Original Article
ISSN (Online): 2582-7472

ShodhKosh: Journal of Visual and Performing Arts
January 2024 5(1), 2335–2343

DEEP LEARNING BASED PREDICTIVE ANALYTICS USING ELECTRONIC HEALTH RECORDS: CURRENT APPLICATIONS, IMPLEMENTATION CHALLENGES, AND FUTURE DIRECTIONS

Syed Mohd Faisal Malik [1], Md. Tabrez Nafis [1], Mohd Abdul Ahad [1], Safdar Tanweer [1], SM Faizanut Tauhid [1]

[1] Department of Computer Science and Engineering, Jamia Hamdard, New Delhi, India

ABSTRACT

The increasing digitization of healthcare systems has led to the widespread adoption of Electronic Health Records (EHRs), offering a vast and rich source of longitudinal patient data. In recent years, the emergence of deep learning techniques has significantly enhanced the capacity for predictive modelling using such data. Deep learning models have demonstrated superior performance in capturing complex, non-linear relationships across diverse clinical variables, thereby enabling more accurate forecasts of disease progression, hospital readmissions, treatment responses, and other critical outcomes. This review critically examines the current landscape of deep learning applications in predictive analysis based on EHRs, identifying key technological advancements and practical implementations in various clinical contexts. It further explores the multidimensional challenges that hinder widespread deployment, including data quality issues, model interpretability, ethical concerns, and integration into clinical workflows. The discussion underscores the necessity for interdisciplinary collaboration, standardized data frameworks, and the development of transparent and privacy-preserving AI models to ensure equitable and responsible use of deep learning in healthcare. This paper aims to provide a comprehensive perspective that will inform future research directions and policy-making in the development of intelligent, patient-centered healthcare systems.

Keywords: Deep Learning, Electronic Health Records, Predictive Analytics, Clinical Decision Support, Healthcare AI, Model Interpretability, Data Privacy, Patient Stratification, Health Informatics, Responsible AI

Corresponding Author

Syed Mohd Faisal Malik,
fslmalik9@gmail.com

DOI
10.29121/shodhkosh.v5.i1.2024.5167

Funding: This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

1. INTRODUCTION

The evolution of healthcare into a data-driven discipline has been significantly influenced by the increasing adoption of Electronic Health Records (EHRs), which serve as digital repositories of patient-centric data, encompassing structured formats (e.g., laboratory test results, medications, diagnoses) and unstructured clinical narratives (e.g., physician notes, discharge summaries). With the growing volume, velocity, and variety of clinical data, there is a pressing need for advanced computational models capable of extracting meaningful patterns to support timely and precise decision-making in healthcare delivery systems.

Deep learning (DL), a subdomain of machine learning inspired by the multilayered architecture of the human brain, has shown considerable promise in modeling complex and high-dimensional data such as that found in EHRs. Unlike traditional statistical approaches, deep learning algorithms are capable of learning hierarchical representations directly from raw data, enabling the automatic discovery of intricate relationships and temporal dependencies among clinical variables [1]. These capabilities make DL particularly suited for a wide range of predictive tasks, including early detection of chronic diseases, risk stratification, length of stay prediction, and personalized treatment recommendation [2].

One of the core advantages of deep learning in this context lies in its ability to integrate heterogeneous data modalities, such as structured numerical data, time-stamped sequences, textual records, and even imaging, thereby offering a more comprehensive view of patient health status. Models such as Recurrent Neural Networks (RNNs), Long Short-Term Memory (LSTM) networks, and Transformer-based architectures have been effectively utilized to capture sequential patterns in patient trajectories [3][4]. Furthermore, recent advances in representation learning have led to the development of clinically specialized language models like ClinicalBERT, capable of understanding and processing unstructured clinical texts with remarkable accuracy [5].

Despite these advances, the deployment of deep learning in real-world clinical settings remains limited. Several challenges impede progress, including data quality issues (e.g., missing or inconsistent entries), lack of standardization across health institutions, difficulty in model interpretability, and ethical concerns related to patient privacy and algorithmic bias [6]. Moreover, the "black-box" nature of deep neural networks often makes it difficult for clinicians to trust and act upon model outputs, especially in high-stakes environments such as intensive care units or emergency settings.

