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

AI in Healthcare: A Review of The Collect–Predict–Care Pipeline for Clinical Decision Support

Doke Shubham Somnath; Prof. N. G. Rokade; Prof. M. C. Kolse; Prof. S. K. Said
Department Of Artificial Intelligence And Data Science Jaihind College Of Engineering Pune, Maharashtra, India

Abstract—Healthcare systems face a shortage of trained doctors and specialists, heavy hospital workloads, rapidly growing volumes of medical data, diagnostic delays and unequal access to quality care, particularly in rural areas. Artificial Intelligence (AI), including machine learning, deep learning, computer vision and natural language processing, is increasingly used to learn patterns from medical data and support clinicians in diagnosis, treatment planning and patient monitoring. This paper reviews AI in healthcare through a three-stage collect–predict–care pipeline. In the collection stage, medical images such as X-ray, CT, MRI, retinal and skin photographs, laboratory results, electronic health records and wearable readings are gathered and prepared. In the prediction stage, convolutional neural networks analyse images, machine learning models estimate patient risk from records and language models interpret clinical notes. In the care stage, results are presented to doctors as decision support, and the doctor makes the final diagnosis and treatment decision. The review surveys landmark studies for each stage, proposes a layered clinical decision support architecture, and discusses challenges of privacy, bias, explainability, infrastructure cost and regulation. It concludes that AI can make healthcare faster, more accurate and more accessible when it is validated in real hospitals and always operates under clinical supervision.

Index Terms—Artificial intelligence, machine learning, deep learning, healthcare systems, clinical decision support, medical data analysis, electronic health records, telemedicine

I. INTRODUCTION

A district hospital in India may have a single radiologist reviewing hundreds of X-rays a day, while a primary health centre in a village may have no specialist at all.

Patients travel long distances, wait for reports and sometimes receive a diagnosis only after their condition has worsened. At the same time, every scan, test and prescription adds to a growing store of digital medical data that no human team can fully study.

AI offers a way to bridge this gap by learning from past cases to support present ones. Topol [1] describes this as the convergence of human and artificial intelligence, in which algorithms take over repetitive pattern recognition and give clinicians more time for patients. The goal is not to replace doctors but to extend their reach.

This paper explains AI in healthcare through a simple three-stage pipeline: collect, predict and care. It reviews representative studies for each stage, proposes an architecture for clinical decision support, identifies research gaps and discusses the challenges that must be resolved for safe adoption.

II. BACKGROUND AND FUNDAMENTAL CONCEPTS

2.1 The Collect–Predict–Care Pipeline

Fig. 1 shows the three stages. Collection gathers and cleans data from multiple sources. Prediction applies trained models to produce probabilities, risk scores or highlighted regions. Care places these outputs in front of a clinician, who interprets them in the context of the patient and takes the decision. Feedback from outcomes flows back to improve the models.

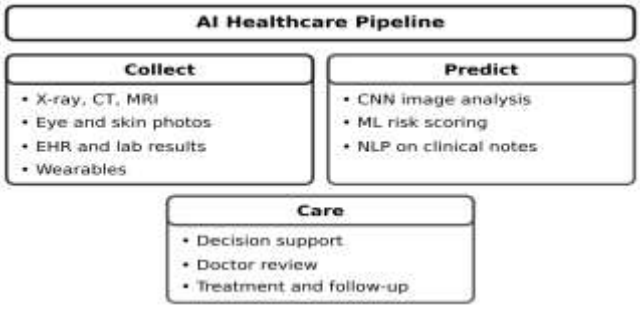


Fig. 1. The collect–predict–care pipeline of AI-based healthcare.

© 2026 IJSRET

2.2 Core Ai Techniques

Convolutional Neural Networks (CNNs) learn visual features from images and are the main tool for radiology, ophthalmology and dermatology. Architectures such as U-Net [2] segment organs or tumours. For tabular data in health records, gradient-boosted trees and deep sequence models estimate risk. Natural Language Processing (NLP) models pre-trained on biomedical text, such as BioBERT [3], extract diagnoses and medications from notes.

III. LITERATURE SURVEY

Collection and datasets. Large, labelled datasets enabled modern medical AI. Wang *et al.* [4] released ChestX-ray8 with over one hundred thousand radiographs labelled for common thoracic diseases by mining radiology reports, and Irvin *et al.* [5] introduced CheXpert, which explicitly models uncertain labels and provides a test set annotated by board-certified radiologists. MIMIC-III [6] provides de-identified intensive-care records that support research on prediction from health records.

Prediction from images. Esteva *et al.* [7] trained a CNN on about 130,000 skin-lesion images and achieved performance on par with dermatologists in distinguishing malignant from benign lesions. Gulshan *et al.* [8] reported high sensitivity and specificity for detecting referable diabetic retinopathy from retinal photographs, a condition that is common among diabetic patients in India and often goes unscreened.

