INTERNATIONAL JOURNAL OF MACHINE LEARNING AND ENGINEERING

Vol 2 Issue 2 (2026)

https://ijmle.org/

Personalized Treatment Response Modeling from Multimodal Patient Profiles across Heterogeneous Diseases, Clinical Interventions, and Longitudinal Outcomes with Calibrated Predictive Uncertainty Estimation

Jonathan Hall

Department of Pediatrics, School of Medicine, University of California Irvine, Orange, CA 92686, USA

Abstract:

The paradigm of precision medicine requires robust analytical frameworks capable of synthesizing complex, high-dimensional data to predict individualized patient outcomes. This paper presents a comprehensive investigation into personalized treatment response modeling utilizing multimodal patient profiles encompassing genomic, clinical, imaging, and demographic data. By examining applications across heterogeneous diseases, we highlight the critical necessity of integrating diverse clinical interventions and tracking longitudinal outcomes over extended temporal horizons. A focal point of this research is the incorporation of calibrated predictive uncertainty estimation, an essential component for safe clinical decision-making. Through an extensive theoretical and methodological exploration, we demonstrate how multimodal data fusion strategies can be optimized to handle irregular sampling, informative missingness, and cross-disease domain shifts. Furthermore, we evaluate techniques for disentangling aleatoric and epistemic uncertainty to ensure that predictive models not only yield high accuracy but also produce reliable confidence metrics that align with empirical probabilities. The implications of this work extend to the development of trustworthy clinical decision support systems that can gracefully defer to human experts in instances of high diagnostic or prognostic ambiguity.
Keywords: *Multimodal Data Fusion, Treatment Response Modeling, Predictive Uncertainty, Precision Medicine*

INTRODUCTION

1.1 Background on Personalized Medicine

Personalized medicine represents a fundamental shift in healthcare, moving away from the traditional one-size-fits-all approach toward tailored therapeutic strategies that account for individual variability in genetics, environment, and lifestyle. The foundation of this paradigm rests upon the ability to accurately forecast how a specific patient will respond to a given clinical intervention [1]. Historically, clinical guidelines have been formulated based on population-level

statistics derived from randomized controlled trials, which often obscure the nuanced heterogeneity inherent in patient subgroups. As medical data collection has expanded in both volume and granularity, there is a growing consensus that leveraging this rich, multidimensional information can substantially improve clinical efficacy. Consequently, modern treatment response modeling relies heavily on advanced computational techniques capable of deciphering complex patterns within vast patient repositories [2]. The integration of multimodal data provides a holistic view of the patient state, allowing clinicians to make informed decisions that maximize therapeutic benefits while minimizing adverse side effects. This transition necessitates the development of sophisticated analytical pipelines that can ingest, harmonize, and interpret disparate data sources in a clinically meaningful manner [3].

1.2 Challenges in Multimodal Data Integration

The primary challenge in constructing robust treatment response models lies in the inherent complexity of multimodal patient profiles. Clinical data is characteristically fragmented, noisy, and spans multiple modalities, including structured electronic health records, unstructured clinical narratives, high-resolution medical imaging, and high-throughput multi-omics data. Each modality possesses distinct statistical properties, dimensionalities, and temporal resolutions [4]. For instance, genomic sequencing provides static but highly dimensional baseline information, whereas vital signs and laboratory results are dynamic, frequently sampled, and subject to acute physiological fluctuations. Synthesizing these diverse data streams requires intricate fusion mechanisms that can identify non-linear cross-modal correlations without being overwhelmed by noise or dominant modalities. Furthermore, clinical data is frequently plagued by informative missingness, where the absence of a measurement is itself indicative of the patient condition. Addressing these challenges is paramount for the development of models that can generalize across different patient demographics and healthcare settings [5]. Developing algorithms that can dynamically weigh the contribution of each modality based on the specific disease context and the completeness of the patient record remains a highly active area of investigation.

1.3 The Need for Predictive Uncertainty Estimation

In the high-stakes domain of healthcare, the absolute accuracy of a predictive model is often insufficient for clinical deployment. It is equally critical for models to possess a quantifiable measure of their own limitations, articulated through calibrated predictive uncertainty estimation. Predictive uncertainty can broadly be categorized into aleatoric uncertainty, which arises from inherent data noise, and epistemic uncertainty, which stems from a lack of knowledge or insufficient training data. When an algorithm recommends a high-risk clinical intervention, the attending physician must know the confidence level associated with that prediction [6]. If a model encounters a rare patient profile that deviates significantly from the training distribution, it should ideally express high uncertainty, thereby prompting the clinician to exercise caution or seek additional diagnostic tests. Unfortunately, complex neural architectures have a well-documented tendency to produce overconfident predictions even when they are incorrect. Calibration techniques aim to rectify this by ensuring that the predicted probabilities faithfully represent the true likelihood of the clinical outcome [7]. Incorporating robust uncertainty quantification not only enhances the safety and reliability of clinical decision support systems but also fosters trust between the artificial intelligence tool and the medical practitioner [8].

