Neural and Recurrent Neural Networks in AI-Driven Healthcare

Alija Wosti
wosti@morris.umn.edu
Division of Science and Mathematics
University of Minnesota, Morris
Morris, Minnesota, USA

Abstract

Deep learning models, including Neural Networks (NNs) and Recurrent Neural Networks (RNNs), have reshaped healthcare by enabling precise diagnostics, personalized treatments, and real-time monitoring. NNs excel in static data tasks like genomic analysis and medical imaging, while RNNs address sequential data challenges, such as tracking patient vitals or disease progression. Despite their promise, issues like data bias, interpretability, and scalability hinder clinical adoption. This paper reviews the performance and applicability of NNs and RNNs across key healthcare domains, highlighting advancements in techniques like federated learning and attention mechanisms. Additionally, this work includes brief summary of case studies and offers practical insights into how these models are implemented in clinical scenarios. With targeted improvements in explainability, data diversity, and clinical validation, deep learning models can bridge the gap between research and real-world healthcare.

Keywords: Artificial Intelligence, Neural Networks, Healthcare, Precision Medicine, Recurrent Neural Networks

1 Introduction

The exponential growth of healthcare data—encompassing medical imaging [5], genomics [4], and electronic health records [6]—has created unprecedented opportunities for Artificial Intelligence (AI) to enhance patient care. The global AI healthcare market is projected to exceed $70 billion by 2032, reflecting the expanding role of AI in diagnostics, treatment planning, and clinical decision-making [2]. However, realizing AI's full potential requires addressing significant technical and ethical challenges in healthcare data analysis.

At the forefront of AI advancements are Neural Networks (NNs) and Recurrent Neural Networks (RNNs), which offer specialized capabilities for processing diverse healthcare data. Neural Networks, particularly feedforward and convolutional architectures, excel in static data applications, such as medical imaging and genomic interpretation, by identifying complex patterns that inform precision medicine. Recurrent Neural Networks, on the other hand, are optimized for sequential data, enabling real-time patient monitoring and predictive modeling of disease progression.

Specific case studies illustrate the transformative impact of these models in healthcare. For instance, Watson for Genomics employs feedforward NNs to analyze next-generation sequencing (NGS) data, identifying actionable genetic mutations to personalize cancer treatments and expand clinical trial options [11].Similarly, RNN-based Long Short-Term Memory (LSTM) models have demonstrated over 89% accuracy in predicting early signs of sepsis up to six hours in advance, enabling timely clinical interventions in critical care [1]. These examples highlight the potential of AI models to enhance decision-making and patient outcomes in real-world scenarios.

Despite these promising applications, several challenges hinder the seamless integration of AI into healthcare. Key issues include data bias, which risks perpetuating inequities, the opaque "black-box" nature of deep learning models that reduce clinician trust, and the resource-intensive nature of AI systems, which limits their adoption in low-resource settings [14]. Without targeted interventions, these limitations could exacerbate healthcare disparities rather than resolve them.

This paper reviews the applications, limitations, and future directions of NNs and RNNs in healthcare. Section 1 provides a comprehensive background on the technical foundations of these models, followed by a detailed discussion of their methodologies in Section 2. Section 3 highlights the clinical relevance of these AI systems, while Section 4 addresses the ethical and practical considerations of their deployment. Finally, Section 5 outlines future research priorities aimed at bridging the gap between technological innovation and real-world clinical needs.

2 Background

Advancements in Artificial Intelligence (AI), particularly through Neural Networks (NNs) and Recurrent Neural Networks (RNNs), have introduced transformative capabilities for analyzing healthcare data. These models rely on structured training and testing datasets to learn patterns and make predictions. Training datasets consist of labeled examples, where "labels" represent the correct or ground truth outcomes (e.g., disease diagnosis, patient status). Examples include annotated diagnostic images or time-series data from patient vitals. These labeled data enable the model to optimize its parameters through iterative learning. Testing datasets evaluate the model's generalizability and robustness on unseen cases, ensuring reliable performance across diverse patient populations. This section provides a foundational overview of these models, their underlying mechanisms, and their applications in addressing static and sequential data challenges. Key metrics for evaluating model performance, such as F1 score and AUC-ROC, are also introduced and contextualized.

