IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)
An International Open Access, Peer-reviewed, Refereed Journal

A Review On Machine Learning Applications In Healthcare

[1]Monika Jadhav, [2]Tejashri Dhomkar, [3]Sayali Kale
[1]*Assistant Professor*, [2]Student, [3]Student
[1]*Department of Computer Science*
[1]*Nowrosjee Wadia College*,Pune,India

Abstract: Machine learning (ML) has emerged as a transformative force in modern healthcare, enabling computers to learn from complex medical data and assist clinicians in making faster, more accurate decisions. This paper provides a structured review of how ML techniques — including supervised learning, unsupervised learning, reinforcement learning, and deep learning — are being applied across key healthcare domains such as disease diagnosis, drug discovery, medical imaging, patient monitoring, and genomics. We examine representative studies from the literature, compare ML algorithms with respect to accuracy and clinical utility, and discuss the practical challenges that currently hinder large-scale adoption, including data privacy, model interpretability, and regulatory compliance. We further identify open research gaps and outline promising directions for future work, including federated learning, explainable AI, and multimodal data fusion. Our aim is to offer researchers and practitioners a clear, accessible overview of the current landscape and the road ahead.

Keywords - machine learning, healthcare, deep learning, medical imaging, disease prediction, drug discovery, electronic health records, federated learning, explainable AI.

I. Introduction

The healthcare industry generates an enormous volume of data every day — from electronic health records (EHRs) and medical imaging scans to genomic sequences and wearable sensor readings. Extracting clinically meaningful insights from this data manually is both time-consuming and error-prone. Machine learning (ML) offers a powerful alternative: algorithms that can identify hidden patterns, build predictive models, and continuously improve with experience [1].

Over the past decade, ML has moved from academic research into real clinical workflows. Tools powered by deep neural networks now assist radiologists in detecting tumours, help pathologists classify tissue samples, and flag high-risk patients before their condition deteriorates [2]. At the same time, the field faces genuine challenges: healthcare data is often scarce, imbalanced, and governed by strict privacy regulations, while clinicians rightly demand explanations for algorithmic recommendations before acting on them [3].

This paper organises the current state of ML in healthcare into a coherent narrative. Section II surveys related work. Section III describes the major ML paradigms. Section IV catalogues key application domains. Section V presents a comparative analysis. Sections VI and VII address challenges and research gaps, respectively. Section VIII outlines future directions, and Section IX concludes the paper.

The paper is organized as follows:

II. Section II discusses background and related work
III. Section III describes the major ML paradigms in healthcare
IV. Section IV catalogues key application domains
V. Section V presents a comparative analysis
VI. Section VI discusses challenges
VII. Section VII describes research gaps
VIII. Section VIII outlines future directions
IX. Section X concludes the paper

II. LITERATURE REVIEW

Early work by Shortliffe et al. [4] demonstrated that rule-based expert systems could assist physicians in antibiotic selection, laying the groundwork for AI-assisted diagnosis. The introduction of support vector machines (SVMs) in the 1990s and random forests in the 2000s broadened the toolkit available to medical data scientists [5].

The deep-learning revolution, triggered by Krizhevsky et al.'s ImageNet results [6], had an immediate impact on medical imaging. Esteva et al. [7] demonstrated that a convolutional neural network (CNN) could classify skin lesions with accuracy matching board-certified dermatologists. Similarly, Rajpurkar et al. [8] showed that a CNN trained on chest X-rays could detect pneumonia more reliably than radiologists in several test scenarios.

Beyond imaging, recurrent neural networks (RNNs) and transformer models have been applied to clinical text mining and longitudinal EHR analysis. Miotto et al. [9] proposed Deep Patient, an unsupervised representation of patient histories that predicted future diagnoses across 78 disease categories. More recently, large language models have been explored for clinical note summarisation and patient triage [10].

Systematic reviews by Topol [11] and Obermeyer and Emanuel [12] confirm that ML is not intended to replace clinicians but to augment them — freeing physicians from repetitive analytical tasks so they can focus on patient interaction and complex clinical reasoning.

III. METHODOLOGY: ML PARADIGMS IN HEALTHCARE

A. Supervised Learning

Supervised learning trains a model on labelled examples (input–output pairs). In healthcare, labels correspond to diagnoses, treatment outcomes, or risk scores. Common algorithms include logistic regression, decision trees, SVMs, and gradient-boosted trees. These methods work well when labelled data are plentiful and the prediction target is clearly defined [1].

B. Unsupervised Learning

Unsupervised methods discover structure in unlabelled data. Clustering algorithms such as k-means and hierarchical clustering are used to identify patient sub-groups with similar clinical profiles. Dimensionality-reduction techniques like principal component analysis (PCA) and autoencoders help visualise high-dimensional genomic or imaging data [5].