Prediction from records and text. Rajkomar *et al.* [9] built deep learning models on entire electronic health records represented in a standard format and accurately predicted mortality, readmission and length of stay. Lee *et al.* [3] showed that BioBERT, pre-trained on biomedical literature, outperformed general language models in recognizing biomedical entities and relations. Singhal *et al.* [10] demonstrated that large language models can answer medical examination questions accurately, while noting gaps in safety and reasoning.

Care and fairness. Obermeyer *et al.* [11] found that a commercial algorithm used to allocate extra care systematically under-served Black patients because it predicted cost rather than illness, illustrating how a poorly chosen target can embed bias at the care stage. Table I summarizes these studies.

TABLE I Summary of Literature Survey

Ref.	Method	Key Findings	Limitations
[4]	ChestX-ray8 dataset with report-mined labels	Large benchmark for thoracic disease detection	Labels from reports are noisy
[5]	CheXpert with uncertainty labels	Expert-annotated test set; handles uncertainty	Single institution source
[6]	MIMIC-III critical-care database	Open de-identified records for research	ICU population only
[7]	Inception CNN on skin-lesion images	Dermatologist-level classification	Mostly lighter skin tones
[8]	Deep CNN on retinal photographs	Accurate diabetic retinopathy screening	Requires good image quality
[9]	Deep learning on full EHR	Predicted mortality and readmission	Complex to deploy
[3]	BioBERT biomedical language model	Better biomedical entity and relation extraction	English text only
[10]	Large language models for medical QA	High accuracy on medical exam questions	Safety and reasoning gaps
[11]	Bias audit of care-allocation algorithm	Revealed bias from cost-based target	One algorithm audited

IV. RESEARCH GAP

The following gaps are identified:

- **G1:** Many models are trained on data from a few large hospitals and are not validated in rural or resource-limited settings.
- **G2:** Integration of AI outputs into routine clinical workflows and electronic records is rarely studied.
- **G3:** Models seldom explain their predictions in a form that busy doctors can verify quickly.
- **G4:** Infrastructure costs for computing, storage and connectivity limit adoption in small hospitals.
- **G5:** Clear rules for approval, liability and post-deployment monitoring of AI tools are still developing.

© 2026 IJSRET

V. OBJECTIVES

The objectives of this review are:

1. 1) To explain the working of AI-based healthcare systems through the collect–predict–care pipeline.

2. 2) To review AI methods for medical images, health records and clinical notes.

3. 3) To propose a layered clinical decision support architecture with the doctor in control.

4. 4) To compare AI techniques in terms of data needs, performance and explainability.

5. 5) To analyse challenges of privacy, bias, cost and regulation and suggest mitigation.

VI. PROPOSED SYSTEM ARCHITECTURE

The architecture in Fig. 2 implements the collect–predict–care pipeline as a clinical decision support system (CDSS) that can serve both hospitals and telemedicine centres.

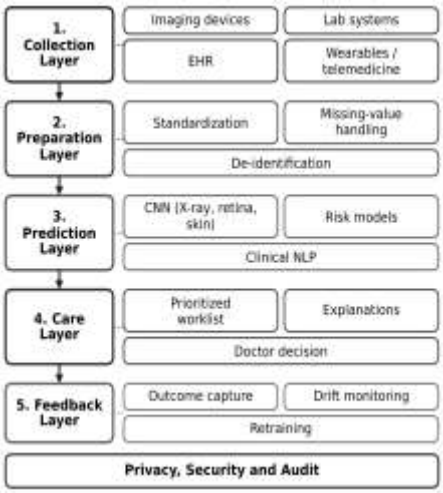


Fig. 2. Proposed clinical decision support architecture based on the collect–predict–care pipeline.

6.1 Collection Layer

Data enters from imaging devices, laboratory systems, electronic health records and wearable monitors. In telemedicine settings, health workers capture retinal or skin photographs with low-cost devices and upload them securely.

6.2 Preparation Layer

Images are standardized in size and intensity, records are cleaned and missing values handled, and identifiers are removed or encrypted. Data is mapped to standard codes so that models receive consistent inputs.

6.3 Prediction Layer

A model registry hosts validated models: CNNs for chest X-rays, retinal and skin images; risk models for conditions such as sepsis or readmission; and NLP models for summarizing notes. Each output includes a confidence score.

6.4 Care Layer

Results appear on the doctor's dashboard as prioritized worklists, highlighted image regions and risk explanations. The doctor accepts, modifies or rejects each suggestion and records the final decision.

6.5 Feedback Layer

Confirmed diagnoses and patient outcomes are collected to monitor model accuracy, detect performance drift and support periodic retraining.

