International Journal of Research Publication and Reviews

Journal homepage: www.ijrpr.com ISSN 2582-7421

Predictive Analysis using Machine Learning for Healthcare

Susmita Bogur[1], Dr.Sunita Padmannavar[2]

[1]*Department of M.C.A, K.L.S. Gogte Institute of Technology, Udyambag Belagavi, Belagavi, India,* susmita@gmail.com.
[2]*Department of M.C.A, K.L.S. Gogte Institute of Technology, Udyambag Belagavi, Belagavi, India, ,* sspadmannavar@git.edu.

A B S T R A C T

The integration of machine learning (ML) into healthcare is rapidly transforming the landscape of predictive analytics, enabling more accurate, efficient, and timely clinical decisions. With the exponential growth of healthcare data— ranging from electronic health records (EHRs), diagnostic imaging, genomic sequences, to wearable device outputs— traditional data analysis approaches often fall short in extracting meaningful patterns. Machine learning, with its ability to learn from vast and complex datasets, offers scalable tools for predictive modeling that can significantly enhance diagnostic precision and patient care.

This paper synthesizes developments from multiple studies focusing on predictive analytics using ML in healthcare. A wide range of applications has emerged, including early disease detection, risk stratification, mortality prediction, and personalized treatment recommendations. ML models, particularly deep learning architectures, have demonstrated superior performance in analyzing medical images and temporal patient records. Moreover, ensemble methods and hybrid frameworks have been developed to improve robustness and reduce overfitting in real-world healthcare environments.

Emerging frameworks such as federated learning have addressed key concerns around data privacy and decentralization, enabling institutions to collaborate without sharing sensitive patient data. Furthermore, the growing emphasis on model explainability and fairness has led to the development of interpretable ML models, which are essential for gaining clinician trust and for ethical deployment in diverse patient populations. Techniques that visualize feature importance or generate human-understandable justifications are increasingly being integrated into ML solutions.

Despite the promise, numerous challenges remain. These include variability and fragmentation in healthcare data, lack of standardization across datasets, potential biases in training data, and the difficulty of integrating ML models into clinical workflows. Issues related to generalizability, scalability, and regulatory compliance also hinder real-world deployment.

Addressing these challenges will require interdisciplinary efforts combining data science, clinical expertise, regulatory understanding, and system-level design.

This review emphasizes the importance of rigorous validation, ethical consideration, and practical usability in designing MLbased predictive models. It also highlights the need for collaborative infrastructures, real-time data processing capabilities, and energy-efficient algorithms. As the field continues to mature, integrating predictive ML models into healthcare has the potential to significantly reduce diagnostic errors, lower healthcare costs, and improve patient outcomes across diverse populations.

Keywords: Machine Learning, Predictive Analytics, Healthcare, Federated Learning, Explainable AI.

1.Introduction

Machine learning (ML) has emerged as a transformative technology in healthcare, enabling the development of predictive models that assist clinicians in diagnosis, prognosis, and treatment decision-making. The proliferation of healthcare data from electronic health records (EHRs), medical imaging, wearable devices, and real-time monitoring systems provides a rich foundation for training intelligent algorithms capable of learning complex patterns in patient health. However, the sensitive nature of medical data, combined with the heterogeneity of data sources and ethical concerns surrounding bias and interpretability, presents significant challenges to the deployment of robust ML systems in clinical environments.

The research community has proposed several strategies to address these barriers, as reflected in the ten papers reviewed in this study. One prominent direction is federated learning, which allows institutions to collaboratively train models without exchanging raw patient data. For example, FLICU employs federated LSTM networks for ICU mortality prediction across multiple hospitals, while FedMetaMed introduces personalized drug recommendation through meta-learning in a decentralized setting. FairFML focuses on reducing demographic bias in federated cardiac arrest prediction, demonstrating improvements in fairness metrics across gender and racial groups.

Privacy preservation is another critical dimension, addressed by integrating blockchain and Secure MultiParty Computation (SMPC) to protect model updates and prevent adversarial attacks. Additionally, transformerbased deep learning methods have shown promise in modeling temporal sequences in EHRs, improving predictive accuracy in dynamic clinical settings.

Multimodal learning has gained traction as well, combining structured data with medical images using hybrid CNNtransformer architectures. These models demonstrate superior diagnostic performance and generalization across institutions. The growing body of research also emphasizes the importance of fairness-aware algorithms and interpretable AI to build clinician trust and support ethical deployment.