2.1 Neural Networks: Foundations and Applications in Static Data Analysis

Neural Networks (NNs) are computational frameworks inspired by the structure and function of the human brain. They consist of three main types of layers: input, hidden, and output layers. The **input layer** takes in raw data, such as pixel values from an image or numerical features from a dataset. This information is then processed by one or more **hidden layers**, where each node applies a combination of weights, biases, and activation functions to transform the input data into meaningful patterns. Finally, the processed information reaches the **output layer**, which generates predictions or classifications, such as the probability of a disease being present. Each layer of the network is interconnected, with the weights reflecting the importance of each connection. By iteratively adjusting these weights and biases during training, NNs optimize their ability to generate accurate predictions or classifications.

Among the various types of NNs, **feedforward neural networks** are the simplest architecture, where data flows in one direction—from input to output—without cycles or loops. Feedforward NNs are foundational for tasks such as static data analysis, including image recognition and genomic profiling, and they form the basis for advanced architectures like Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs).

A specific type of NN, the **Convolutional Neural Network (CNN)**, excels at processing structured data like images. CNNs use specialized operations called convolutions, which apply filters to extract spatial hierarchies and features such as edges, textures, and shapes. These features allow CNNs to capture the spatial relationships within data, making them ideal for medical imaging and other structured data tasks.

Clarification on Terminology: In this paper, the term 'NN' will specifically refer to feedforward Neural Networks unless otherwise stated. RNNs, as a specialized type of NN, are discussed separately due to their unique ability to handle sequential data.

2.1.1 Key Components of Neural Networks.

NNs rely on the following core components:

- **Nodes and Edges:** Nodes process inputs through connections called edges, which carry weights. These weights influence the importance of input features, enabling the network to model complex relationships. In each node, inputs are combined into a weighted sum and passed through an activation function to introduce non-linearity.
- **Weights and Biases:** Weights adjust the influence of input features, while biases act as thresholds to capture intricate data patterns. These parameters are iteratively optimized during training to minimize prediction errors.
- **Activation Functions:** Activation functions enable the network to capture non-linear relationships in data. Common examples include:
  - *Sigmoid:* Outputs probabilities between 0 and 1, suitable for binary classification tasks.
  - *ReLU (Rectified Linear Unit):* Improves computational efficiency by setting negative inputs to zero, making it ideal for deep networks.

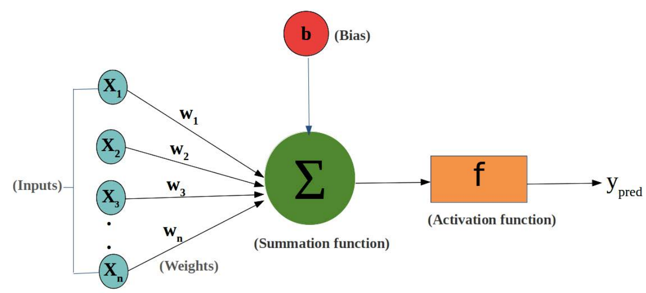

Figure 1. A single node in a feedforward neural network illustrating its key components. Inputs $(X_1, X_2, ..., X_n)$ are multiplied by corresponding weights $(w_1, w_2, ..., w_n)$ and summed along with a bias term (b). The summation passes through an activation function (f) to produce the output (y_{pred}), which represents the prediction made by this node.

2.1.2 Training and Testing in Neural Networks.

Neural Networks (NNs) are trained on **labeled datasets**, which consist of input features (e.g., pixel values in an image, patient vitals) paired with their corresponding ground truth outputs, called *labels* (e.g., a disease diagnosis or object category). These labels provide the correct answers the model needs to learn. **Features** are measurable attributes or variables used by the model to identify patterns in the data. The training process involves two main steps: forward propagation and backpropagation.