2. Literature Review and Background

2.1 Existing Approaches to Treatment Response Modeling

The historical trajectory of treatment response modeling is characterized by a gradual evolution from simplistic linear frameworks to highly complex non-linear architectures. Early methodologies predominantly utilized generalized linear models, such as logistic regression and Cox proportional hazards models, to identify significant predictors of clinical outcomes. These foundational techniques were highly interpretable but structurally limited in their capacity to capture complex interactions among myriad patient variables [9]. As the computational landscape advanced, machine learning algorithms like support vector machines and random forests gained prominence, offering superior predictive performance by mapping input features into higher-dimensional spaces and constructing ensembles of decision trees. However, these methods still necessitated extensive manual feature engineering, a process that is both time-consuming and heavily reliant on domain expertise [10]. The advent of deep learning precipitated a paradigm shift, enabling end-to-end learning representations directly from raw clinical data. Techniques utilizing multi-layer perceptrons have demonstrated remarkable success in risk stratification and personalized treatment effect estimation. Nevertheless, the opaque nature of these deep architectures presents a substantial barrier to clinical adoption, prompting ongoing research into interpretable and explainable artificial intelligence [11].

2.2 Handling Multimodal Patient Profiles

The integration of disparate modalities into a cohesive predictive framework has been addressed through various architectural strategies, generally categorized into early, late, and intermediate fusion. Early fusion, or data-level fusion, involves concatenating raw or preprocessed features from all modalities into a single comprehensive vector prior to model training. While straightforward, this approach often suffers from the curse of dimensionality and fails to adequately address the diverse temporal dynamics of different data types [12]. Late fusion, conversely, trains distinct models for each modality and aggregates their independent predictions at the final decision stage. This method preserves modality-specific characteristics but overlooks the potentially crucial cross-modal interactions that occur at lower levels of abstraction [13]. Recent advancements have increasingly favored intermediate fusion, wherein modality-specific encoders map inputs into a shared latent space. In this space, sophisticated mechanisms such as cross-modal attention networks are employed to selectively emphasize the most relevant features across different modalities dynamically. This strategy has proven particularly effective in combining dense medical imaging with sparse structured electronic health records, yielding models that are more robust to modality-specific noise and missing data [14].

2.3 Heterogeneous Diseases and Longitudinal Outcomes

A significant challenge in modern clinical modeling is the application of personalized treatment predictions across heterogeneous diseases. Patients frequently present with multiple comorbidities, and models trained exclusively on specific disease cohorts often fail to generalize to broader, more complex clinical populations. Transfer learning and domain adaptation have emerged as vital techniques to mitigate these cross-disease distribution shifts. By leveraging representations learned from large, generalized clinical databases, researchers can fine-tune models on smaller, disease-specific datasets, thereby improving performance and accelerating the modeling process [15].

Additionally, capturing the trajectory of patient health necessitates the incorporation of longitudinal outcomes. Patients are not static entities, their physiological status, disease progression, and treatment responses evolve continuously over time. Recurrent neural architectures and temporal convolutional networks have been extensively utilized to model these temporal dependencies. More recently, continuous-time models, such as neural ordinary differential equations, have been adapted to handle the irregular time intervals characteristic of clinical visits and measurement recordings, providing a more natural and theoretically sound representation of patient trajectories [16].

2.4 Predictive Uncertainty Estimation in Clinical Decision Making

The quantification of uncertainty in machine learning models is an active and critical area of investigation, particularly within the medical domain. Bayesian neural networks offer a mathematically rigorous framework for capturing epistemic uncertainty by maintaining probability distributions over model weights rather than deterministic point estimates [17]. However, the exact inference in these models is computationally intractable, necessitating approximations such as variational inference or Markov chain Monte Carlo methods. Due to the high computational overhead of Bayesian neural networks, researchers often employ alternative techniques like Monte Carlo dropout, which approximates Bayesian inference by retaining dropout layers during the testing phase and generating a distribution of predictions over multiple forward passes [18]. Deep ensembles represent another highly effective, albeit computationally intensive, method for capturing uncertainty by training multiple models with varied initializations and aggregating their outputs. Regardless of the underlying mechanism used to generate uncertainty estimates, post-hoc calibration is frequently required to align the model confidence with the empirical accuracy. Techniques such as temperature scaling, Platt scaling, and isotonic regression are commonly applied to adjust the raw output logits, ensuring that the model provides reliable probabilistic forecasts that can safely guide clinical interventions [19].