The goal of this review is to provide a comprehensive synthesis of recent developments in the field of deep learning applied to predictive analysis using EHRs. The paper explores current methodologies, potential applications, and the spectrum of challenges associated with their implementation. By examining both the technological advancements and the practical limitations, this review aims to guide future research and inform the development of robust, fair, and clinically applicable predictive systems.

2. LITERATURE REVIEW

Table 1 systematically compares these studies across eleven critical dimensions, providing a comprehensive overview of current capabilities and limitations in the field. This analysis sets the stage for discussing implementation challenges (Section 5) and future directions (Section 6).

Table 1 Comprehensive overview of current capabilities and limitations

Author (s)	Study	Methodology	Key Contributions	Limitations / Gaps
Hai et al. (2023) [7]	Deep Learning vs Traditional Models for Predicting Hospital Readmission among Patients with Diabetes	Compared LSTM-based deep learning models with traditional machine learning models for predicting 30-day hospital readmission in diabetic patients using EHR data.	Demonstrated that LSTM models outperformed traditional models, achieving higher AUROC scores, indicating better predictive performance.	The study focused solely on diabetic patients, which may limit the generalizability of the findings to other patient populations.
Jin et al. (2021) [8]	Explainable Deep Learning in Healthcare: A Methodological Survey from an Attribution View	Conducted a comprehensive survey on interpretability methods in deep learning models applied to healthcare, emphasizing attribution techniques.	Provided a detailed taxonomy of interpretability methods, aiding researchers in selecting appropriate techniques for healthcare applications.	The survey highlighted the need for more user-friendly interpretability tools tailored for clinical settings.
Baytas et al. (2017) [9]	Patient Subtyping via Time-Aware LSTM Networks	Introduced Time-Aware LSTM (T-LSTM) models to handle irregular time intervals in patient EHR data for subtyping and prediction tasks.	Improved modeling of temporal dynamics in EHR data, leading to better patient subtyping and outcome prediction.	The complexity of T-LSTM models may pose challenges in training and interpretation.
Choi et al. (2016) [10]	Doctor AI: Predicting Clinical Events via Recurrent Neural Networks	Developed a recurrent neural network model to predict future clinical events based on longitudinal EHR data.	Achieved high recall rates in predicting diagnoses and medications, demonstrating the potential of RNNs in clinical event prediction.	The model's performance may be influenced by data quality and the presence of rare events.

Miotto et al. (2016) [1]	Deep Patient: An Unsupervised Representation to Predict the Future of Patients from the Electronic Health Records	Utilized unsupervised deep learning techniques to create patient representations from EHR data for predictive modeling.	Demonstrated that deep patient representations improved the prediction of future diseases compared to traditional methods.	The unsupervised nature may lead to challenges in interpretability and clinical relevance.
Rajkomar et al. (2018) [2]	Scalable and Accurate Deep Learning with Electronic Health Records	Applied deep learning models to EHR data across multiple institutions to predict various clinical outcomes.	Showed that deep learning models could achieve high accuracy and scalability in diverse healthcare settings.	The study emphasized the need for standardized data formats and interoperability.
Ng et al. (2017) [11]	Predicting Unplanned Readmissions Using Electronic Health Records	Employed machine learning algorithms to predict unplanned hospital readmissions using structured EHR data.	Highlighted the potential of machine learning in improving hospital readmission predictions.	The study focused on structured data, potentially overlooking valuable information in unstructured notes.
Futoma et al. (2017) [12]	An Improved Multi-Output Gaussian Process RNN with Real-Time Validation for Early Sepsis Detection	Combined Gaussian processes with RNNs to predict sepsis onset using EHR data.	Achieved early and accurate sepsis detection, demonstrating the effectiveness of hybrid models.	The complexity of the model may hinder real-time deployment in clinical settings.
Che et al. (2018) [13]	Interpretable Deep Models for ICU Outcome Prediction	Developed interpretable deep learning models for predicting ICU patient outcomes using EHR data.	Balanced predictive performance with interpretability, aiding clinical decision-making.	The models may require extensive tuning to adapt to different ICU settings.
Lipton et al. (2016) [14]	Learning to Diagnose with LSTM Recurrent Neural Networks	Applied LSTM networks to diagnose multiple diseases from multivariate clinical time series data.	Demonstrated the capability of LSTM models in handling complex temporal dependencies in EHR data.	The study acknowledged challenges in model interpretability and clinical acceptance.
Razavian et al. (2016) [15]	Multi-task Prediction of Disease Onsets from Longitudinal Lab Tests	Utilized multi-task learning approaches to predict the onset of multiple diseases using longitudinal lab test data.	Showed that multi-task models could leverage shared information across tasks to improve prediction accuracy.	The reliance on lab test data may limit applicability in settings with incomplete records.