C. Reinforcement Learning

Reinforcement learning (RL) trains an agent to optimise a reward signal through trial and error. In healthcare, RL has been applied to personalised treatment planning, dynamic drug-dosing, and robotic surgery assistance, where the reward can be defined as improved patient outcomes [13].

D. Deep Learning

Deep learning uses multi-layer neural networks to learn hierarchical feature representations directly from raw data. CNNs excel at image analysis; RNNs and long short-term memory (LSTM) networks handle sequential data such as time-series vital signs; and transformer architectures process free-text clinical notes. Deep learning typically requires large labelled datasets and significant computational resources, but delivers state-of-the-art accuracy in many medical tasks [6][7].

IV. APPLICATIONS

A. Disease Diagnosis and Prediction

ML models trained on EHR data can predict the onset of conditions such as diabetes, heart disease, and sepsis hours or even days before clinical signs become obvious. Random forests and gradient boosting (e.g., XGBoost) have achieved AUC scores above 0.90 on several benchmark datasets [14]. Early warning systems built on such models are now deployed in ICUs at several major hospitals.

B. Medical Imaging

CNN-based systems assist radiologists by detecting and localising anomalies in X-rays, CT scans, MRIs, and histopathology slides. Google's LYNA system detects metastatic breast cancer in lymph-node biopsies with 99% accuracy [15]. In ophthalmology, deep-learning systems detect diabetic retinopathy from fundus photographs with sensitivity and specificity comparable to trained ophthalmologists.

C. Drug Discovery and Development

Developing a new drug traditionally takes 10–15 years and costs over a billion US dollars. ML accelerates this pipeline by predicting molecule bioactivity, identifying drug–target interactions, and repurposing existing approved drugs for new indications. Graph neural networks and recurrent architectures have shown particular promise in molecular property prediction [16].

D. Genomics and Personalised Medicine

Genome-wide association studies (GWAS) generate millions of genetic variants per patient. ML methods condense this information into polygenic risk scores and identify variants linked to rare diseases. Combined with clinical data, these models support truly personalised treatment regimens — selecting therapies most likely to benefit a specific patient's genetic profile [17].

E. Remote Patient Monitoring

Wearable devices continuously capture heart rate, oxygen saturation, blood glucose, and activity levels. ML algorithms process these streams in real time to detect arrhythmias, hypoglycaemic events, and deteriorating respiratory function. This capability is especially valuable for managing chronic diseases outside the hospital setting and reducing unnecessary readmissions [18].

F. Natural Language Processing in Clinical Text

A large proportion of medical knowledge is locked inside free-text notes, discharge summaries, and radiology reports. NLP techniques — from named-entity recognition to transformer-based language models — extract structured information, identify medication errors, and assist in clinical coding and billing [10].

V. COMPARATIVE ANALYSIS

Table I summarises commonly used ML algorithms across key healthcare tasks, comparing their typical accuracy range, data requirements, interpretability, and computational cost. The values are drawn from aggregated results in the literature and are intended as indicative benchmarks rather than absolute figures.TABLE I. COMPARISON OF ML ALGORITHMS FOR HEALTHCARE APPLICATIONS

Study	**Technique**	**Application Area**	**Advantages**	**Accuracy**	**Limitations**
Chen et al. [2]	Logistic Regression	Disease Prediction	Simple, interpretable	~75% – 85%	Limited to linear relationships
Hussain et al. [4]	Random Forest	Clinical Decision Support	High accuracy	~75% – 85%	Less interpretable
Chen et al. [2]	SVM	Classification tasks	Effective in high dimensions	~75% – 85%	Computationally expensive
Esteva et al. [1]	CNN	Medical Imaging	Automatic feature extraction	~75% – 85%	Requires large datasets
Miotto et al. [5]	RNN / LSTM	Time-series health data	Good for sequential data	~80% – 88%	Training complexity

VI. CHALLENGES

A. Data Privacy and Security

Medical records are among the most sensitive personal data. Regulations such as HIPAA (USA) and GDPR (Europe) impose strict requirements on how patient data may be collected, stored, and shared for research. These constraints limit the volume of data available for training and complicate multi-institutional collaboration [3].

B. Data Quality and Imbalance

Real-world clinical datasets are often incomplete, inconsistently labelled, and heavily imbalanced — rare diseases, by definition, have few positive examples. Training robust models under these conditions requires careful preprocessing, resampling strategies, and evaluation metrics beyond simple accuracy (e.g., F1-score, AUC-ROC) [5].

C. Interpretability and Clinical Trust

Deep learning models are powerful but opaque. Clinicians are understandably reluctant to act on recommendations they cannot understand or verify. Explainable AI (XAI) methods such as SHAP, LIME, and Grad-CAM aim to provide post-hoc explanations, but these are often approximate and may not fully satisfy clinical requirements [19].

D. Regulatory and Ethical Barriers

Deploying an ML model as a clinical decision-support tool requires regulatory approval (e.g., FDA 510(k) clearance in the USA). The approval pathway for continuously learning systems remains unclear. Ethical concerns around algorithmic bias — particularly the risk that models perform worse for under-represented patient groups — also require systematic attention [20].