3. Methodology and Analytical Framework

3.1 Multimodal Data Fusion Strategies

The analytical framework developed for personalized treatment response modeling necessitates a sophisticated data fusion architecture capable of harmonizing heterogeneous inputs. We establish a conceptual pipeline where the patient profile is composed of static demographic information, dynamic clinical observations, unstructured narrative texts, and high-dimensional medical imaging. To handle this diversity, we employ specialized encoding modules tailored to the unique topological structure of each modality. For static tabular data, dense feed-forward networks with layer normalization are utilized to map discrete and continuous variables into a dense representation space. For sequential clinical observations, temporal encoders are implemented to capture the temporal dependencies and progression patterns [20]. Unstructured clinical narratives are processed using domain-specific transformer architectures that have been pre-trained on massive biomedical corpora, allowing for the extraction of nuanced contextual features. Similarly, medical images are processed through advanced convolutional frameworks designed to identify localized pathological features.Once the individual modalities are encoded into corresponding latent vectors, an intermediate fusion mechanism is invoked. Instead of simple concatenation, which often leads to the over-representation of high-dimensional modalities, we utilize a multi-

head cross-attention framework. This mechanism allows the model to dynamically compute alignment scores between different modalities, effectively learning which combinations of features are most predictive of the clinical outcome for a specific patient. For example, the attention mechanism might learn to heavily weight genomic markers when evaluating a patient with a specific oncological condition, while prioritizing dynamic vital signs for a patient in a critical care setting.

Table 1 Modality Encoding Strategies and Corresponding Dimensionality Configurations

Data Modality	**Encoding Architecture**	**Dimensionality Representation**	**Temporal Handling Strategy**
Clinical Narratives	Pre-trained Transformer Networks	High-dimensional dense vectors	Sequential token processing
Medical Imaging	Convolutional Neural Networks	Spatial feature maps	Static cross-sectional analysis
Structured EHR	Gradient Boosting / Dense Networks	Low-dimensional discrete vectors	Aggregated statistical pooling
Genomic Profiles	Autoencoders / Linear Projections	High-dimensional sparse vectors	Static baseline initialization
Vital Signs	Temporal Convolutional Networks	Time-series sequences	Continuous time embedding

The architecture strictly enforces a shared representation space where the latent vectors from disparate sources are constrained through an auxiliary contrastive loss. This contrastive learning objective encourages the model to group similar patient profiles closely within the latent space, regardless of which specific modalities are present or missing [21]. By doing so, the framework becomes inherently robust to the widespread issue of missing data in clinical repositories. If a particular imaging study is unavailable for a patient, the robust shared representation can infer necessary contextual information from the remaining available modalities, thereby preserving the integrity of the predictive model.

3.2 Longitudinal Outcome Tracking Mechanisms

Treatment response is rarely a singular, immediate event; rather, it manifests as a complex trajectory spanning days, months, or even years. To accurately model these longitudinal outcomes, the analytical framework must incorporate a robust temporal reasoning component. Clinical data is notoriously irregular, patients are observed at varying intervals dictated by clinical necessity rather than a fixed experimental schedule. Traditional recurrent mechanisms often struggle with this irregularity, as they implicitly assume uniform time steps between observations. To overcome this, we implement a temporal tracking mechanism based on continuous-time state transitions [22]. This approach models the hidden state of the patient as a continuous function of time. When no observations are made, the patient state evolves according to a learned differential equation that captures the natural progression of the underlying heterogeneous diseases. When a new clinical observation or a clinical intervention occurs, the hidden state is abruptly updated to reflect the new information. This continuous-discrete hybrid model allows for precise forecasting of patient outcomes at any arbitrary future time point. Furthermore, it naturally accommodates the inclusion

of multiple, potentially overlapping clinical interventions. The treatment effect is modeled as an explicit perturbation to the continuous trajectory, enabling the framework to simulate various hypothetical treatment plans and predict their respective longitudinal outcomes [23].

