjens *et al.* [11] catalogued more than three hundred deep learning studies in medical image analysis and found that CNNs had largely replaced hand-crafted features by 2017. Gulshan *et al.* [2] trained an Inception-based network on over one hundred thousand fundus photographs and reported high sensitivity for referable diabetic retinopathy, a result that later informed screening programmes. De Fauw *et al.* [12] extended this idea to three-dimensional optical coherence tomography with a two-stage design that first segments tissue and then classifies referral urgency, improving robustness across scanner types.

In radiology, CheXNet [13] demonstrated pneumonia detection on chest radiographs at a level comparable to practising radiologists, while McKinney *et al.* [14] evaluated a breast-cancer screening system across the United Kingdom and the United States and observed fewer false positives and false negatives than single readers. These studies are compelling, but most are retrospective and rely on curated test sets.

For drug discovery, Stokes *et al.* [15] used a graph neural network to screen millions of molecules and identified halicin, a structurally novel antibiotic, and AlphaFold [16] solved protein structure prediction to near-experimental accuracy, opening new routes for target identification. In the EHR domain, Rajkomar *et al.* [17] showed that deep models built on raw, standardized records can predict in-hospital mortality and readmission without extensive manual feature engineering.

Equally important are studies that expose risk. Obermeyer *et al.* [18] revealed that a widely used commercial risk algorithm under-estimated the needs of Black patients because it used past cost as a proxy for illness. The WHO guidance on AI for health [19] and the ICMR ethical guidelines [20] respond to such findings with principles of autonomy, safety, transparency, accountability and equity. Table I summarizes the key works.

TABLE I Summary of Literature Survey

Ref.	Author(s) / Year	Technique / Approach	Key Contribution	Limitation / Gap
[11]	Litjens et al., 2017	Survey of deep learning in medical image analysis	Catalogued over 300 studies; showed CNN dominance across organs and modalities	Mostly retrospective studies; little clinical-outcome evidence
[2]	Gulshan et al., 2016	Inception CNN on retinal fundus photographs	High sensitivity for referable diabetic retinopathy	Limited to one disease; image-quality dependent
[12]	De Fauw et al., 2018	Two-stage segmentation + classification on 3-D OCT	Expert-level referral decisions; device-independent tissue maps	Requires costly OCT scanners
[13]	Rajpurkar et al., 2017	121-layer DenseNet (CheXNet) on chest X-rays	Radiologist-level pneumonia detection	Labels mined from reports are noisy
[14]	McKinney et al., 2020	Deep ensemble for mammography	Reduced false positives and negatives versus single readers	Retrospective; two countries only
[15]	Stokes et al., 2020	Message-passing graph neural network	Discovered antibiotic halicin from large virtual screen	Requires extensive wet-lab confirmation
[16]	Jumper et al., 2021	AlphaFold attention-based structure network	Near-experimental protein structure accuracy	Static structures; dynamics not modelled
[17]	Rajkomar et al., 2018	Deep models on FHIR-formatted EHR	Accurate mortality and readmission prediction	Tested in US hospitals only
[18]	Obermeyer et al., 2019	Audit of commercial risk algorithm	Exposed racial bias from cost-as-proxy label	Single algorithm studied

© 2026 IJSRET

3.1 Research Gap

Analysis of the surveyed literature reveals the following gaps:

- Most high-performing models are validated retrospectively on data from a few institutions; prospective, multi-centre trials in Indian hospitals remain rare.
- Training datasets under-represent rural populations, regional languages in clinical notes and low-cost imaging devices common in district hospitals.
- Explainability is often added after training rather than designed into the clinical workflow, limiting clinician trust.
- Drug-discovery gains are concentrated in early stages; evidence that AI improves late-phase trial success is still limited.
- 5) Governance frameworks such as the DPDP Act 2023 [21] address data protection, but dedicated rules for validating and monitoring clinical AI devices are still evolving.

3.2 Objectives

The objectives of this review are:

- To study and classify AI techniques used in medical imaging, drug discovery and personalized medicine.
- To compare representative methods in terms of data needs, accuracy, interpretability and deployment readiness.
- To identify ethical, privacy and bias-related risks and relate them to Indian policy initiatives.
- To propose a layered reference architecture for responsible deployment of clinical AI.
- To outline future research directions that link AI adoption to measurable patient outcomes.

3.3 Proposed System Architecture

Fig. 2 presents a five-layer reference architecture for responsible healthcare AI. It is designed so that every prediction can be traced back to its data source, model version and the clinician who acted on it.