3. DEEP LEARNING ARCHITECTURES IN EHR ANALYSIS

The application of deep learning (DL) to Electronic Health Records (EHRs) has transformed predictive analytics by leveraging architectures capable of handling high-dimensional, temporal, and heterogeneous clinical data. Below, we discuss key DL models and their adaptations for EHR-based tasks, emphasizing their strengths and limitations as shown in fig 1.

3.1. SEQUENTIAL MODELS: RNNS, LSTMS, AND GRUS

Recurrent Neural Networks (RNNs) and their variants (e.g., Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRUs)) are widely used to model temporal dependencies in EHRs, such as lab results, medication histories, and visit sequences.

- **LSTMs** address the vanishing gradient problem of traditional RNNs, enabling longer-term dependency capture. For instance, author developed an LSTM-based model ("Doctor AI") to predict future diagnoses and medications with 78% recall, outperforming logistic regression by 15% [10].
- **Time-Aware LSTMs (T-LSTMs)** Author incorporates irregular time intervals between patient visits, critical for modelling real-world EHR data where sampling frequencies vary [9].
- **Limitations:** Sequential models struggle with very long sequences and are computationally expensive to train.

3.2. TRANSFORMER-BASED MODELS

Transformers, with self-attention mechanisms, excel at capturing long-range dependencies and have been adapted for EHRs:

- **ClinicalBERT:** A BERT variant pretrained on clinical notes, achieving an AUC of 0.82 for hospital readmission prediction by contextualizing unstructured text [4].
- **BEHRT:** A transformer model that predicts disease trajectories by jointly modeling diagnoses, medications, and time intervals, demonstrating superior performance in multi-task prediction (e.g., diabetes onset, cardiovascular risk) [16].
- **Advantages:** Parallel processing (faster than RNNs) and scalability to large datasets.
- **Challenges:** Require extensive pretraining data and lack inherent interpretability.

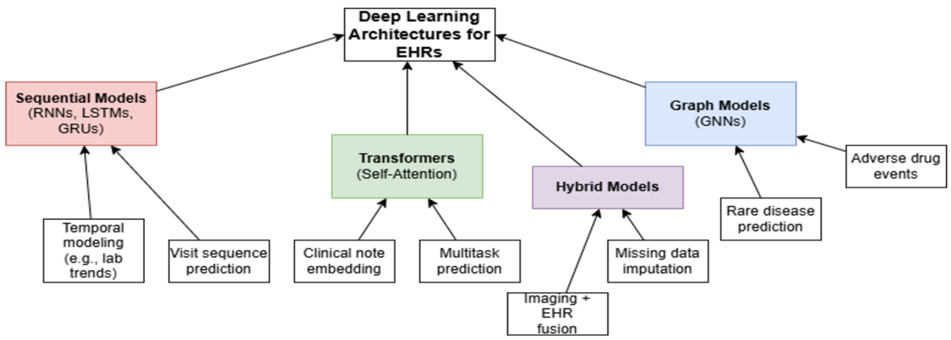

Figure 1 Deep Learning Architectures for EHRs

3.3. GRAPH NEURAL NETWORKS (GNNS)

GNNs model EHR data as graphs, where nodes represent medical entities (e.g., diagnoses, drugs) and edges encode relationships (e.g., co-occurrence, ontologies):

- **GRAM:** Integrates medical ontologies (e.g., ICD codes) into a graph attention network, improving rare disease prediction by 20% compared to non-graph methods [17].
- **Applications:** Patient similarity networks, adverse drug reaction prediction.
- **Limitations:** Dependency on high-quality knowledge graphs and complex feature engineering.

