Multimodal Fusion for EHR Data: A Term Paper Overview

Zhaokun Wang

March 31, 2025

Abstract

Integrating diverse multimodal data (structured, unstructured, temporal) from Electronic Health Records (EHRs) is challenging due to heterogeneity and temporal irregularities, yet crucial for enhancing clinical predictions. This paper reviews and compares recent multimodal fusion techniques for temporal EHR data. The paper analyzed methods employing architectures like LSTMs with attention, intermediate fusion modules (using Bio-BERT and LSTMs), and multimodal Transformers for tasks including mortality and sepsis prediction. The comparison covers fusion strategies, data handling, and architectures, highlighting improved predictive accuracy from multimodal integration over unimodal approaches. Key challenges involve data sparsity, optimal fusion design, and interpretability. Future directions include self-supervised learning, advanced large language models, and developing robust, explainable fusion frameworks.

1 Introduction: The Significance of Multimodal Fusion in Temporal EHR Analysis

Electronic Health Records (EHRs) are now a key part of healthcare today, collecting a ton of patient information. These records hold different kinds of data, like structured data such as lab results and vital signs, unstructured data like clinical notes and radiology reports, and time-based details that show how a patient's health changes over time(Ma et al., 2024). EHRs get tricky because these data types are so different, they're collected at uneven times, and often some info is missing(Niu et al., 2023). This mix creates both a great chance and a tough problem. The chance comes from pulling together the best parts of each data type to get a fuller view of a patient's health. But the problem is figuring out how to sblend these varied sources into one solid model without technical headaches.

A patient's health doesn't stay still—it's always shifting, and the time patterns hidden in EHR data play a big role in making good diagnoses and treatment plans(Ma et al., 2024). Spotting long-term connections and small but important changes in a patient's condition over time is quite important for creating models that predict what's next(Ma et al., 2024). If we only look at patient data as a still snapshot, we will miss the time angle, which might hide key signs of how a disease is growing. Studying data with time in mind really matters in healthcare area.

Mixing info from multiple data types can seriously boost the accuracy of clinical prediction models compared to just using one source(Ma et al., 2024). Each type of data in EHRs often adds extra pieces to the puzzle, giving a richer and more complete picture of a patient's health(Wang et al., 2022). For example, time series data gives hard numbers on things like heart rate, while clinical notes bring in deeper meaning and word-based insights about the patient's state. Combining these well can tackle the weak spots each type has on its own, helping doctors make smarter choices. This term paper will wrap up and look at the methods and results from several research papers discussed in the seminar, all digging into how to fuse multimodal temporal EHR data.

2 Foundational Concepts in Multimodal EHR Data Integration

Electronic Health Records (EHRs) hold a big mix of patient info, grabbed through different data types. Structured Data makes up a hefty chunk, usually shown in tables. This covers basics like age and gender, lab test numbers for things like blood markers, vital signs such as heart rate or blood pressure, diagnosis codes (like ICD codes), and meds info (like drug codes) (Ma et al., 2024). These bits get logged at set times and give a number-based peek into a patient's body stats. Unstructured Data, though, is mostly the story-like clinical notes written by a doctor. Think doctor updates, discharge write-ups, and radiology reports—these pack in deep details about symptoms, past health, and why certain medical calls were made (Ma et al., 2024). Then there's Temporal Data, super key for tracking how health stuff changes over time. It's got time series of body measurements, event logs tagging medical moments with timestamps, and the time-ordered stack of clinical notes from a patient's visits (Ma et al., 2024). Besides these main types, EHRs also have things like medical images, a more detailed EHR data is list on Table 1. With all these data flavors being so different, we need special tricks to handle and mix them up right—each type's got its own vibe and needs its own way of being coded and sorted out.

Integrating multimodal temporal data from Electronic Health Records (EHRs) presents significant hurdles. Firstly, Data Heterogeneity is a major obstacle, as different data types vary greatly in format, scale, and meaning, such as numerical lab values versus free-text clinical notes. Secondly, Temporal Irregularity and Sparsity complicate integration, because data like vital signs and clinical notes are often recorded at inconsistent intervals and frequently contain missing values. These and other challenges associated with multimodal EHR data analysis will be discussed in greater detail in the subsequent sections.