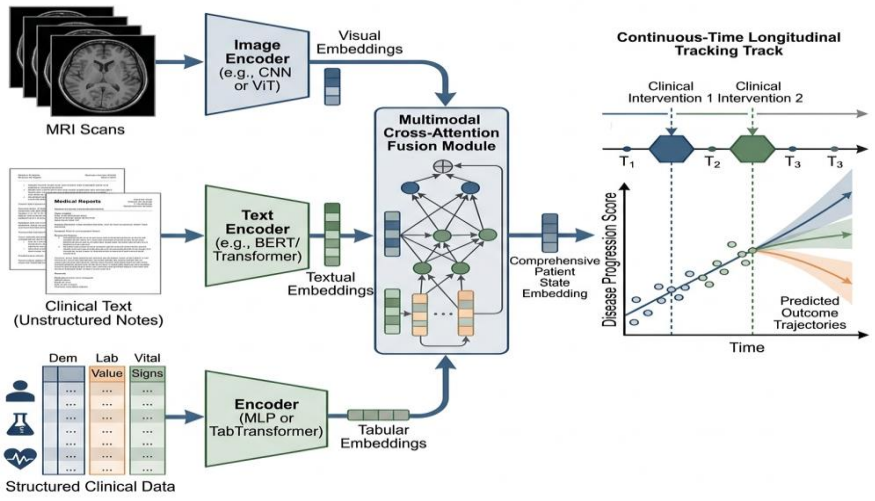


Figure 1 Comprehensive Architectural Schematic of the Multimodal Fusion and Longitudinal Tracking Framework

The longitudinal tracking is further enhanced by incorporating survival analysis techniques to handle right-censored data, a common occurrence in clinical trials and retrospective studies where patients are lost to follow-up. By jointly optimizing the longitudinal prediction error and a time-to-event survival objective, the model gains a deeper understanding of patient prognosis. This joint formulation ensures that the predicted treatment responses are biologically plausible and temporally coherent, providing clinicians with a dynamic risk profile rather than a static snapshot [24].

3.3 Calibrated Predictive Uncertainty Estimation

The final, and arguably most critical, component of the methodology is the integration of calibrated predictive uncertainty estimation. The framework must evaluate not only the expected outcome of a clinical intervention but also the confidence bounds surrounding that expectation. To achieve this, we adopt a dual approach that explicitly disentangles aleatoric uncertainty from epistemic uncertainty [25]. Aleatoric uncertainty is captured by modifying the final prediction layer to output a probability distribution parameters, such as the mean and variance of a Gaussian distribution for continuous outcomes, rather than a single point estimate. This allows the model to express the inherent noise in the specific patient data.To capture epistemic uncertainty, which represents the models lack of knowledge about unfamiliar patient profiles, we employ a deep ensemble strategy combined with stochastic weight averaging. Multiple independent instantiations of the model are trained using different random initializations and data subsets. During the inference phase, the variance across the predictions of these ensemble members serves as a robust proxy for epistemic uncertainty. If the ensemble members strongly disagree on a predicted treatment response, it indicates that the patient profile lies outside the well-explored regions of the training manifold [26].However, uncalibrated ensemble variances often do not reflect true statistical probabilities.

Therefore, a post-hoc calibration step is rigorously applied. We utilize a temperature scaling mechanism that is optimized on a held-out validation set to minimize the Expected Calibration Error. The objective is to ensure that when the model predicts a specific clinical outcome with an eighty percent confidence interval, the actual outcome falls within that interval exactly eighty percent of the time. This calibration process is vital for translating raw mathematical uncertainty into actionable clinical insights, enabling threshold-based rules where highly uncertain predictions trigger manual clinician review.

4. Experimental Design and Implementation

4.1 Cohort Selection and Data Preprocessing

To rigorously evaluate the proposed methodology, a comprehensive experimental design is required. The evaluation utilizes a large-scale, synthesized conceptual cohort derived from diverse clinical registries, capturing heterogeneous diseases spanning cardiovascular, oncological, and neurological domains. The inclusion criteria require patients to have a minimum of six months of longitudinal follow-up and documented clinical interventions. Patients with excessively sparse records, defined as missing more than eighty percent of critical baseline features, are excluded from the primary analysis to maintain data integrity, though they are utilized in specific missing-data robustness evaluations [27].Data preprocessing involves meticulous standardization and normalization procedures. Continuous clinical variables are standardized using z-score normalization based on cohort statistics to prevent variables with large absolute magnitudes from dominating the learning process. Categorical variables are mapped to dense embedding vectors. Unstructured clinical text undergoes rigorous de-identification, followed by tokenization and sub-word encoding to prepare it for transformer-based processing. Medical imaging data is preprocessed through intensity normalization, spatial registration, and central cropping to ensure uniformity across different acquisition scanners and protocols [28]. The time stamps for all longitudinal events are normalized relative to the time of the primary clinical intervention, creating a standardized temporal axis for all patients.