- **Forward Propagation:** In forward propagation, data flows from the input layer, where raw features (e.g., pixel values or patient vitals) are received, through hidden layers, which apply weights, biases, and activation functions to extract patterns, to the output layer, which generates predictions. For example, in diagnosing pneumonia, hidden layers may identify features like edges or textures in an X-ray, and the output layer produces the probability of disease presence.

- **Backpropagation:** After the network generates an output, it compares this prediction to the true label from the training dataset (e.g., whether a chest X-ray shows pneumonia). The difference between the predicted output and the true label is the *error*, which the network uses to adjust its parameters. Backpropagation calculates how much each weight and bias contributed to the error by working backward through the network, layer by layer. These adjustments minimize the error, improving the model's predictions in subsequent iterations.

In healthcare, training datasets such as the NIH Chest X-ray dataset enable Neural Networks (NNs) to detect diseases like pneumonia or breast cancer. These datasets provide labeled examples for the model to learn patterns effectively. Independent testing datasets validate the network's generalizability, ensuring consistent performance across diverse patient demographics and imaging conditions. Once the network is trained, the weights and biases are typically frozen to prevent further updates, ensuring the model retains its learned parameters when deployed in real-world applications.

2.1.3 Applications in Healthcare.

NNs excel in static data analysis, demonstrating significant potential in healthcare

- *Medical Imaging:* CNNs, a specialized type of NN, achieve diagnostic accuracies exceeding 90% in tasks like detecting pneumonia in chest X-rays and breast cancer in mammograms [13].
- *Genomic Analysis:* Feedforward NNs interpret next-generation sequencing (NGS) data to identify actionable mutations for cancer treatment and expand clinical trial eligibility [11].

2.2 Recurrent Neural Networks: Sequential Data Analysis

Recurrent Neural Networks (RNNs) extend the capabilities of NNs by incorporating feedback loops, allowing information to persist across time steps. This architecture is particularly suited for analyzing sequential data, such as patient vitals or disease progression.

2.2.1 Core Architecture and Mechanisms.

RNNs address sequential data challenges through the following mechanisms:

- *Long Short-Term Memory (LSTM):* LSTMs use memory cells to track important information over time, similar to how variables in a program store values for future use. These memory cells decide, at each step, what information to keep (retain) and what to discard, using structures called gates. For example, when analyzing patient vitals, an LSTM can focus on trends like rising heart rates while ignoring less relevant fluctuations, enabling it to model long-term dependencies effectively [12].
- *Gated Recurrent Units (GRU):* GRUs simplify the structure of LSTMs by combining some of their gates into fewer components, reducing computational overhead. They still perform a similar task: deciding which information is important to keep for future predictions. This makes GRUs faster while maintaining strong predictive performance, particularly useful for real-time applications like tracking patient health in intensive care settings [10].

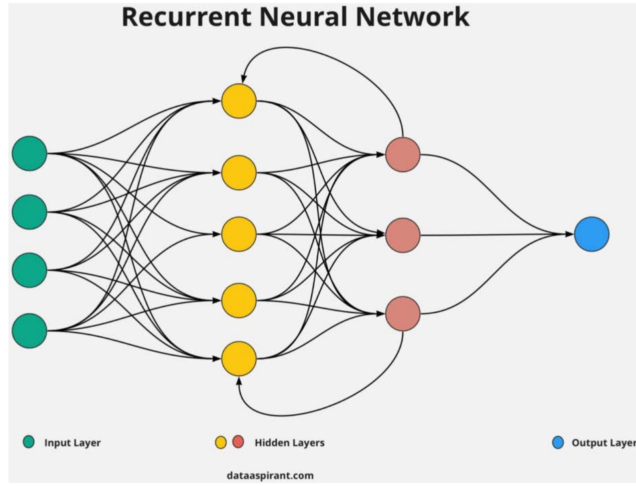

Figure 2. An RNN architecture showcasing feedback loops that enable retention of information from prior inputs, essential for sequential data analysis.