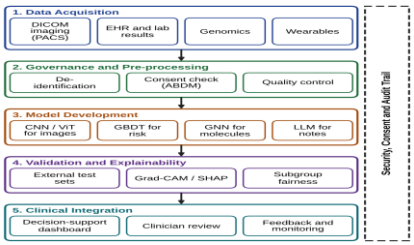


Fig. 2. Proposed five-layer reference architecture for responsible clinical AI deployment.

3.4 Data Acquisition Layer

This layer collects imaging studies in DICOM format, EHR entries, laboratory results, genomic files and wearable streams. Under the Ayushman Bharat Digital Mission, each record can be linked to a health ID with patient consent, which provides a lawful basis for secondary use.

3.5 Data Governance And Pre-Processing Layer

Records are de-identified, harmonized to common vocabularies and checked for quality. Consent status is verified before data enters any training pipeline, and differential access controls restrict who can view identifiable fields.

3.6 Model Development Layer

Task-specific models are trained here: CNNs and vision transformers for images, gradient-boosted trees for tabular risk prediction, graph networks for molecules and language models for notes. Each model is versioned and accompanied by a datasheet that documents its training population.

3.7 Validation And Explainability Layer

Models are evaluated on external test sets stratified by age, sex, region and device. Saliency maps such as Grad-CAM [22] and feature-attribution methods such as SHAP [23] provide case-level explanations, while subgroup metrics expose potential bias.

3.8 Clinical Integration Layer

Validated models are exposed through decision-support dashboards inside the hospital information system. Outputs are presented as suggestions with confidence levels, and the final decision, together with any override, is recorded for audit and continuous learning.

IV. METHODOLOGY AND WORKING

The working of the architecture follows the workflow in Fig. 3. After a clinical question is defined, data is gathered and curated, a model is trained, and performance is measured with metrics appropriate to the clinical risk. For a binary diagnostic task, sensitivity and specificity are computed as

$$Sensitivity = TP / (TP + FN) \quad (1)$$

$$Specificity = TN / (TN + FP) \quad (2)$$

where TP, TN, FP and FN denote true positives, true negatives, false positives and false negatives respectively. In screening, high sensitivity is preferred to avoid missed disease, while in confirmatory testing specificity is emphasized to limit unnecessary procedures. Overall discrimination is summarized by the area under the receiver operating characteristic curve (AUC).

© 2026 IJSRET

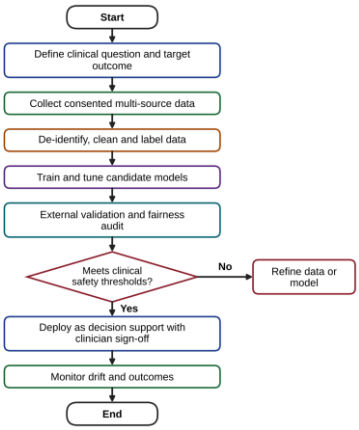


Fig. 3. Workflow for developing and deploying a clinical AI model.

Once deployed, performance is monitored for data drift, for example when a hospital replaces an X-ray machine. A drift indicator can be computed by comparing the distribution of incoming inputs with the training distribution; when it exceeds a threshold, the model is re-validated before further use.

V. COMPARATIVE ANALYSIS

Table II compares the main AI approaches by application area. Imaging CNNs are the most mature but require large expert-labelled datasets. Graph neural networks have accelerated virtual screening, whereas EHR models face the greatest challenges from missing and inconsistent records. Language models are promising for summarization but still require safeguards against fabricated content.

TABLE II Comparison of AI Techniques by Application Area

Area	Typical Model	Maturity	Main Limitation
Medical imaging	CNN, ViT, U-Net	High	Needs expert labels; scanner shift
Drug discovery	GNN, generative models	Medium	Late-phase benefit unproven
EHR risk	GBDT, transformers	Medium	Missing and biased records
Clinical text	LLMs, BioBERT	Emerging	Hallucination; privacy
Personalized therapy	Multi-omics ML	Early	Small cohorts; cost

VI. CHALLENGES AND FUTURE SCOPE

Table III maps the principal challenges to practical mitigation strategies. Beyond these, future research should focus on federated learning [24] so that hospitals can collaborate without moving patient data, on low-cost models that run on edge devices in primary health centres, and on regional-language interfaces that allow community health workers to use AI tools directly. National strategy documents [25] already identify healthcare as a priority sector; the next step is to establish shared Indian benchmark datasets and transparent post-market surveillance of approved AI tools.