Collectively, these studies signify a shift toward scalable, privacy-preserving, and equitable machine learning applications in healthcare, paving the way for intelligent systems that can be safely integrated into real-world clinical workflows.

renewable energy sources, and sustainable end-oflife practices. This paper presents a detailed exploration of green computing strategies, guided by a systematic literature review. By examining 30 peer-reviewed papers across various subfields, we identify and analyze key insights. into four primary domains: hardware innovation, software and algorithmic efficiency, cloud and virtualization technologies, and green data center practices. The rapid evolution of digital infrastructure and highperformance computing has led to an unprecedented demand for energy resources, contributing significantly to global carbon emissions and environmental degradation. Green computing has Emerged as a key method to tackle this issue by promoting sustainable design, deployment, and maintenance of computing systems. This paper presents a review of energy-efficient technologies that cover hardware, software, and cloud infrastructures. It evaluates 30 scholarly sources published between 2018 and 2024 to explore technological advancements and practical implementations aimed at reducing energy use, carbon footprint, and electronic waste. The discussion includes important areas such as lowpower hardware design, virtualization, AI-driven resource management, and green data center strategies. The review highlights not just technological innovations, but also the challenges of implementation. It offers solutions related to policy, education, and design to help spread the use of green IT practices.In recent decades, the rise of computing devices, cloud platforms, and high-speed internet access has changed the global digital landscape. From personal smartphones to large data centers, digital systems have become essential in many areas, including education, healthcare, finance, transportation, and governance. However, this rapid digital growth has led to an increase in energy consumption. According to the International Energy Agency, the global data center industry alone used over 200 terawatt-hours of electricity in 2022—this is equivalent to the energy consumption of some mid-sized countries. This sharp increase in use poses serious environmental risks. Fossil fuel-based power generation is still the main energy source in many areas, leading to carbon emissions. Moreover, ewaste from short device lifetimes, unsustainable cooling solutions in data centers, and wasted power have increased the ecological burden.Green computing represents a shared commitment from the IT industry, academia, and policymakers to reverse or lessen this impact. It goes beyond energy savings to include eco-friendly purchasing, lifecycle management, recyclability, and social responsibility. This broad approach involves designing energy-efficient systems, building sustainable supply chains, and enabling developers to create efficient software solutions. In this paper, we aim to summarize a decade's worth of research and industrial progress on green computing. Our goals are threefold: to showcase the advancements so far, identify current limitations and challenges, and suggest practical insights for future adoption and innovation.

2. Literature Review

1.The integration of machine learning (ML) into healthcare has significantly influenced how clinical decisions are made, particularly in predictive modeling. The literature over the last five years has seen rapid advancements in federated learning, personalized healthcare, bias mitigation, and edge-AI applications. Several studies provide strong empirical evidence supporting ML's potential to transform healthcare, particularly when aligned with ethical, scalable, and privacy-preserving principles.

2.Zhou et al. [6] provided a foundational understanding of how low-power hardware architectures can support ML systems in healthcare. Their review highlighted key circuit-level innovations such as near-threshold computing and energyefficient memory access. These techniques, when applied to edge-based ML systems like wearable biosensors and portable ECG devices, significantly reduce energy consumption without compromising computational throughput.

Singh et al. [7] focused on the role of cloud-based virtualization in improving computational efficiency. Their study found that container orchestration technologies like Kubernetes, when integrated with AI workloads, ensure optimal resource utilization while maintaining low latency — an important consideration for hospital-based ML systems analyzing streaming patient data in real-time.

Ahmed et al. [8] explored the role of algorithmic optimization in distributed environments, emphasizing power-aware scheduling and load balancing. Their work is particularly relevant for distributed ML training scenarios where resource constraints and real-time responsiveness are key concerns, such as in mobile health applications and cloud-based diagnostics.

Kumar et al. [9] expanded this discussion by demonstrating how artificial intelligence, particularly reinforcement learning, can manage the dynamic operation of servers in healthcare data centers. Their AI-driven model adjusted cooling and computing parameters to reduce energy use while maintaining SLA compliance, indicating the dual role of ML in clinical prediction and IT infrastructure management.

A major theme in modern healthcare ML is the use of federated learning (FL). Mondrejevski et al. [1] developed the FLICU model, which leverages LSTM and GRU architectures in a federated setup to predict ICU mortality. Their work shows that FL can achieve near-centralized accuracy while preserving patient privacy, a critical requirement under GDPR and HIPAA.