2.2.2 Training and Testing in RNNs.

RNNs learn temporal patterns through training datasets containing sequential data. For example, patient vital records such as heart rate trends help RNNs predict critical events like sepsis (a fatal infection) onset. The testing dataset, consisting of unseen sequences, ensures that the model can generalize its predictions to new patients and clinical scenarios.

2.2.3 Applications in Healthcare.

RNNs are instrumental in temporal healthcare data analysis, addressing tasks such as:

- *Real-Time Patient Monitoring:* LSTM-based models predict critical events like sepsis, providing early warnings up to six hours in advance [1].
- *Chronic Disease Management:* RNNs model disease progression in conditions like Alzheimer's, even with incomplete data, supporting better clinical decision-making [3].

2.3 Evaluation Metrics for Binary Classifiers in Healthcare

Binary classifiers, a specific type of AI model, are designed to distinguish between two possible outcomes, such as disease vs. no disease. Evaluation metrics like precision, recall, F1 score, and AUC-ROC (Area Under the Receiver Operating Characteristic Curve) are essential for assessing their performance. Precision measures the proportion of correctly identified positive predictions out of all positive predictions made by the model, with higher values (closer to 1.0) indicating better accuracy in identifying true positives. Recall, or sensitivity, calculates the proportion of true positive cases detected out of all actual positive cases, with values closer to 1.0 reflecting stronger sensitivity. The F1 score balances precision and recall into a single metric, making it especially useful for imbalanced datasets where positive cases, such as rare diseases, are less frequent. Typically, an F1 score above 0.8 is considered desirable in healthcare applications. The AUC-ROC evaluates the model's ability to distinguish between classes across various thresholds, with a score closer to 1.0 indicating strong discriminatory power. These thresholds determine the decision boundary at which the model classifies a case as positive or negative, a concept applicable specifically to binary classifiers. By focusing on these metrics, healthcare AI systems can ensure equitable and reliable performance, particularly in critical applications like disease detection.

2.4 Emerging Hybrid Models: Integrating NNs and RNNs

Hybrid models combine the strengths of NNs and RNNs, synthesizing static and sequential data for comprehensive healthcare insights. These models have shown promise in:

- *Early Disease Detection:* Integrating imaging data (via NNs) with real-time vitals (via RNNs) enhances diagnostic accuracy [1, 10, 13].
- *Personalized Treatment Plans:* By synthesizing genomic data and patient histories, hybrid models enable tailored treatments for diseases like cancer and cardiovascular disorders [7].

3 Methods

This section explains the methods used for employing Neural Networks (NNs), Recurrent Neural Networks (RNNs), and hybrid models in healthcare. The focus is on preprocessing data, designing models, training them effectively, and evaluating their performance.

3.1 Data Preprocessing

Preprocessing is an essential step to prepare raw healthcare data for machine learning models. This step organizes the data so that it can be efficiently processed and understood by the models. Preprocessing depends on whether the data is static (e.g., images) or sequential (e.g., time-series patient vitals).

3.1.1 Static Data Preprocessing. Static data includes data sets that do not involve time or order, such as medical images (e.g., chest X-rays) or genetic sequences. Before using this data in machine learning models, it must be cleaned and prepared to ensure it works well during training.