TABLE III Challenges and Mitigation Strategies

Challenge	Mitigation
Patient-data privacy	Consent management, de-identification, federated learning
Dataset bias	Representative sampling, subgroup metrics, bias audits
Opacity of models	Saliency maps, SHAP, model documentation
Accountability	Human-in-the-loop sign-off and audit logs
Regulatory uncertainty	Alignment with ICMR guidelines and DPDP Act; post-market monitoring

VII. CONCLUSION

AI has moved from research curiosity to practical tool in several areas of healthcare, most clearly in medical imaging, and it is beginning to reshape drug discovery and personalized treatment. The reviewed evidence shows, however, that technical accuracy is only the first requirement. Reliable clinical evidence, representative data, explainable outputs, secure data handling and continuous human oversight are equally necessary. For India, the combination of a digital health infrastructure, a data protection law and national AI strategy creates a strong foundation, but dedicated validation and monitoring rules for clinical AI are still needed. The success of healthcare AI should ultimately be measured by better patient outcomes and more equitable access to care, not by the number of algorithms developed.

7.1 Acknowledgment

The authors express their sincere gratitude to the Department of Artificial Intelligence and Data Science, Jaihind College of Engineering, Pune, for providing the facilities, guidance

© 2026 IJSRET

and encouragement needed to carry out this work. The authors also thank the faculty members and colleagues whose valuable suggestions helped improve the quality of this review.

REFERENCES

1. Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, no. 7553, pp. 436–444, 2015.
2. V. Gulshan et al., "Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs," JAMA, vol. 316, no. 22, pp. 2402–2410, 2016.
3. A. Esteva et al., "Dermatologist-level classification of skin cancer with deep neural networks," Nature, vol. 542, no. 7639, pp. 115–118, 2017.
4. E. J. Topol, "High-performance medicine: The convergence of human and artificial intelligence," Nat. Med., vol. 25, no. 1, pp. 44–56, 2019.
5. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2016, pp. 770–778.
6. O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in Proc. MICCAI, 2015, pp. 234–241.
7. A. Vaswani et al., "Attention is all you need," in Proc. Adv. Neural Inf. Process. Syst. (NeurIPS), vol. 30, 2017.
8. K. Singhal et al., "Large language models encode clinical knowledge," Nature, vol. 620, no. 7972, pp. 172–180, 2023.
9. A. E. W. Johnson et al., "MIMIC-III, a freely accessible critical care database," Sci. Data, vol. 3, Art. no. 160035, 2016.
10. X. Wang et al., "ChestX-ray8: Hospital-scale chest X-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases," in Proc. IEEE CVPR, 2017, pp. 2097–2106.
11. G. Litjens et al., "A survey on deep learning in medical image analysis," Med. Image Anal., vol. 42, pp. 60–88, 2017.
12. J. De Fauw et al., "Clinically applicable deep learning for diagnosis and referral in retinal disease," Nat. Med., vol. 24, no. 9, pp. 1342–1350, 2018.
13. P. Rajpurkar et al., "CheXNet: Radiologist-level pneumonia detection on chest X-rays with deep learning," arXiv:1711.05225, 2017.
14. S. M. McKinney et al., "International evaluation of an AI system for breast cancer screening," Nature, vol. 577, no. 7788, pp. 89–94, 2020.
15. J. M. Stokes et al., "A deep learning approach to antibiotic discovery," Cell, vol. 180, no. 4, pp. 688–702, 2020.
16. J. Jumper et al., "Highly accurate protein structure prediction with AlphaFold," Nature, vol. 596, no. 7873, pp. 583–589, 2021.
17. A. Rajkomar et al., "Scalable and accurate deep learning with electronic health records," npj Digit. Med., vol. 1, Art. no. 18, 2018.
18. Z. Obermeyer, B. Powers, C. Vogeli, and S. Mullainathan, "Dissecting racial bias in an algorithm used to manage the health of populations," Science, vol. 366, no. 6464, pp. 447–453, 2019.
19. World Health Organization, Ethics and Governance of Artificial Intelligence for Health: WHO Guidance. Geneva, Switzerland: WHO, 2021.
20. Indian Council of Medical Research, Ethical Guidelines for Application of Artificial Intelligence in Biomedical Research and Healthcare. New Delhi, India: ICMR, 2023.
21. Government of India, The Digital Personal Data Protection Act, 2023 (Act No. 22 of 2023). New Delhi: Ministry of Law and Justice, 2023.
22. R. R. Selvaraju et al., "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in Proc. IEEE ICCV, 2017, pp. 618–626.
23. S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in Proc. NeurIPS, vol. 30, 2017.
24. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proc. AISTATS, 2017, pp. 1273–1282.
25. NITI Aayog, National Strategy for Artificial Intelligence #AIforAll. New Delhi, India: Government of India, 2018.

© 2026 IJSRET