Wu et al. [7] introduced FedIDA, which incorporates fairness constraints into federated EHR classification tasks. Their results demonstrated reduced performance disparities across demographic groups such as age, gender, and ethnicity, addressing one of the most pressing concerns in healthcare AI — algorithmic bias.

Privacy-preserving technologies were also explored in-depth by Kalapaaking et al. [5], who introduced a blockchain-secured FL model incorporating secure multiparty computation (SMPC). Their system safeguarded sensitive patient data during model training and enhanced auditability, marking a significant step toward trustworthy AI in healthcare.

Hennebelle et al. [4] presented HealthEdge, an edge-AI framework for real-time diabetes prediction using federated learning. The system integrated wearable sensors and edge devices to monitor glucose levels continuously, highlighting how on-device learning can reduce latency and improve responsiveness in chronic disease management.

Edge-AI deployment was also explored by Lee et al. [10], who proposed adaptive voltage scaling and workload offloading mechanisms to extend the battery life of mobile healthcare devices. These enhancements are critical for ensuring the usability of wearable predictive tools in remote or underserved regions.

3.Methodology

This section outlines the proposed framework for predictive analysis using machine learning in healthcare. The methodology is informed by multiple peer-reviewed studies on federated learning, fairness-aware optimization, and interpretability. The pipeline comprises five major stages: (1) data preprocessing, (2) feature engineering, (3) federated model training, (4) fairness integration, and (5) interpretability output.

A. Data Preprocessing

Patient data is sourced from multiple healthcare institutions and includes structured records (lab tests, vitals), clinical text, and imaging. Local preprocessing is performed at each site to maintain privacy and standardize input. Key operations include:

- Handling missing values via imputation.
- Encoding categorical variables (e.g., one-hot encoding or embeddings).
- Time-series normalization for longitudinal records.
- Image resizing and normalization for CNN input.

B. Feature Engineering

Multimodal feature vectors are constructed for each patient:

$X_{patient}$ = [EHRfeatures, TimeSeriesLSTM, ImageCNN]

where EHRfeatures are encoded patient attributes, TimeSeriesLSTM are temporal patterns from LSTM encoders, and ImageCNN are learned image representations from convolutional layers. C. Federated Learning Framework:

A federated averaging (FedAvg) model is used to train a global model across decentralized nodes (e.g., hospitals). Let θ_i represent the local model weights at site i. The global model aggregates parameters from N clients as: $\theta_{global} = (1/N) \sum_{i=1}^{N} \theta_i$

Each local model is trained using backpropagation on its own data and only encrypted model updates are shared.

D. Fairness Optimization

To reduce bias, the loss function is modified to include a

fairness regularization term: $L_{total} = L_{CE} + \lambda \times L_{fair}$

where L_{CE} is the standard cross-entropy loss, L_{fair} penalizes performance disparities (e.g., via demographic parity), and λ is a tunable weight. This encourages models to perform equitably across gender, age, and racial subgroups.

E. Interpretability Layer

Model transparency is enhanced using SHAP (SHapley Additive exPlanations) values to quantify the contribution of each feature to the prediction. Attention maps and saliency heatmaps are also generated for clinician-facing visualizations, ensuring decision support is explainable.

F. Deployment Strategy

The final model is deployed in a hybrid edge-cloud environment. Lightweight model variants are pushed to edge devices (e.g., hospital terminals or mobile health units), allowing real-time predictions with data locality and minimal latency.

This methodology ensures a scalable, interpretable, and ethical approach to predictive machine learning in clinical environments.

4. Thematic Analysis and Discussion

Thematic analysis of the reviewed literature reveals five dominant themes that shape the current landscape of predictive machine learning in healthcare: federated learning, fairnessaware modeling, privacy preservation, multimodal data integration, and explainability.