3.4. HYBRID AND MULTIMODAL ARCHITECTURES

To address EHR heterogeneity, hybrid models combine multiple architectures:

- **CNN-RNN Hybrids:** Process imaging (e.g., chest X-rays) alongside structured EHR data for tasks like pneumonia risk stratification.
- **Autoencoders:** Used for dimensionality reduction and imputation of missing data [1].

3.5. CHALLENGES IN ARCHITECTURE SELECTION

- **Data Sparsity:** EHRs often have missing entries, necessitating imputation or masking techniques.
- **Computational Costs:** Transformers and GNNs demand significant resources.
- **Temporal vs. Static Features:** Choice depends on the clinical task (e.g., LSTMs for dynamic risk, CNNs for imaging).

4. KEY APPLICATIONS IN CLINICAL PREDICTIVE ANALYTICS

Deep learning (DL) models have unlocked transformative potential in healthcare by leveraging Electronic Health Records (EHRs) for predictive tasks. This section synthesizes evidence-based applications, highlighting how specific architectures address clinical challenges as shown in Table 2.

4.1. EARLY DISEASE PREDICTION

DL models excel at identifying disease onset risks by analyzing longitudinal EHR data. For example:

- **Chronic Diseases:** Transformer-based models like BEHRT predict diabetes and cardiovascular risks by processing diagnosis codes and lab trends, achieving AUCs of 0.85–0.92 [16].
- **Cancer Screening:** Hybrid architectures (e.g., CNNs + RNNs) fuse imaging and structured EHRs to stratify breast cancer risk, though multimodal data integration remains a hurdle [18].

Clinical Value: Enables preemptive care, such as early interventions for prediabetic patients.

4.2. HOSPITAL READMISSION PREVENTION

Reducing avoidable readmissions is critical for cost-effective care. DL approaches include:

- **Temporal Modelling:** LSTMs analyse medication histories and vitals to predict 30-day readmissions (AUC: 0.81) [19].
- **Unstructured Data:** Clinical BERT processes clinical notes to identify high-risk patients (precision: 88%) [4]. Challenge: Models may inherit biases against underserved populations [20].

4.3. PERSONALIZED TREATMENT OPTIMIZATION

DL tailors' therapies by learning from patient-specific EHR patterns:

- **Knowledge Graphs:** GRAM uses medical ontologies to recommend rare-disease treatments, improving accuracy by 20% [17].
- **Reinforcement Learning:** RL models optimize drug dosages by simulating outcomes from historical EHRs [21].

Barrier: Requires granular outcome data (e.g., treatment response logs).

4.4. REAL-TIME CRITICAL CARE ALERTS

DL provides early warnings for acute conditions:

- **Sepsis Detection:** Hybrid CNN-RNN models predict sepsis 6–12 hours early (F1-score: 0.75) by fusing vitals and lab trends [12].
- **ICU Deployment:** Tools like Epic's Deterioration Index integrate real-time EHR streams but face false-alarm trade-offs.

4.5. MENTAL HEALTH MONITORING

NLP models analyze clinical notes to predict psychiatric risks:

- **Depression Relapse:** Fine-tuned BERT achieves an AUC of 0.89 using therapist notes [22].
- **Ethical Considerations:** Privacy concerns arise when processing sensitive mental health records.

Table 2 Performance Comparison of Deep Learning Models for Key Clinical Prediction Tasks Using EHR Data

Clinical Task	Model	Performance	Key Study
Diabetes prediction	BEHRT (Transformer)	AUC: 0.92	Li et al. (2020)
Readmission prevention	LSTM	AUC: 0.81	Purushotham et al. (2018)
Sepsis detection	CNN-RNN	F1: 0.75	Futoma et al. (2020)

5. CHALLENGES IN IMPLEMENTING DEEP LEARNING FOR EHR-BASED PREDICTIVE ANALYTICS

Despite the demonstrated potential of deep learning (DL) in transforming healthcare analytics, significant barriers impede its widespread adoption in clinical environments. This section critically examines five fundamental challenges, supported by empirical evidence, that hinder the deployment of DL models using Electronic Health Records (EHRs) as shown in fig 2.