E. Generalisation Across Settings

A model trained on data from one hospital may perform poorly at another due to differences in patient demographics, clinical workflows, and equipment. This distribution shift problem is a significant barrier to the broad deployment of ML tools in healthcare [11].

VII. RESEARCH GAPS

Despite substantial progress, several important gaps remain in the literature. First, most published studies evaluate ML models on retrospective datasets; prospective clinical trials that measure real-world patient outcome improvements are far less common [12]. Second, benchmark datasets for many clinical tasks remain small, single-site, and demographically homogeneous, making it difficult to assess true generalisability [5].

Third, the majority of ML papers in healthcare focus on prediction accuracy while largely ignoring calibration — the alignment between a model's predicted probabilities and actual observed outcomes. A model that is accurate but poorly calibrated can be misleading in clinical practice [14]. Fourth, few studies examine the long-term performance of deployed models as clinical practice and patient populations evolve over time — a phenomenon known as model drift. Robust monitoring and retraining pipelines are needed but rarely described [18].

Finally, multi-modal learning — integrating imaging, genomics, clinical text, and sensor data — remains an open research challenge. Most existing work considers a single data modality, missing the richer insights that come from combining complementary information sources [17].

VIII. FUTURE SCOPE

A. Federated Learning

Federated learning allows multiple hospitals to collaboratively train a shared model without transferring patient data. Each site trains locally and shares only model gradients, preserving privacy while enabling learning from diverse populations. This paradigm has strong potential to overcome data-sharing barriers and improve model robustness [21].

B. Explainable AI (XAI)

Advances in XAI will be critical for clinical adoption. Future research should move beyond post-hoc explanation methods toward inherently interpretable architectures that maintain high accuracy. Standardised evaluation frameworks for explanation quality — as perceived by clinicians — are also needed [19].

C. Multimodal Data Integration

Combining imaging data, genomic profiles, clinical text, and wearable sensor streams into unified ML models represents a frontier with great clinical promise. Transformer-based architectures and cross-modal attention mechanisms show early potential for tackling this integration challenge [17].

D. AI-Assisted Drug Discovery

Generative models such as variational autoencoders and diffusion models are beginning to design novel drug candidates with desired properties from scratch. Combined with RL-based optimisation, these approaches could dramatically shorten the drug development timeline [16].

E. Real-Time Clinical Decision Support

The convergence of edge computing and lightweight ML models will enable real-time decision support at the bedside. Tiny neural networks capable of running on wearable devices could detect acute events — such as atrial fibrillation or hypoglycaemia — and alert both patients and clinicians instantaneously [18].

REFERENCES

[1] A. Esteva, B. Kuprel, R. A. Novoa et al., “Dermatologist-level classification of skin cancer with deep neural networks,” Nature, vol. 542, no. 7639, pp. 115– 118, 2017.
[2] M. Chen, Y. Hao, K. Hwang, L. Wang, and L. Wang, “Disease prediction by machine learning over big healthcare data,” IEEE Access, vol. 5, pp. 8869– 8879, 2017.
[3] J. Choi, E. Schuetz, and J. Stewart, “Machine learning for healthcare: Review, opportunities, and challenges,” Briefings in Bioinformatics, vol. 21, no. 5, pp. 1667–1678, 2020.
[4] S. A. Hussain, M. A. Khan, and M. B. Alam, “A review on machine learning algorithms for healthcare data analysis,” IEEE Access, vol. 9, pp. 110246– 110269, 2021.
[5] R. Miotto, F. Wang, S. Wang, X. Jiang, and J. T. Dudley, “Deep learning for healthcare: Review, opportunities and challenges,” Briefings in Bioinformatics, vol. 19, no. 6, pp. 1236–1246, 2018.
[6] M. A. Albahli and A. A. Alhindi, “Applications of machine learning in healthcare and medical diagnosis: A review,” Health Informatics Journal, vol. 29, no. 1, pp. 1–18, 2023.
[7] P. Rajpurkar, J. Irvin, K. Zhu et al., “CheXNet: Radiologist-level pneumonia detection on chest X-rays with deep learning,” arXiv preprint arXiv:1711.05225, 2017.
[8] N. S. Islam and T. Hoque, “AI and machine learning in drug discovery and development: Current trends and future prospects,” Computers in Biology and Medicine, vol. 160, 2023.
[9] A. R. Mamoshina, L. Ojomoko, and A. Zhavoronkov, “Applications of deep learning in biomedicine,” Molecular Pharmaceutics, vol. 13, no. 5, pp. 1445– 1454, 2018. 21
[10] S. H. Lee, Y. J. Kim, and H. S. Park, “Deep learning in medical imaging: General overview,” Korean Journal of Radiology, vol. 21, no. 4, pp. 505–517, 2020.