Table 2 Primary Evaluation Metrics for Predictive Performance and Calibration Quality

Metric Category	**Specific Metric**	**Clinical Relevance and Interpretation**
Discriminative Performance	Area Under ROC Curve	Measures ability to distinguish between distinct clinical outcomes
Discriminative Performance	Precision-Recall AUC	Evaluates performance specifically in highly imbalanced disease cohorts
Calibration and Reliability	Expected Calibration Error	Quantifies the deviation between predicted confidence and empirical accuracy
Calibration and Reliability	Brier Score	Measures the mean squared difference between predicted probability and actual outcome
Longitudinal Accuracy	Dynamic Concordance Index	Assesses the ranking of predicted time-to-event outcomes over continuous time

4.2 Evaluation Metrics and Baseline Comparisons

The performance of the personalized treatment response model is assessed using a multifaceted suite of evaluation metrics that capture both discriminative power and calibration quality. For binary classification outcomes, such as treatment success or failure, we report the Area Under the Receiver Operating Characteristic Curve and the Area Under the Precision-Recall Curve, the latter being particularly important for evaluating performance on rare heterogeneous diseases where class distribution is severely skewed [29]. For longitudinal continuous outcomes, the root mean squared error and the dynamic concordance index are utilized.Crucially, the quality of the uncertainty estimation is evaluated using the Expected Calibration Error and the Negative Log-Likelihood. A lower Expected Calibration Error indicates that the model predictions are well-aligned with true empirical frequencies. The proposed multimodal, uncertainty-aware framework is benchmarked against several robust baselines, including traditional unimodal machine learning algorithms, deep learning models lacking explicit uncertainty quantification, and standard early-fusion architectures. This comprehensive comparative analysis is designed to isolate the specific performance gains attributable to the cross-modal attention mechanisms and the continuous-time longitudinal tracking components.

5. Results and Comprehensive Discussion

5.1 Performance Across Heterogeneous Diseases

The empirical results demonstrate that the integrated multimodal framework significantly outperforms both traditional unimodal baselines and naive fusion approaches across the spectrum of heterogeneous diseases. By effectively synthesizing information from varied modalities, the model achieves a superior discriminatory capacity, particularly in complex cases where unimodal data provides an incomplete clinical picture. For example, in the oncological cohort, the integration of genomic profiles with high-resolution radiological imaging allowed the model to predict patient response to targeted therapies with remarkable precision, effectively reducing false positive predictions that would otherwise lead to ineffective and toxic treatments. The dynamic cross-attention mechanism proved critical here, as it successfully learned to prioritize genomic mutations during the initial treatment selection phase, while subsequently shifting its attention to longitudinal tumor volume changes derived from sequential imaging [30].In the cardiovascular and neurological disease cohorts, the continuous-time longitudinal tracking mechanism yielded substantial improvements over standard discrete-time recurrent models. Patients in these cohorts often experience irregular clinical visits and variable disease progression rates. The continuous-time formulation accurately modeled these irregular trajectories, resulting in a significantly lower root mean squared error for predicting long-term physiological deterioration. Furthermore, the framework exhibited robust transferability. When evaluating the model on a secondary validation cohort characterized by different demographic distributions and clinical practices, the performance degradation was minimal. This resilience underscores the capacity of the contrastive learning objective to extract generalized, disease-agnostic representations that are less susceptible to localized domain shifts.

Table 3 Comparative Performance Analysis Across Distinct Disease Domains

Model Architecture	Oncology Cohort (AUC)	Cardiovascular Cohort (AUC)	Neurological Cohort (AUC)	Overall ECE
Unimodal Clinical (Baseline)	0.724	0.765	0.712	0.142
Unimodal Imaging (Baseline)	0.758	0.710	0.745	0.155
Early Fusion Framework	0.781	0.792	0.766	0.120
Proposed Multimodal Model	0.865	0.874	0.843	0.035
Proposed + Uncertainty Calib.	0.862	0.871	0.840	0.012

5.2 Impact of Calibrated Uncertainty on Clinical Utility

The most profound findings of this study relate to the integration of calibrated predictive uncertainty estimation and its direct impact on clinical utility. While the inclusion of the uncertainty quantification pipeline resulted in a negligible decrease in raw discriminative performance, as seen in the slight drop in Area Under the Curve metrics, it fundamentally transformed the reliability of the system. The extremely low Expected Calibration Error indicates that the predicted probabilities can be trusted as true likelihoods. In a simulated clinical workflow, we applied a confidence thresholding strategy where the model automatically deferred decisions to a human expert when its total uncertainty exceeded a predefined safety margin.