Table 1: Simplified Modalities for Temporal Fusion

Modality Type	Examples	Characteristics	Challenges in Temporal Fusion
Structured Data	ICD codes, drug codes, lab values, CPT codes	Discrete, organized, quantifiable	Irregular intervals, unit variations, coding inconsistencies
Unstructured Notes	Progress notes, discharge summaries	Free-text, contextual, subjective	Temporal extraction, language variability, alignment with structured data
Medical Imaging	X-rays, MRIs, CT scans	Visual, high-dimensional	Dimensionality, processing needs, temporal alignment
Genomic Information	Mutations, SNPs	Complex, predisposition-related	Longitudinal integration, temporal marker-disease links
Physiological Signals	ECG, EEG, blood pressure	Continuous, time-series	Noise, event alignment, pattern extraction
Patient-Reported Outcomes	Symptom surveys, QoL scores	Subjective, patient-driven	Reporting variability, alignment with clinical data

3 Methodology Summary and Comparison of Core Papers

This section provides a detailed summary and comparison of three primary research papers from seminar sessions that address the challenges of analysing EHR data.

3.1 Research on Multimodal Fusion of Temporal Electronic Medical Records

The study titled "Research on Multimodal Fusion of Temporal Electronic Medical Records"(Ma et al., 2024) investigates a new approach to fuse temporal and non-temporal clinical notes along with tabular data to enhance prediction tasks using EHRs. The researchers identified that while deep learning has significantly impacted EHR research, the integration of diverse modalities within time series data still remain relatively underexplored. They proposed Time Series Multimodal Adaptation Gate (T-MAG) model to address this gap.

The T-MAG model processes four modalities of EHR data: static notes, time series notes, static tabular data, and time series tabular data. For temporal data, a preprocessing phase involving padding to a 30-day interval and segmenting into 3-day sub-sequences is applied. These sub-sequences are then fed into a Long Short-Term Memory (LSTM) network to generate sub-sequence representations.

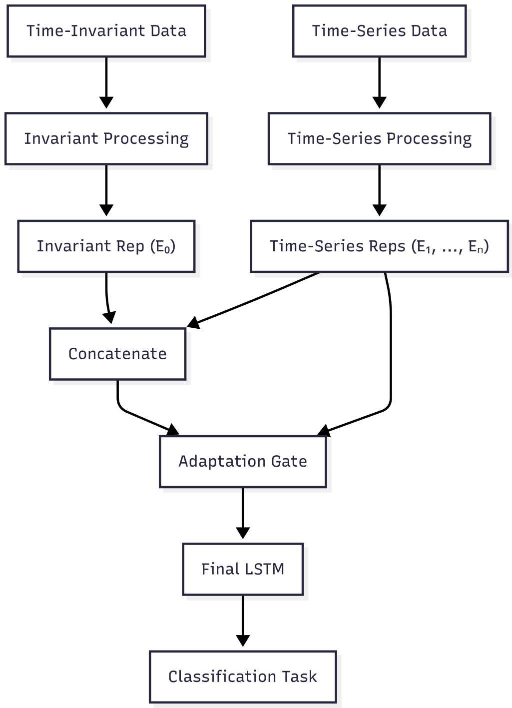

Figure 1: T-MAG's Architecture.

The model employs Multimodal Attention Gates (MAG) to fuse both static and temporal subsequence representations. Notably, an attention-backtracking module is introduced specifically for temporal fused representations to capture long-term dependencies. The concatenated results from the static and temporal branches are further processed by another LSTM to yield the final fused representation.

The model was evaluated on a dataset comprising 1271 myocardial infarction and 6450 stroke inpatients from a Beijing tertiary hospital. The study compared the predictive performance of T-MAG against several baseline models, including Crossformer. The key findings demonstrated that the proposed T-MAG model consistently achieved superior predictive accuracy in both in-hospital mortality and longer hospital stay prediction tasks for both myocardial infarction and stroke patients. Furthermore, the ablation study revealed that removing the attention-backtracking module led to a significant decline in performance, showing the importance of temporal data. The authors concluded that their method effectively integrates data from all four modalities and exhibits a good understanding of how to handle irregular time series data and lengthy clinical texts.