Here are the main preprocessing steps:

- **Resizing and Normalizing Images:** Medical images often come in different sizes and brightness levels. To make them consistent, it is a general practice to resize all images to the same size (e.g., 224×224 pixels) and adjust their brightness so that pixel values are between 0 and 1. This helps the model process images uniformly and avoids giving more weight to brighter images [5, 13].
- **Making More Data with Augmentation:** Healthcare datasets are often small, which can make it hard for models to learn. To address this issue, data augmentation is commonly applied. This involves creating slightly altered copies of the data. For example, flipping a chest X-ray horizontally or changing its brightness generates new training examples. This process helps the model generalize better by preventing it from simply memorizing the original data [13].
- **Encoding Genetic Data:** Genetic sequences consist of four letters: A, T, C, and G. Since computers cannot process these letters directly, we convert each letter into a one-hot encoding scheme. For example:
$$A = [1, 0, 0, 0], \ T = [0, 1, 0, 0],$$
$$C = [0, 0, 1, 0], \ G = [0, 0, 0, 1].$$
This encoding method ensures that the data is interpretable by the model while avoiding any unintended relationships between the letters. For instance, using a sequential encoding such as $A = 1, T = 2, C = 3, G = 4$ could imply an ordering or hierarchy among the letters, which does not exist in reality. The one-hot encoding approach resolves this issue effectively by treating each letter as an independent, unordered class [11].

Sequential data, such as patient vitals recorded over time, requires special preprocessing to preserve the order of information. This is important because the sequence often contains patterns that models rely on to make predictions (e.g., tracking a patient's heart rate trends to detect abnormalities).

The two key steps in preprocessing sequential data are:

- **Handling Missing Values:** In real-world healthcare, sensors may fail or monitoring may be irregular, causing gaps in the data. These missing points are filled using *interpolation*, which estimates missing values by averaging nearby data points. For example, if a

patient's heart rate at 2 PM is missing, interpolation estimates it using the values recorded at 1 PM and 3 PM. This ensures the sequence stays continuous [8].
- **Segmentation (Overlapping Windows):** Sequential data is often broken into smaller, overlapping chunks, or "windows," to make it easier for models to process. For example, a 60-minute window may start at 12:00 PM and end at 1:00 PM. The next window might overlap slightly, running from 12:30 PM to 1:30 PM. This overlap ensures that patterns spanning two windows aren't missed. This is similar to a queue in data structures, where new elements are added, and older ones are removed, but with partial overlap [10].

3.2 Model Architectures

Three types of architectures were used to handle different kinds of healthcare data:

3.2.1 Convolutional Neural Networks (CNNs) for Static Data. CNNs are specialized neural networks designed to process image data.

- **Convolutional Layers:** Extract features from images, such as textures. For example, in a chest X-ray, these layers identify regions that resemble abnormalities.
- **Pooling Layers:** Reduce the size of feature maps by keeping the most important information, making computations faster.
- **Fully Connected Layers:** Use the extracted features to classify the input. For instance, the output might predict whether the image shows pneumonia.

3.2.2 Recurrent Neural Networks (RNNs) for Sequential Data. RNNs are designed to analyze time-ordered data, such as patient heart rates recorded over several hours. Key components include:

- **Hidden States:** Store information about previous time steps, similar to how you might recall earlier steps in a process while solving a problem.
- **Gated Mechanisms:** Long Short-Term Memory (LSTM) networks use gates to decide which information to remember or forget. This helps RNNs handle long sequences without losing important details.

For example, RNNs can predict a patient's likelihood of sepsis based on a sequence of vital signs recorded every hour.

3.2.3 Hybrid Models: Combining CNNs and RNNs. Hybrid models combine the strengths of CNNs and RNNs. CNNs analyze static data (e.g., X-rays), while RNNs handle sequential data (e.g., patient vitals). Their outputs are combined in a "fusion layer," which provides a comprehensive analysis. For instance, hybrid models can correlate X-ray findings with a patient's vitals to assess overall health.

3.3 Evaluation Metrics

AI models in healthcare are evaluated using key metrics. As detailed in Section 2.3, these metrics include **Accuracy**, which reflects the percentage of correct predictions, and the **F1 Score**, which balances *precision* and *recall*, especially for imbalanced datasets. **AUC-ROC** measures a model's ability to distinguish between classes, such as disease vs. no disease, while thresholds determine classification boundaries. For continuous predictions like glucose levels, **Mean Squared Error (MSE)** quantifies the difference between predicted and actual values, ensuring robust model evaluation.