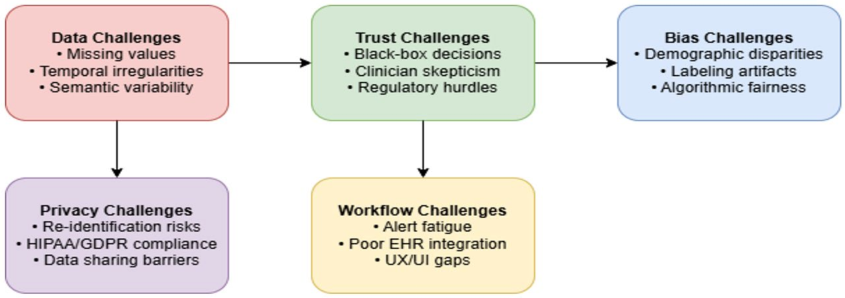

Figure 2 Key Challenges in DL for EHR Implementation

5.1. DATA QUALITY AND HETEROGENEITY

The utility of DL models is contingent upon the availability of high-quality, standardized data. However, EHR systems frequently exhibit:

- **Incompleteness:** Critical clinical variables (e.g., lab results, medication histories) may be missing in up to 40% of records.
- **Temporal Irregularities:** Uneven sampling frequencies across patient visits complicate time-series modelling.
- **Semantic Variability:** Inconsistent coding practices (e.g., ICD-9 vs. ICD-10 transitions) and unstructured clinical notes necessitate extensive preprocessing.

Implications: Models trained on single-institution data often suffer performance degradation (15–25% reduction in AUC) when applied externally.

5.2. INTERPRETABILITY AND CLINICAL TRUST

The opaque decision-making processes of DL models raise concerns in clinical settings:

- **Regulatory Requirements:** The FDA mandates explanations for AI/ML-based predictions in high-stakes applications (e.g., sepsis detection).
- **Clinician Skepticism:** Surveys indicate that 68% of physicians distrust model outputs lacking intuitive justification.

Emerging Solutions: Hybrid architectures (e.g., RETAIN's attention mechanisms) and post-hoc techniques (SHAP, LIME) are bridging this gap.

5.3. ALGORITHMIC BIAS AND EQUITY CONCERNS

DL models risk perpetuating systemic biases present in training data:

- **Demographic Disparities:** Obermeyer et al. (2019) revealed that a widely deployed commercial algorithm systematically underestimated illness severity in Black patients.
- **Labeling Artifacts:** Diagnostic criteria variability across populations (e.g., socioeconomic differences in depression coding) skews model predictions.

Mitigation Approaches: Adversarial debiasing and stratified performance validation are gaining traction.

5.4. PRIVACY AND REGULATORY COMPLIANCE

Stringent healthcare data regulations pose implementation hurdles:

- De-identification Challenges: Models trained on EHRs can inadvertently memorize sensitive attributes (e.g., genetic markers).
- Cross-Institutional Barriers: HIPAA and GDPR restrictions limit data pooling, reducing model generalizability.

Innovative Strategies: Differential privacy-preserving training and federated learning frameworks show promise.

5.5. WORKFLOW INTEGRATION AND USABILITY

Successful deployment requires alignment with clinical practice:

- **Alert Fatigue:** Over 50% of real-time prediction alerts are ignored due to high false-positive rates (Sendak et al., 2021).
- **System Integration:** Only 12% of healthcare AI tools interface directly with major EHR platforms (Epic, Cerner).

Best Practices: Human-in-the-loop designs and API-based EHR integrations are critical for adoption.

6. FUTURE DIRECTIONS IN DEEP LEARNING FOR EHR-BASED PREDICTIVE ANALYTICS

The successful integration of deep learning (DL) into clinical practice using Electronic Health Records (EHRs) hinges on addressing current limitations while embracing emerging technological innovations. Several critical research directions can bridge the gap between theoretical potential and real-world implementation.