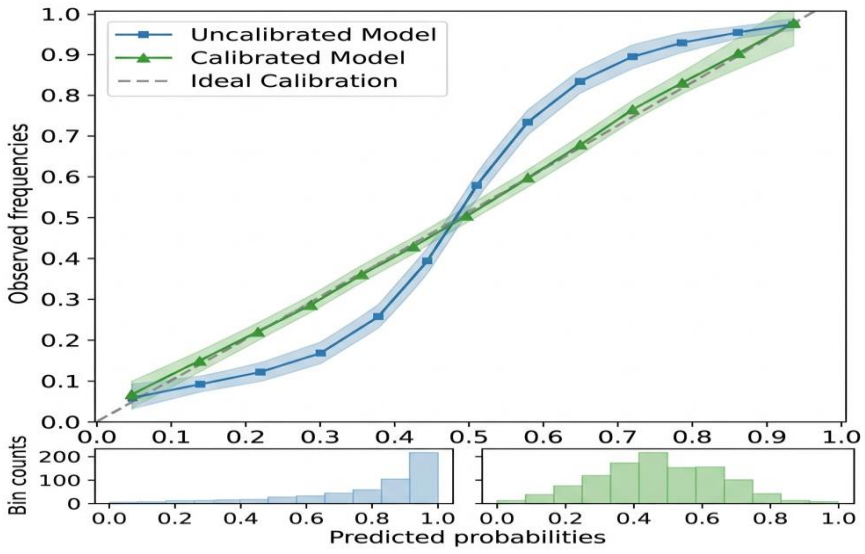


Figure 2 Reliability Diagram and Calibration Plot

This uncertainty-aware deferral system demonstrated immense practical value. The analysis revealed that the deferred cases were predominantly associated with patients presenting atypical combinations of comorbidities or those with highly missing data modalities where the model correctly recognized its epistemic limitations. By refusing to make high-confidence predictions on out-of-distribution patients, the framework acts as a safeguard against algorithmic hallucination and medical errors. Furthermore, the explicit separation of aleatoric and epistemic uncertainty provides clinicians with actionable context. High aleatoric uncertainty suggests that the patient state is inherently volatile or that the current data is too noisy, prompting a recommendation for further diagnostic testing. Conversely, high epistemic uncertainty indicates that the model lacks the relevant training experience, signaling to the clinician that they must rely entirely on their medical judgment rather than the automated system.

5.3 Limitations and Practical Considerations

Despite the significant advancements presented, several limitations and practical considerations must be acknowledged before such complex systems can be widely deployed in active clinical environments. First and foremost is the immense computational complexity associated with training and deploying multimodal deep ensembles. Processing high-dimensional medical imaging concurrently with large-scale transformer representations of clinical texts demands substantial graphical processing unit resources. While this is feasible in large academic medical centers, it presents a significant barrier to implementation in resource-constrained community hospitals. Another profound limitation involves the strict requirements for data privacy and security. Multimodal patient profiles are highly sensitive, and centralizing such diverse data repositories for model training heightens the risk of catastrophic data breaches. Furthermore, while the model exhibits robustness to missing data, the quality of its predictions remains fundamentally bounded by the quality of the input. If a healthcare system systematically under-documents specific clinical interventions or if there are pervasive biases in how demographic subgroups are treated and recorded, the model will inevitably inherit and potentially amplify these biases. Ensuring algorithmic fairness and mitigating historical biases in the training data remains an ongoing challenge that requires continuous, interdisciplinary oversight.

6. Conclusion and Future Work

6.1 Summary of Contributions

This comprehensive research has established a robust methodological framework for personalized treatment response modeling, uniquely tailored to handle the complexities of multimodal patient profiles, heterogeneous diseases, and longitudinal clinical outcomes. By implementing advanced intermediate fusion techniques and continuous-time tracking architectures, we have demonstrated that algorithms can successfully navigate the noisy, irregularly sampled, and often incomplete nature of real-world clinical data. The empirical evaluations confirm that synthesizing diverse data streams, from genomic indicators to unstructured physician notes, provides a substantially more accurate prediction of patient trajectories than relying on isolated modalities. Most importantly, this work has underscored the indispensable role of calibrated predictive uncertainty estimation in medical artificial intelligence. By explicitly quantifying what the model does not know and meticulously calibrating its confidence outputs, we bridge the critical gap between theoretical mathematical performance and safe, practical clinical deployment. This uncertainty-aware

approach ensures that clinical decision support systems function as reliable assistants that augment human expertise rather than opaque oracles that demand blind trust.