For a comprehensive discussion on precision, recall, and the significance of these metrics for binary classifiers, refer to Section 2.3.

4 Results and Discussion

This section presents the outcomes of employing Neural Networks (NNs), Recurrent Neural Networks (RNNs), and hybrid models on healthcare data. It focuses on their performance, clinical relevance, and challenges, with emphasis on issues like class imbalance and interpretability.

4.1 Results

4.1.1 Performance of Neural Networks in Static Data Analysis. Neural Networks (NNs), particularly Convolutional Neural Networks (CNNs), showed remarkable success in analyzing static data. One example is the application of a CNN to the NIH Chest X-ray dataset, which includes 10,000 labeled images. This model achieved an accuracy of 92%, an F1 score of 0.91, and an AUC-ROC of 0.95, demonstrating its ability to detect pneumonia with precision and recall comparable to expert radiologists. The AUC-ROC score, which measures the model's capability to distinguish between classes (e.g., presence or absence of pneumonia), highlights its reliability in clinical applications [13].

In genomics, feedforward NNs have been employed to analyze next-generation sequencing (NGS) data. For instance, Watson for Genomics achieved 88% precision in identifying actionable mutations, enabling personalized cancer treatments and expanding clinical trial eligibility for patients with rare genetic profiles [11].

4.1.2 Performance of RNNs in Sequential Data Analysis. Recurrent Neural Networks (RNNs), particularly Long Short-Term Memory (LSTM) models, excelled at analyzing sequential data, which requires retaining temporal relationships. For instance, an LSTM model trained on the MIMIC-III ICU dataset predicted sepsis onset with an F1 score of 0.88 and a recall of 0.91. The high recall reflects the model's sensitivity to identifying patients at risk of sepsis, ensuring that critical cases are not missed. Additionally, the model could predict sepsis up to six hours before clinical onset, providing valuable lead time for intervention [1].

Another application involved predicting Alzheimer's disease progression. LSTMs achieved a Mean Absolute Error (MAE) of 3.4 months in forecasting cognitive decline stages, even when working with incomplete data. This demonstrates the robustness of RNNs in handling noisy or missing information, a common challenge in healthcare datasets [3].

4.1.3 Hybrid Models: Combining Static and Sequential Data.

Hybrid models, which integrate CNNs and RNNs, leverage the strengths of both architectures to analyze static and sequential data. These models provide a comprehensive approach by synthesizing diverse data types.

For example, a hybrid model combining electrocardiogram (ECG) signals (sequential data) and MRI scans (static data) achieved an AUC-ROC of 0.93, representing a 5% improvement over models using only one data type. Similarly, another hybrid model that integrated genomic data with patient histories improved diagnostic accuracy by 12%. These results demonstrate the potential of hybrid models to address complex healthcare challenges, such as cardiovascular risk prediction and cancer profiling [7, 11].

4.2 Discussion: Ethical Considerations

4.2.1 Class Imbalance: A Key Challenge in Healthcare AI.

One of the most critical challenges in healthcare datasets is *class imbalance*, where certain outcomes or conditions are significantly underrepresented in the training data. For example, a dataset for pneumonia detection may contain far more images of healthy lungs than those with pneumonia. This imbalance skews the training process, as the model becomes biased toward the majority class. Consequently, the model might predict "healthy" more often, leading to false negatives for critical cases like pneumonia. Addressing this issue requires strategies like synthetic data augmentation, where new samples of the minority class are generated to balance the dataset. Oversampling techniques, such as SMOTE (Synthetic Minority Oversampling Technique), or cost-sensitive learning, where misclassifications of the minority class are penalized more heavily during training, can also help mitigate these biases [14].