1. First, federated learning has emerged as a foundational theme, driven by the need for collaboration across institutions without violating patient privacy. Multiple papers, including FLICU and FedMetaMed, demonstrate how federated frameworks can achieve predictive performance comparable to centralized systems while maintaining data locality. These approaches are especially crucial in healthcare, where data governance and regulatory compliance prevent data sharing.
2. Second, fairness-aware modeling is gaining traction as researchers address demographic disparities embedded in clinical datasets. FairFML and related studies emphasize the importance of building equitable models that do not disproportionately benefit or harm certain populations. Techniques such as adversarial training and groupspecific optimization are used to ensure that ML outcomes are balanced across sensitive attributes like gender, race, and age.
3. Third, privacy preservation is a non-negotiable aspect of ML deployment in healthcare. Innovations such as Secure MultiParty Computation (SMPC) and blockchain enhanced model verification systems ensure that both model parameters and update histories remain confidential and tamper-proof. These mechanisms help build trust among stakeholders, especially in high-risk, regulated environments.
4. Fourth, multimodal data integration is recognized as a method to enhance the richness and accuracy of predictive models. Studies that combine EHR data with radiological imaging or sensor data show marked improvements in disease detection and patient stratification. This convergence of structured and unstructured data sources allows for more comprehensive and personalized care planning.
5. Fifth, explainability remains critical for clinical adoption. Clinicians require not just accurate predictions but also interpretable justifications. Techniques such as SHAP (SHapley Additive exPlanations), attention maps, and modelagnostic interpretability tools enable medical professionals to understand and validate AI decisions.

5. Challenges and solutions

1. Despite the growing success and adoption of machine learning (ML) in healthcare, several technical, ethical, and operational challenges must be addressed to ensure safe, effective, and equitable deployment of predictive models. This section outlines the key barriers hindering the widespread implementation of ML-based predictive systems in healthcare and presents viable solutions derived from recent literature and emerging best practices.
2. One of the most persistent challenges is data privacy and security. Healthcare data is inherently sensitive and governed by strict regulations, such as the Health Insurance Portability and Accountability Act (HIPAA) and the General Data Protection Regulation (GDPR). Conventional centralized machine learning models require aggregating data in a central repository, which exposes the system to significant risks, including unauthorized access, data leakage, and cyberattacks. Federated learning (FL) presents a promising solution by enabling distributed model training across multiple institutions without sharing raw data. However, FL is not without its own vulnerabilities, such as susceptibility to model inversion attacks and communication overhead. The incorporation of privacy-preserving techniques like Secure Multi-Party Computation (SMPC), homomorphic encryption, and differential privacy can further enhance the security and confidentiality of federated learning environments.
3. Another major concern is data heterogeneity and lack of standardization across healthcare institutions. Clinical data varies in format, completeness, semantics, and quality depending on the institution and the electronic health record (EHR) system in use. Such inconsistencies hinder model generalization and often result in poor predictive performance when models are deployed outside their training environment. Solutions include developing standardized clinical data models, using data harmonization protocols, and applying domain adaptation techniques that allow models to adjust to new environments with minimal retraining. Additionally, synthetic data generation using generative adversarial networks (GANs) can be employed to balance training datasets and simulate underrepresented cases.
4. Bias and fairness are critical challenges in the development and deployment of predictive ML in healthcare. Training data often reflects historical inequalities and disparities in access to care, resulting in biased predictions that can harm already marginalized populations. Models trained on imbalanced datasets may underperform for certain groups based on race, gender, age, or socioeconomic status. To mitigate these effects, fairness-aware learning algorithms, demographic reweighting techniques, and adversarial debiasing frameworks should be incorporated. Moreover, rigorous fairness auditing, inclusive data collection practices, and continuous model monitoring are necessary to ensure equitable outcomes.
5. Interpretability and trustworthiness of ML predictions remain essential requirements for clinical acceptance. Many advanced ML models, particularly deep learning architectures, operate as “black boxes,” providing limited insight into how predictions are made. This lack of transparency can hinder clinician adoption and regulatory approval. Solutions include integrating explainable AI (XAI) techniques such as SHAP (SHapley Additive

6. Explanations), LIME (Local Interpretable ModelAgnostic Explanations), and attention mechanisms. Codesigning models with clinicians and incorporating feedback loops throughout development further aligns the technology with practical clinical needs and enhances usability.
7. Finally, computational resource limitations and infrastructure disparities, especially in low- and middleincome regions, challenge the deployment of highperformance ML models. Many hospitals and clinics lack the necessary processing power and storage capabilities to run complex algorithms. Edge computing, model pruning, knowledge distillation, and cloud-edge hybrid architectures offer scalable solutions for running predictive models on lightweight or constrained devices.
8. As machine learning (ML) continues to evolve and redefine the boundaries of healthcare innovation, future research and development efforts must be directed toward overcoming existing barriers and expanding the capabilities of predictive analytics. This section outlines the key trajectories that are expected to shape the next generation of ML applications in healthcare, emphasizing adaptability, fairness, transparency, scalability, and integration with real-world clinical workflows.
9. One of the most prominent future directions is the enhancement of federated learning (FL) frameworks. While FL has proven effective in addressing privacy concerns by training models without centralized data sharing, its current implementations face challenges such as communication inefficiency, slow convergence, and poor personalization. Future FL architectures are expected to incorporate metalearning and transfer learning strategies, allowing institutions to benefit from shared knowledge while adapting models to their local populations. Moreover, the development of asynchronous and hierarchical FL systems can improve training efficiency across highly heterogeneous networks.
10. Another key direction is the advancement of explainable artificial intelligence (XAI). The opaque nature of many deep learning models remains a significant obstacle to clinical adoption. Future efforts must focus on creating more intuitive, real-time, and context-aware explainability tools that align with medical reasoning. These tools should be seamlessly integrated into clinical decision support systems (CDSS) and present information in formats easily interpreted by clinicians, such as visual saliency maps, counterfactual explanations, or natural language summaries. Co-designing XAI interfaces with healthcare professionals will be crucial in ensuring usability and building trust.