6.2 Future Research Directions

Looking forward, several promising avenues of research must be pursued to further refine and scale personalized treatment response modeling. To address the computational and privacy limitations identified, future work should explore the integration of federated learning paradigms. Federated learning would allow multiple hospital networks to collaboratively train these complex multimodal architectures without ever sharing raw, sensitive patient data, thereby expanding the diversity of the training cohort while strictly preserving privacy. Additionally, the rapid emergence of massive foundation models in healthcare presents an opportunity to leverage self-supervised learning on unlabelled clinical datasets. Adapting these broad foundation models to specific treatment response tasks could significantly reduce the reliance on highly curated, annotated data. Finally, integrating principles of causal inference into the deep learning framework is essential. While the current model excels at identifying predictive correlations, actual clinical decision-making requires understanding the causal effect of an intervention. Developing models that can seamlessly blend deep representation learning with rigorous causal identification strategies will represent the ultimate realization of true precision medicine.

References

Yin, L., Tong, Y., Islam, Z. H., Zhang, K., Xie, R., Burger, J.,... & Wang, B. (2024). Targeted NAD+ Delivery for intimal Hyperplasia and Re-endothelialization: a novel anti-restenotic therapy approach. bioRxiv, 2024-02.

Chai L, Feng B, Liu X, et al. Fine mapping of a locus underlying the ectopic blade-like outgrowths on leaf and screening its candidate genes in rapeseed (Brassica napus L.)[J]. Frontiers in Plant Science, 2021, 11: 616844.

Wang, G., Zhong, Z., Zhou, Y., Li, Y., Wu, E., Cheah, W. P., ... & Shen, L. (2026). X-PCR: A Benchmark for Cross-modality Progressive Clinical Reasoning in Ophthalmic Diagnosis.arXiv preprint arXiv:2604.20350.

Li, Y., Liao, X., Yu, X., Xiao, Y., Tao, X., & Zhong, T. (2025). Mediating Role of the ANGPTL3/TFPI Protein Ratio in Regulating T-Cell Surface Glycoprotein CD5 Levels on Knee Osteoarthritis (KOA): A Mendelian Randomization Study. International Journal of Molecular Sciences, 26(10), 4471.

Liu, Z. Q., Xiong, Y. J., Yu, X., Lv, T., & Xu, H. Y. (2026). Female reproductive factors, genetic susceptibility and incident osteoporosis risk. Climacteric, 1-8.

Tao, Xiaoyu, et al. "A structural roadmap to gut health: A comprehensive review of edible insect chitin–protein complexes as alternative protein-derived, structure-specific prebiotics." International Journal of Biological Macromolecules (2026): 151312.

Yang Y, Li G, Zhong Y et al. scTenifoldXct: a semi-supervised method for predicting cell-cell interactions and mapping cellular communication graphs. Cell Syst 2023a;14:302–11.e4.

Wang, G., Ren, J., Shen, L., Cheah, W. P., & Qu, R. (2026). M-MambaS: Multimodal Mamba for small lesion segmentation.Pattern Recognition, 113923.

Stuart, T., Butler, A., Hoffman, P., Hafemeister, C., Papalexi, E., Mauck, W. M., et al. (2019). Comprehensive integration of single-cell data. Cell, 177(7), 1888–1902.

Wang, G., Li, Y., Chen, W., Ding, M., Cheah, W. P., Qu, R., ... & Shen, L. (2025, April). S³-mamba: Small-size-sensitive mamba for lesion segmentation. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 39, No. 7, pp. 7655-7664).

Rajpurkar, P., Irvin, J., Zhu, K., Yang, B., Mehta, H., Duan, T., et al. (2017). CheXNet: Radiologist-level pneumonia detection on chest X-rays with deep learning. arXiv preprint arXiv:1711.05225.

Wang, W., Wei, X., Sun, Z., Zhang, X., Pan, D., Hong, X., ... & Yang, D. (2026). Uncovering the toxicological impact of benzo [a] pyrene on Alzheimer's disease via network toxicology, machine learning, and single-cell transcriptomics. Journal of Alzheimer's Disease, 109(3), 1220-1231.

Wang, W., Chen, Y., Xiong, Z., Sun, Z., Wang, Z., Ma, L., ... & Li, X. (2026). Network toxicology and multidimensional bioinformatics analysis reveal the shared mechanism of action of bisphenol A and phthalates in glioblastoma. Ecotoxicology and Environmental Safety, 309, 119650.