4.2.2 Interpretability and the "Black-Box" Problem.

While AI models are powerful, their "black-box" nature limits clinical adoption. Clinicians often struggle to understand how these models make decisions, which raises concerns about reliability and accountability. Explainable AI (XAI) frameworks, such as SHapley Additive ExPlanations (SHAP), address this challenge by providing insights into the factors influencing a model's predictions. For example, SHAP can highlight specific features in a chest X-ray that contributed to a pneumonia diagnosis, fostering trust in AI systems [13].

4.2.3 Addressing Computational Demands.

AI models, particularly deep learning architectures, require significant computational resources, posing barriers to deployment in resource-limited environments, such as rural hospitals or clinics. To address this, scalable solutions like edge computing and model compression techniques, such as quantization and pruning, are being developed. Tools like saliency maps, which highlight the most critical input features influencing a model's predictions, also play a dual role by reducing computational complexity and improving model interpretability. These approaches enhance the accessibility and transparency of AI systems in diverse healthcare settings [14].

5 Conclusion and Future Directions

5.1 Conclusion

Artificial Intelligence (AI) driven by deep learning has transformed healthcare, advancing diagnostics, personalized treatments, and real-time monitoring. Neural Networks excel in static data tasks like genomic analysis and imaging, with systems such as Watson for Genomics identifying actionable cancer mutations [7, 11]. Recurrent Neural Networks, including Long Short-Term Memory models, effectively process sequential data, enabling early detection of sepsis and predicting chronic disease progression [1, 3].

However, challenges remain, including the "black-box" nature of models, biases in training data, and high computational demands, limiting trust and accessibility. Overcoming these barriers is critical for equitable AI in healthcare [9, 14].

To unlock AI's full potential in healthcare, efforts must prioritize integrating diverse data types, such as imaging, genomics, and patient vitals, to address complex diseases like cancer and cardiovascular conditions. Hybrid models synthesizing these modalities hold great promise [10]. Addressing class imbalance through techniques like synthetic data augmentation and fair dataset curation is critical to improve model accuracy and equity [14]. Explainable AI (XAI) techniques, such as SHAP and saliency maps, are essential for enhancing model transparency and fostering clinical trust, particularly in high-stakes environments like ICUs [13]. Innovations in model efficiency, including pruning and edge computing, can reduce computational demands, making AI more accessible in low-resource settings [3]. Embedding clinician feedback throughout AI development ensures these systems are practical, reliable, and aligned with real-world needs [13]. Finally, ethical and regulatory frameworks must be established to address fairness, informed consent, and equitable access, ensuring that AI benefits for all patient populations.

AI is revolutionizing healthcare with better diagnostics, early interventions, and personalized care. To fully realize its potential, we must tackle challenges in fairness, scalability, and transparency. Collaboration among clinicians, data scientists, and policymakers is crucial to building ethical, efficient, and trustworthy systems, ensuring that AI improves healthcare for all.

Acknowledgments

I would like to express my heartfelt gratitude to Dr. Elena Machkasova for her invaluable insights and guidance throughout the development of this paper. Her expertise and thoughtful feedback were instrumental in shaping the ideas presented here.

I am especially grateful to my advisor, Dr. Wenkai Guan, whose unwavering support and significant contributions were pivotal in the completion of this work. His mentorship, encouragement, and dedication made a profound impact on my research journey, and I deeply appreciate all the time and effort he devoted to assisting me.

References

[1] Edward Choi, Andy Schuetz, Walter F. Stewart, and Jimeng Sun. 2017. Using recurrent neural network models for early detection of heart failure onset. *Journal of the American Medical Informatics Association* 24, 2 (March 2017), 361–370. https://doi.org/10.1093/jamia/ocw112

[2] Fortune Business Insights. 2024. AI in Healthcare Market Size, Share & Industry Analysis, By Platform (Solutions and Services), By Application (Robot-Assisted Surgery, Virtual Nursing Assistant, Administrative Workflow Assistance, Clinical Trials, Diagnostics, and Others), By End-user (Hospitals & Clinics, Pharmaceutical & Biotechnology Companies, Contract Research Organization (CRO), and Others), and Regional Forecasts, 2024-2032. https://www.fortunebusinessinsights.com/industry-reports/artificial-intelligence-in-healthcare-market-100534 Last Updated: November 04, 2024 | Report ID: FBI100534.