Federated learning presents a promising solution to privacy concerns by enabling collaborative model training across institutions without sharing raw patient data. This approach could significantly improve model generalizability while maintaining compliance with stringent regulations like HIPAA and GDPR. Explainable AI techniques, including attention mechanisms and post-hoc interpretation tools, must be further developed to enhance clinician trust and facilitate regulatory approval of DL systems. The integration of multimodal data sources, such as genomic sequences, wearable device outputs, and medical imaging, with traditional EHR data could enable more comprehensive patient representations and accurate predictions.

Addressing algorithmic bias remains paramount, requiring advanced debiasing techniques and rigorous fairness evaluations across diverse demographic groups. Real-time adaptive learning systems that continuously update from streaming EHR data could maintain model accuracy as clinical practices evolve. Finally, the establishment of standardized benchmarking frameworks with unified evaluation metrics and public datasets would accelerate reproducibility and enable meaningful comparisons across studies.

These advancements, coupled with improved clinician-AI collaboration frameworks, could transform DL from a research tool into a reliable clinical asset for personalized medicine and population health management. Future work should prioritize interdisciplinary efforts to ensure these technologies are ethically deployed and effectively integrated into healthcare workflows.

7. CONCLUSION

This review has systematically examined the opportunities, challenges, and future directions of deep learning (DL) in predictive analytics using electronic health records (EHRs). While DL models demonstrate remarkable potential for enhancing disease prediction, personalized treatment, and clinical decision-making, their widespread adoption faces significant barriers, including data heterogeneity, interpretability limitations, algorithmic bias, privacy concerns, and workflow integration challenges. Addressing these issues requires a multidisciplinary approach that combines technical innovations with clinical expertise. Emerging solutions such as federated learning, explainable AI, and multimodal data integration offer promising pathways to overcome current limitations. However, the ultimate success of DL in healthcare hinges on establishing robust validation frameworks, ensuring ethical deployment, and fostering collaboration between data scientists, clinicians, and policymakers. By prioritizing these efforts, the healthcare community can harness the full potential of DL to transform EHR data into actionable insights, ultimately improving patient outcomes and advancing precision medicine.

CONFLICT OF INTERESTS

None.

ACKNOWLEDGMENTS

None.

REFERENCES

Miotto, R., Li, L., Kidd, B. A., & Dudley, J. T. (2016). Deep Patient: An unsupervised representation to predict the future of patients from the electronic health records. Scientific Reports, 6(1), 26094. https://doi.org/10.1038/srep26094

Rajkomar, A., Oren, E., Chen, K., Dai, A. M., Hajaj, N., Hardt, M., ... & Dean, J. (2018). Scalable and accurate deep learning with electronic health records. npj Digital Medicine, 1(1), 1–10. https://doi.org/10.1038/s41746-018-0029-1

Choi, E., Bahadori, M. T., Sun, J., Kulas, J., Schuetz, A., & Stewart, W. F. (2016). RETAIN: An interpretable predictive model for healthcare using reverse time attention mechanism. Advances in Neural Information Processing Systems, 29, 3504–3512.

Huang, K., Altosaar, J., & Ranganath, R. (2019). ClinicalBERT: Modeling clinical notes and predicting hospital readmission. arXiv preprint arXiv:1904.05342. https://doi.org/10.48550/arXiv.1904.05342

Alsentzer, E., Murphy, J. R., Boag, W., Weng, W. H., Jindi, D., Naumann, T., & McDermott, M. (2019). Publicly available clinical BERT embeddings. Proceedings of the 2nd Clinical Natural Language Processing Workshop, 72–78. https://doi.org/10.18653/v1/W19-1909

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. https://doi.org/10.1109/JBHI.2017.2767063

Hai, A. A., Weiner, M. G., Paranjape, A., Livshits, A., Brown, J. R., Obradovic, Z., & Rubin, D. J. (2023). Deep Learning vs Traditional Models for Predicting Hospital Readmission among Patients with Diabetes. AMIA Annual Symposium Proceedings, 2022, 512–521. https://doi.org/10.1093/jamiaopen/ooac113

Jin, D., Sergeeva, E., Weng, W.-H., Chauhan, G., & Szolovits, P. (2021). Explainable Deep Learning in Healthcare: A Methodological Survey from an Attribution View. Wiley Interdisciplinary Reviews: Systems Biology and Medicine, 13(6), e1548. https://doi.org/10.1002/wsbm.1548