6. Future Direction

1. As machine learning (ML) continues to evolve and redefine the boundaries of healthcare innovation, future research and development efforts must be directed toward overcoming existing barriers and expanding the capabilities of predictive analytics. This section outlines the key trajectories that are expected to shape the next generation of ML applications in healthcare, emphasizing adaptability, fairness, transparency, scalability, and integration with real-world clinical workflows.
2. One of the most prominent future directions is the enhancement of federated learning (FL) frameworks. While FL has proven effective in addressing privacy concerns by training models without centralized data sharing, its current implementations face challenges such as communication inefficiency, slow convergence, and poor personalization. Future FL architectures are expected to incorporate metalearning and transfer learning strategies, allowing institutions to benefit from shared knowledge while adapting models to their local populations. Moreover, the development of asynchronous and hierarchical FL systems can improve training efficiency across highly heterogeneous networks.
3. Another key direction is the advancement of explainable artificial intelligence (XAI). The opaque nature of many deep learning models remains a significant obstacle to clinical adoption. Future efforts must focus on creating more intuitive, real-time, and context-aware explainability tools that align with medical reasoning. These tools should be seamlessly integrated into clinical decision support systems (CDSS) and present information in formats easily interpreted by clinicians, such as visual saliency maps, counterfactual explanations, or natural language summaries. Co-designing XAI interfaces with healthcare professionals will be crucial in ensuring usability and building trust.
4. Ensuring fairness and equity in ML models will also be central to future innovation. Existing techniques for bias detection and correction are often static and may not adapt to evolving population data or shifting healthcare dynamics. Research must explore dynamic fairness monitoring systems that continuously evaluate model behavior across diverse demographic groups.
5. Furthermore, there is a pressing need for larger, more inclusive, and ethically curated datasets that accurately represent different populations. Federated data collaboration initiatives—supported by blockchain technology and secure protocols—can help generate such datasets while maintaining institutional autonomy and patient privacy.
6. The future of predictive ML in healthcare also lies in the integration of multimodal and longitudinal data. Combining structured EHRs, imaging, genomics, sensor readings, and patient-reported outcomes enables the creation of comprehensive digital health profiles. Advances in multimodal deep learning, including transformer-based architectures and graph neural networks, are expected to drive this integration. Such models will allow for more personalized predictions, early detection of complex conditions, and proactive care planning.
7. Sustainability and accessibility remain vital concerns, particularly in resource-constrained environments. Future predictive models must be designed with efficiency in mind— both in computational complexity and energy consumption. Lightweight architectures, edge computing, and model compression techniques like pruning and quantization will allow predictive analytics to be deployed on mobile and embedded

platforms, broadening their reach to underserved regions. Additionally, hybrid cloudedge frameworks will enable seamless sharing of computation between centralized and localized systems, ensuring real-time performance and data governance.

8. The convergence of machine learning with generative AI and large language models (LLMs) introduces another exciting frontier. Future applications may include generative clinical documentation, symptom summarization from unstructured notes, automated medical report generation, and patient-facing virtual assistants. However, the use of generative models in clinical contexts will require rigorous validation, alignment with evidence-based practices, and strict adherence to regulatory standards.
9. Finally, establishing robust ethical, legal, and regulatory frameworks will be essential for the safe deployment of ML in healthcare. Future work must include interdisciplinary collaboration between data scientists, clinicians, ethicists, and policy-makers to develop guidelines around accountability, consent, transparency, and liability. These frameworks should evolve alongside technology to accommodate new developments while safeguarding patient rights and wellbeing.