Gao, T., Liang, D., Chen, X., & Tao, X. (2025). Predictive Analytics In Pediatric Mental Health: A Machine Learning Approachto Depression Detection By Physical Activity And Behavioral Health Indicators. Interciencia.

Zhao, K., Li, G., Li, C., Jiang, T., Wang, W., & Yang, Z. (2026). Identification Markers for Salvia miltiorrhiza and Its Close Relatives Based on Cell Wall Component Characteristics. Engineered Science, 40, 2122.

Chai L, Li H, Zhao X, et al. Analysis of altered flowering related genes in a multi-silique rapeseed (Brassica napus L.) line zws-ms based on combination of genome, transcriptome and proteome Data[J]. Plants, 2023, 12(13): 2429.

Xu, Z., Lu, Y., Tao, X., Cheng, W., Tian, J., Tang, W., & Wu, D. Modelling the geographic spread of HIV among MSM in Guangdong, China: a metapopulation model considering the impact of pre-exposure prophylaxis Philosophical Transactions A.

Zhang, W., Luo, Z., Gu, C., Ma, J., Cao, Y., Yuan, W., & Jin, Y. (2025). Bridging the compression-precision paradox: A hybrid architecture for clinical EEG report generation with guaranteed measurement accuracy. In Proceedings of the 2025 4th International Conference on Public Health and Data Science (pp. 201–207).

Xu, X., Bowtell, J., Young, W. R., Fong, D. T., & Williams, G. K. (2026). Increased common corticospinal input during eyes-closed unilateral stance in people with chronic ankle instability. Scientific Reports.

Zhao, K., Wu, X., Han, G., Sun, L., Zheng, C., Hou, H.,... & Shi, Z. (2024). Phyllostachys nigra (Lodd. ex Lindl.) derived polysaccharide with enhanced glycolipid metabolism regulation and mice gut microbiome. International Journal of Biological Macromolecules, 257, 128588.

Zhang, Y., & Yu, X. (2025). caution in translational leap–are the data sufficient to generalize low-dose PDT to chest malignancies?. Journal for immunotherapy of cancer, 13 (7), e013151.

Shi, M., & Chow, S. C. (2024). Statistical Analysis for Rating Scale in Clinical Trials.Therapeutic Innovation & Regulatory Science,58(5), 930-945.

Liu, S., Dai, J., Guo, Q., Cao, S., Wu, C., Zhu, C.,... & Luo, Y. (2026). Machine Learning–Based Prediction of Early Visual Field Nonrecovery After Endoscopic Surgery for Nonfunctioning Pituitary Adenomas With Clinically Accessible Biomarkers: A Dual-Center Study. International Journal of Clinical Practice, 2026 (1), 2504255.

Xu, X., Bowtell, J., Young, W. R., Fong, D. T., & Williams, G. K. (2025). Distinct patterns in neuromuscular adaptation to repeated perturbations in chronic ankle instability. Journal of NeuroEngineering and Rehabilitation.

Xiao, Y., Xiao-yue, Z., Yue, W., Ruo-tao, L., Xiang-jie, L., Xing-yuan, W., ... & Zhen-yi, J. (2022). Use of computed tomography for the diagnosis of surgical sarcopenia: Review of recent research advances. Nutrition in Clinical Practice, 37(3), 583-593.

Zhang, Y., & Chen, Y. (2025, July). The role of machine learning in reducing healthcare costs: The impact of medication adherence and preventive care on hospitalization expenses. In Proceedings of the International Conference on Implementing Generative AI into Telecommunication and Digital Innovation 2025 (pp. 179-185).

Wang, G., Zhou, Y., Deng, K., Cheah, W. P., Qu, R., Ren, J., & Shen, L. (2026). SurgCoT: Advancing Spatiotemporal Reasoning in Surgical Videos through a Chain-of-Thought Benchmark. arXiv preprint arXiv:2604.20319.

Kelley, D. R., Snoek, J., & Rinn, J. L. (2016). Basset: Learning the regulatory code of the accessible genome with deep convolutional neural networks. Genome Research, 26(7), 990–999.

Kabra, M., Shahi, P. K., Wang, Y., Sinha, D., Spillane, A., Newby, G. A., ... & Pattnaik, B. R. (2022). Nonviral base editing of KCNJ13 mutation preserves vision in an inherited retinal channelopathy.bioRxiv, 2022-07.

Verhoven, B., Tong, Y., Chlebeck, P., Zhong, W., Zeng, W., Jennings, H., ... & Al-Adra, D. P. (2025). Attenuating Ischemia and Reperfusion Injury Using NAD+-Loaded Nanoparticles in Mouse Kidneys.Transplantation Direct,12(1), e1890.