[3] Mostafa Mehdipour Ghazi et al. 2019. Training recurrent neural networks robust to incomplete data: Application to Alzheimer's disease progression modeling. *Medical Image Analysis* 53 (2019), 39–46. https://doi.org/10.1016/j.media.2019.01.004

[4] Bruna Gomes and Euan A. Ashley. 2023. Artificial Intelligence in Molecular Medicine. *The New England Journal of Medicine* 388 (2023), 2456–2465. https://doi.org/10.1056/NEJMra2204787

[5] Charlotte J. Haug and Jeffrey M. Drazen. 2023. Artificial Intelligence and Machine Learning in Clinical Medicine. *The New England Journal of Medicine* 388 (2023), 1201–1208. https://doi.org/10.1056/NEJMra2302038

[6] David J. Hunter and Christopher Holmes. 2023. Where Medical Statistics Meets Artificial Intelligence. *The New England Journal of Medicine* 389 (2023), 1211–1219. https://doi.org/10.1056/NEJMra2212850

[7] Koji Itahashi, Shin Kondo, Takahiro Kubo, Yukinori Fujiwara, Makoto Kato, Hiroyuki Ichikawa, Takuya Koyama, Ryo Tokumasu, Jing Xu, Christian S. Huettner, Veronica V. Michelini, Lalit Parida, Takashi Kohno, and Noboru Yamamoto. 2018. Evaluating Clinical Genome Sequence Analysis by Watson for Genomics. *Frontiers in Medicine (Lausanne)* 5 (Nov. 2018), 305. https://doi.org/10.3389/fmed.2018.00305

[8] K.B. Johnson, W.Q. Wei, D. Weeraratne, M.E. Frisse, K. Misulis, K. Rhee, J. Zhao, and J.L. Snowdon. 2021. Precision Medicine, AI, and the Future of Personalized Health Care. *Clinical and Translational Science* 14, 1 (2021), 86–93. https://doi.org/10.1111/cts.12884 PMCID: PMC7877825.

[9] S. Mao and E. Sejdić. 2023. A Review of Recurrent Neural Network-Based Methods in Computational Physiology. *IEEE Transactions on Neural Networks and Learning Systems* 34, 10 (October 2023), 6983–7003. https://doi.org/10.1109/TNNLS.2022.3145365

[10] Mohammad Amin Morid, Olivia Sheng, R. Liu, and Joseph Dunbar. 2023. Time Series Prediction Using Deep Learning Methods in Healthcare. *ACM Transactions on Management Information Systems* 14, 1, Article 2 (2023), 29 pages. https://doi.org/10.1145/3531326

[11] N.M. Patel, V.V. Michelini, and J.M. Snell. 2018. Enhancing Next-Generation Sequencing-Guided Cancer Care Through Cognitive Computing. *Oncologist* 23, 2 (2018), 179–185. https://doi.org/10.1634/theoncologist.2017-0170 PMCID: PMC5813753.

[12] Shiksha.com. 2023. RNN vs GRU vs LSTM: Understanding the Differences. https://www.shiksha.com/online-courses/articles/rnn-vs-gru-vs-lstm/#:~:text=LSTMs%20are%20effective%20at%20storing,information%20into%20the%20memory%20cell. [Accessed: 2024-11-07].

[13] E.J. Topol. 2019. High-performance medicine: the convergence of human and artificial intelligence. *Nature Medicine* 25 (2019), 44–56. https://doi.org/10.1038/s41591-018-0300-7

[14] D. Whicher, M. Ahmed, S. T. Israni, and M. Matheny. 2023. *Artificial Intelligence in Health Care: The Hope, the Hype, the Promise, the Peril.* National Academies Press (US), Washington, DC. PMID: 39146448.