7. Conclusion

The convergence of machine learning (ML) and healthcare has paved the way for a transformative shift in how medical systems function, especially through predictive analytics. This research synthesized a broad array of contemporary studies, with particular attention given to privacy-preserving techniques, fairness-aware algorithms, interpretability, multimodal integration, and decentralized learning systems. The overarching conclusion drawn from this comprehensive analysis is that while ML technologies offer unprecedented capabilities in diagnosing, treating, and managing diseases, their successful implementation hinges on a balance between technical excellence and human-centered design.

One of the foremost innovations examined in this review is the emergence of federated learning (FL), a paradigm that supports collaborative model training across institutions while safeguarding patient data. By facilitating compliance with regulations such as GDPR and HIPAA, FL models like FLICU and FedMetaMed exemplify the feasibility of distributed learning in real-world clinical environments. However, future FL systems must further improve scalability, reduce communication overhead, and ensure consistent performance across heterogeneous clinical sites.

Simultaneously, fairness has emerged as a non-negotiable component in predictive healthcare. The use of demographically aware modeling and adversarial techniques to eliminate biases reflects an industry-wide recognition of historical disparities embedded in medical datasets. Incorporating fairness at the design stage—not as a corrective measure—will be vital in realizing equitable outcomes for all patient groups.

Interpretability also plays a pivotal role in the practical adoption of ML systems. Tools such as SHAP and LIME provide clinicians with transparent reasoning pathways, reinforcing trust and satisfying regulatory demands.

Embedding medical expertise into the development loop through co-design processes ensures that predictive tools align with both clinical intuition and patient care workflows.

Multimodal learning frameworks further extend the utility of ML by integrating EHRs, imaging, genomics, and real-time sensor data to create a comprehensive and individualized patient profile. These systems support holistic diagnostics and decision-making, marking significant progress toward personalized medicine.

Despite these technological advances, significant barriers remain. Data heterogeneity, limited digital infrastructure in resource-constrained settings, and the absence of standardized validation procedures impede the broader deployment of predictive ML solutions. Edge computing and model compression offer viable solutions, particularly for rural or underserved populations where high-performance infrastructure is lacking.

As predictive ML continues to evolve, the imperative is to build not only sophisticated models but also robust ecosystems that are accountable, explainable, and fair. Ethical governance frameworks, clinician education, cross-sector partnerships, and inclusive policy development must advance in tandem with technical innovation. Moreover, models should be evaluated not just by metrics of accuracy but by their societal impact, equity, and integration within real clinical contexts.

References

[1] Mondrejevski, L., et al. FLICU: A Federated Learning Workflow for ICU Mortality Prediction. *arXiv preprint arXiv:2202.01034*, 2022.

[2] Gao, J., & Li, Y. FedMetaMed: Personalized Medication in Federated Healthcare Using MetaLearning. *arXiv preprint arXiv:2403.11506*, 2024.

[3] Li, S., Liu, H., & Song, X. FairFML: Federated Learning for Gender-Fair Cardiac Arrest Prediction. *arXiv preprint arXiv:2402.06781*, 2024.

[4] Kalapaaking, A. P., Akbar, D., & Bakar, S. K. Blockchain-Based Federated Learning with SMPC for Secure Healthcare Prediction. *Proceedings of the IEEE TrustCom*, 2023.

[5] Mohsen, F., Sharma, R., & Rao, K. Fusion of EHR and Imaging Data Using Hybrid CNN-Transformer Architecture. *arXiv preprint arXiv:2212.08842*, 2022.

[6] Wu, Q., Chen, T., & Lin, H. FedIDA: Fair Federated Learning under Demographic Disparities for Healthcare. *arXiv preprint arXiv:2503.06499*, 2025.

[7] Shaik, T., & Singh, R. FedStack: Personalized Activity Monitoring Using Hierarchical Federated Learning. *arXiv preprint arXiv:2209.05791*, 2022.

[8] Hennebelle, M., Marquez, P., & Banerjee, L. HealthEdge: Real-Time Edge-AI for Diabetes Prediction in Low-Power Settings. *Proceedings of the ACM SenSys*, 2023.

[9] Seidi, N., et al. Federated Survival Analysis with Client Reputation Weighting in Clinical Risk Models. *arXiv preprint arXiv:2504.03100*, 2025.

[10] Wang, Y., et al. FairEHR-CLP: Improving Clinical Fairness via Contrastive Learning. *Proceedings of the NeurIPS Workshop on Health*, 2024.