3.2 Deep Multi-Modal Intermediate Fusion of Clinical Record and Time Series Data in Mortality Prediction

Another contribution is the research on "Deep Multi-Modal Intermediate Fusion of Clinical Record and Time Series Data in Mortality Prediction"(Niu et al., 2023). This study addresses the challenge

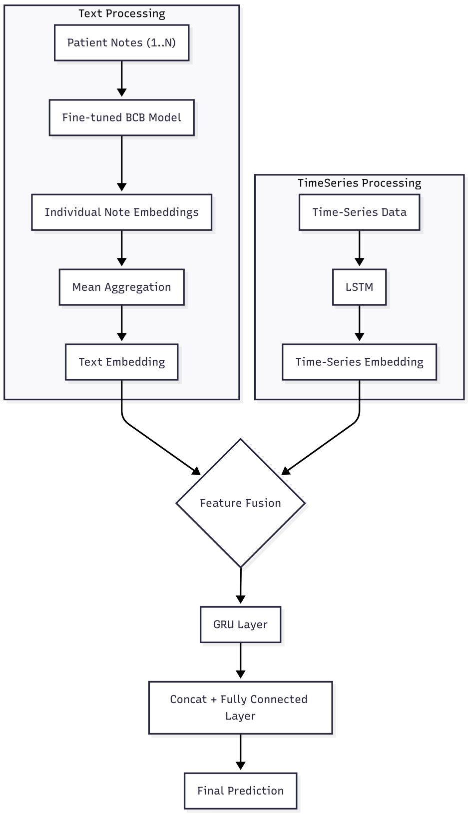

Figure 2: The overall architecture of the model.

of predicting mortality in Intensive Care Units (ICUs) by effectively combining information from two primary sources: time series data from continuous monitoring and clinical records containing physician diagnostic summaries. The authors observed that the majority of the existing mortality prediction studies primarily cascades multimodal features in the late stage, which can overlook valuable cross-modal correlations between the underlying features of different data types.

To overcome this limitation, the researchers proposed a novel multimodal fusion neural network model that incorporates an “intermediate fusion” module. The model first processes each modality separately. For clinical records, a fine-tuned Bio-BERT model is utilized to generate holistic embeddings of the textual data. For time series data, collected within the first 48 hours of ICU admission and pre-processed to handle missing values, a Long Short-Term Memory (LSTM) network is employed to capture temporal dependencies and generate time series embeddings.

The core innovation lies in the intermediate fusion module. The feature vectors from both modalities are transformed into the same dimensionality and then divided into equal channel feature blocks. The fusion module concatenates the corresponding channel blocks from each modality to learn global contextual information within each modality and, more importantly, to capture the correlations between them. A global representation for each modality is created by summing the feature blocks and applying global average pooling. These unimodal global representations are then combined element-wise to create a multimodal global representation, which is passed through a ReLU activation function to further capture dependencies. Block-level attention weights are generated and used to weight the feature blocks of each modality, enhancing the correlation between the patient's health condition changes and clinical data. Finally, the optimized feature blocks are concatenated to produce a fused multimodal feature matrix. This fused representation is then fed into a Gated Recurrent Unit (GRU) layer to further model dependencies, and its output is concatenated with low-level time series features before being passed through a fully connected layer with a sigmoid activation function for the final mortality risk prediction.

The model was trained and evaluated on the publicly available MIMIC-III dataset, which contains data from 18904 ICU patients. The results demonstrated that the proposed intermediate fusion model achieved superior performance in mortality prediction compared to various baseline methods, highlighting the effectiveness of jointly modeling time series data and clinical records and the importance of capturing cross-modal correlations at an intermediate stage. This suggests that understanding the intricate relationships between different data modalities beyond simple concatenation can significantly improve prediction accuracy.