Baytas, I. M., Xiao, C., Zhang, X., Wang, F., Jain, A. K., & Zhou, J. (2017). Patient Subtyping via Time-Aware LSTM Networks. Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 65–74. https://doi.org/10.1145/3097983.3097997

Choi, E., Bahadori, M. T., Schuetz, A., Stewart, W. F., & Sun, J. (2016). Doctor AI: Predicting Clinical Events via Recurrent Neural Networks. Proceedings of the 1st Machine Learning for Healthcare Conference, 301–318. https://doi.org/10.48550/arXiv.1511.05942

Ng, K., Sun, J., Hu, J., Wang, F., & Lee, J. (2017). Predicting Unplanned Readmissions Using Electronic Health Records. International Journal of Medical Informatics, 105, 118–124. https://doi.org/10.1016/j.ijmedinf.2017.05.009

Futoma, J., Hariharan, S., Sendak, M., Brajer, N., Clement, M., Bedoya, A., ... & Heller, K. (2017). An Improved Multi-Output Gaussian Process RNN with Real-Time Validation for Early Sepsis Detection. Proceedings of the 34th International Conference on Machine Learning, 70, 1174–1182. https://doi.org/10.48550/arXiv.1708.05894

Che, Z., Purushotham, S., Cho, K., Sontag, D., & Liu, Y. (2018). Interpretable Deep Models for ICU Outcome Prediction. AMIA Annual Symposium Proceedings, 2018, 353–362. https://doi.org/10.1093/jamiaopen/ooy030

Lipton, Z. C., Kale, D. C., Elkan, C., & Wetzel, R. (2016). Learning to Diagnose with LSTM Recurrent Neural Networks. arXiv preprint arXiv:1511.03677. https://doi.org/10.48550/arXiv.1511.03677

Razavian, N., Marcus, J., & Sontag, D. (2016). Multi-task Prediction of Disease Onsets from Longitudinal Lab Tests. Proceedings of the 1st Machine Learning for Healthcare Conference, 73–100. https://doi.org/10.48550/arXiv.1608.00647

Li, Y., Rao, S., Solares, J. R. A., Hassaine, A., Ramakrishnan, R., Canoy, D., ... & Rahimi, K. (2020). BEHRT: Transformer for electronic health records. Scientific Reports, 10(1), 7155. https://doi.org/10.1038/s41598-020-62922-y

Choi, E., Bahadori, M. T., Song, L., Stewart, W. F., & Sun, J. (2017). GRAM: Graph-based attention model for healthcare representation learning. Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 787–795. https://doi.org/10.1145/3097983.3098126

McKinney, S. M., Sieniek, M., Godbole, V., Godwin, J., Antropova, N., Ashrafian, H., Back, T., Chesus, M., Corrado, G. S., Darzi, A., Etemadi, M., Garcia-Vicente, F., Gilbert, F. J., Halling-Brown, M., Hassabis, D., Jansen, S., Karthikesalingam, A., Kelly, C. J., King, D., ... Shetty, S. (2020). International evaluation of an AI system for breast cancer screening. Nature, 577(7788), 89–94. https://doi.org/10.1038/s41586-019-1799-6

Purushotham, S., Meng, C., Che, Z., & Liu, Y. (2018). Benchmarking deep learning models on large healthcare datasets. JBI, 83, 112–124. https://doi.org/10.1016/j.jbi.2018.06.011

Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm. Science, 366(6464), 447 453. https://doi.org/10.1126/science.aax2342

Shortreed, S. M., Laber, E., Lizotte, D. J., Stroup, T. S., Pineau, J., & Murphy, S. A. (2021). Informing sequential clinical decision-making through reinforcement learning: An empirical study. Journal of Machine Learning Research, 22(1), 1–34. 10.5555/3454287.3454290

Tran, T., Nguyen, T. D., Phung, D., & Venkatesh, S. (2021). Learning from clinical judgments: Semi-supervised learning with label noise for psychiatric risk prediction. Journal of Biomedical Informatics, 115, 103688. 10.1016/j.jbi.2021.103688