VII. METHODOLOGY AND WORKING

The workflow is shown in Fig. 3. For an image classifier, the network outputs a probability for each condition through the softmax function

$$P(y_k \mid x) = exp(z_k) / \Sigma_j \, exp(z_j) \qquad (1)$$

where z_k is the score for class k. Cases with high probability of a serious condition are moved to the top of the doctor's worklist. Diagnostic quality is evaluated with sensitivity, specificity and the area under the ROC curve, and the positive predictive value is checked for the local disease prevalence:

$$PPV = TP / (TP + FP) \qquad (2)$$

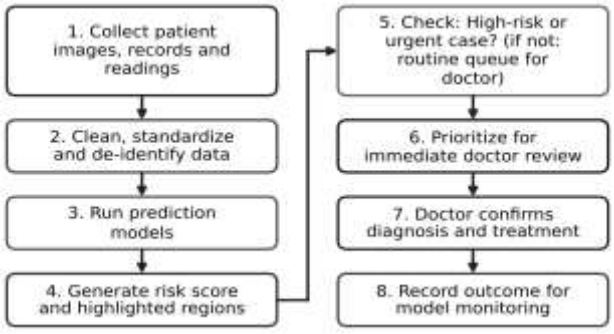


Fig. 3. Working of the AI-based clinical decision support system.

Saliency maps such as Grad-CAM [12] are shown with each image prediction, and SHAP values [13] explain risk scores, allowing the doctor to verify that the model attended to clinically meaningful features.

© 2026 IJSRET

VIII. COMPARATIVE ANALYSIS

Table II compares AI techniques used at the prediction stage. Image models are the most mature and have the strongest evidence; record-based risk models require careful handling of missing data; and language models are powerful but need safeguards against inaccurate statements.

TABLE II Comparison of AI Techniques at the Prediction Stage

Technique	Data Type	Evidence	Explainability
CNN classifier	X-ray, retina, skin	Strong	Saliency maps
U-Net segmentation	CT, MRI	Strong	Visual masks
Gradient boosting	EHR, lab values	Moderate	SHAP values
Sequence models	Time-series vitals	Moderate	Limited
Clinical NLP / LLM	Notes, reports	Emerging	Limited

IX. CHALLENGES AND FUTURE SCOPE

Table III maps challenges to mitigation measures. In the future, federated learning [14] can allow hospitals to train shared models without moving patient data, portable AI-enabled devices can bring screening to villages, and multilingual conversational interfaces can help patients understand their reports. WHO guidance [15] emphasizes that these advances must protect autonomy, ensure transparency and promote equity.

TABLE III Challenges and Mitigation

Challenge	Mitigation
Patient privacy	De-identification, encryption, consent, federated learning
Bias in data	Diverse datasets; subgroup evaluation
Hard-to-explain results	Saliency maps and SHAP explanations
Infrastructure cost	Cloud services; lightweight edge models
Regulation and liability	Clinical validation; doctor retains final decision

X. CONCLUSION

The collect–predict–care pipeline provides a clear way to understand how AI supports healthcare: data is gathered and prepared, models detect patterns and estimate risk, and doctors use these insights to make better decisions. The reviewed studies show that AI can match specialist performance on specific tasks and can process far more data than human teams. To realize this potential, AI tools must be tested thoroughly in real hospitals, protect patient privacy, be checked for bias and always work under the supervision of doctors. With responsible design, AI can make healthcare faster, more accurate and more accessible.

10.1 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 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. E. J. Topol, “High-performance medicine: The convergence of human and artificial intelligence,” Nat. Med., vol. 25, no. 1, pp. 44–56, 2019.
2. O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional networks for biomedical image segmentation,” in Proc. MICCAI, 2015, pp. 234–241.
3. J. Lee et al., “BioBERT: A pre-trained biomedical language representation model for biomedical text mining,” Bioinformatics, vol. 36, no. 4, pp. 1234–1240, 2020.
4. 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.
5. J. Irvin et al., “CheXpert: A large chest radiograph dataset with uncertainty labels and expert comparison,” in Proc. AAAI Conf. Artif. Intell., vol. 33, 2019, pp. 590–597.
6. A. E. W. Johnson et al., “MIMIC-III, a freely accessible critical care database,” Sci. Data, vol. 3, Art. no. 160035, 2016.
7. A. Esteva et al., “Dermatologist-level classification of skin cancer with deep neural networks,” Nature, vol. 542, no. 7639, pp. 115–118, 2017.
8. 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.

© 2026 IJSRET

9. A. Rajkomar et al., “Scalable and accurate deep learning with electronic health records,” npj Digit. Med., vol. 1, Art. no. 18, 2018.
10. K. Singhal et al., “Large language models encode clinical knowledge,” Nature, vol. 620, no. 7972, pp. 172–180, 2023.
11. 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.
12. R. R. Selvaraju et al., “Grad-CAM: Visual explanations from deep networks via gradient-based localization,” in Proc. IEEE ICCV, 2017, pp. 618–626.
13. S. M. Lundberg and S.-I. Lee, “A unified approach to interpreting model predictions,” in Proc. NeurIPS, vol. 30, 2017.
14. 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.
15. World Health Organization, Ethics and Governance of Artificial Intelligence for Health: WHO Guidance. Geneva, Switzerland: WHO, 2021.

© 2026 IJSRET