3.3 Integrating Physiological Time Series and Clinical Notes with Transformer for Early Prediction of Sepsis

The paper "Integrating Physiological Time Series and Clinical Notes with Transformer for Early Prediction of Sepsis"(Wang et al., 2022) presents a multimodal Transformer model designed for the early prediction of sepsis in ICU patients. Recognizing sepsis as a leading cause of death in ICUs, the authors emphasize the critical need for early detection to improve patient survival. Their model integrates physiological time series data and clinical notes collected within the first 36 hours of a patient's ICU admission to predict the onset of sepsis.

The proposed framework comprises two main components: a Physiological Time Series Model (PTSM) and a Clinical Notes Model (CNM). The CNM utilizes ClinicalBERT, a pre-trained language model for clinical text, to generate contextualized embeddings of clinical notes. The output from the ClinicalBERT model, specifically the token representation, is then fed into a feedforward neural network (FNN). The PTSM, inspired by the standard Transformer architecture, processes physiological time series data. It begins with sequence embeddings and positional encoding to capture the temporal order of the measurements. This is followed by a stack of Transformer encoder layers, each consisting of a multi-head self-attention sublayer and a position-wise FNN sublayer with residual connections. Dense interpolation is used to handle irregularly sampled time series data before the final output of the PTSM is generated through an FNN.

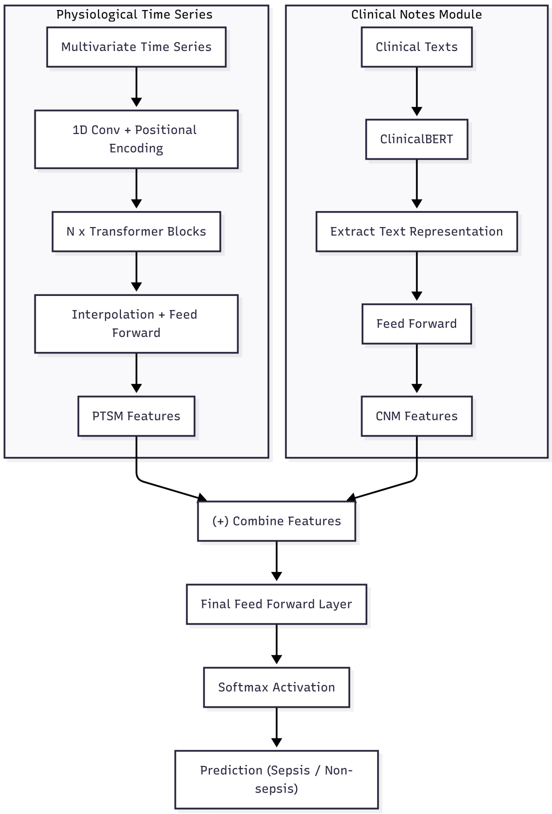

Figure 3: An overview of the multimodal Transformer framework.

The output representations from the PTSM and the CNM are then concatenated and fed into another FNN, followed by a Softmax layer for the binary classification of sepsis. The model was evaluated on two large critical care datasets: MIMIC-III and eICU-CRD, aiming to predict sepsis using data from different time windows within the first 36 hours (12, 18, 24, 30, and 36 hours). The proposed method was compared against six baseline models that combined LSTMs, BiLSTMs, or GRUs for time series data with Word2Vec, FastText, or ELMo for text representations. The experimental results demonstrated that the multimodal Transformer model consistently outperformed the competitive baselines across all evaluation metrics, including AUROC, F1 score, recall, and precision. Ablation analysis further confirmed the importance of incorporating both physiological time series data and clinical notes for achieving optimal performance. Case studies utilizing attention visualization on clinical notes and density plots of physiological features illustrated the unique and complementary information captured by each modality. This research highlights the effectiveness of Transformer-based architectures in integrating diverse EHR data types for a critical clinical prediction task.

4 Comparative Analysis

To facilitate a clearer understanding of the similarities and differences between the three primary papers, a comparative analysis is presented in Table 2.

Table 2: Comparison of Multimodal Fusion Approaches in EHR Research

Feature	Research on Multimodal Fusion of Temporal EHRs	Deep Multi-Modal Intermediate Fusion	Integrating Time Series & Notes for Sepsis
Model	T-MAG (LSTM with attention)	Bio-BERT + LSTM + Fusion Module + GRU	Multimodal Transformer (PTSM + CNM)
Fusion Technique	Multimodal Attention Gates, Attention-Backtracking	Intermediate Fusion Module	Feature Concatenation in Transformer
Data Modalities	Temporal/Static Notes, Temporal/Static Tabular	Clinical Records (Text), Time Series	Physiological Time Series, Clinical Notes
Temporal Data Handling	Padding to 30 days, segmentation into 3-day subsequences, LSTM	First 48 hours of ICU admission, LSTM	First 36 hours (various windows), dense interpolation
Prediction Task	Mortality & LOS (AMI, Stroke)	In-Hospital Mortality (ICU)	Early Sepsis Prediction
Key Results (AUROC)	AMI Mortality: 0.928, Stroke Mortality: 0.954	MIMIC-III: Improved over baselines	MIMIC-III: 0.96–0.99, eICU-CRD: High

The three papers discussed above present distinct yet complementary approaches to the multimodal fusion of EHR data for clinical prediction. While all aim to improve prediction accuracy by leveraging multiple data modalities, they differ in their fusion techniques, data preprocessing strategies, model architectures, and the specific clinical prediction tasks they address.

In terms of fusion techniques, the first paper, focusing on T-MAG, employs multimodal attention gates and an attention-backtracking module to integrate static and temporal data(Ma et al., 2024). This approach utilizes the concept of a primary modality within the attention gate framework, allowing the model to focus on the most informative data streams while incorporating auxiliary information(Ma et al., 2024). The second paper introduces an intermediate fusion module that operates on the feature representations generated by unimodal encoders (Bio-BERT for text and LSTM for time series)(Niu et al., 2023). This module uses channel-wise concatenation and a soft attention mechanism to capture cross-modal correlations at a deeper level of abstraction(Niu et al., 2023). The third paper, employing a multimodal Transformer, utilizes feature concatenation of the outputs from the Physiological Time Series Model (PTSM) and the Clinical Notes Model (CNM) before feeding them into a final feedforward neural network for prediction(Wang et al., 2022). This can be viewed as a form of late fusion at the feature level, enhanced by the Transformer's inherent capability to model dependencies within each individual modality(Wang et al., 2022). The progression in fusion methodologies, from simpler concatenation to more sophisticated attention-based mechanisms and intermediate fusion modules, signifies an evolving understanding of the intricate interactions between different EHR modalities and the growing need to model these interactions with greater efficacy.

The data preprocessing steps undertaken in each study also reflect the unique characteristics of the data modalities involved. The T-MAG paper preprocesses temporal data by padding and segmenting it into fixed-length subsequences(Ma et al., 2024). The intermediate fusion paper resamples time series data to hourly intervals and imputes missing values, while clinical notes are processed using Bio-BERT(Niu et al., 2023). The multimodal Transformer paper performs outlier removal, resampling, and imputation for physiological data, and applies cleaning and removal of sepsis-related terms to clinical notes(Wang et al., 2022). These preprocessing steps underscore the inherent challenges associated with each data type. Time series data often requires strategies for handling irregular sampling and missing values, whereas clinical notes necessitate text-specific preprocessing techniques such as cleaning and tokenization, along with careful consideration to avoid introducing bias or information leakage.

The model architectures employed in these papers also vary. The T-MAG model primarily relies on LSTM networks augmented with custom attention mechanisms(Ma et al., 2024). The intermediate fusion model combines LSTM for time series data, Bio-BERT for clinical text, and a GRU layer for processing the fused representation(Niu et al., 2023). The multimodal Transformer model utilizes the Transformer architecture for both time series (PTSM) and text (ClinicalBERT within CNM), with a final feedforward neural network for the classification task(Wang et al., 2022). The choice of deep learning architecture, whether RNN-based or Transformer-based, appears to be influenced by the specific task and the nature of the data being analyzed. Transformers are increasingly being adopted for their ability to handle long sequences and model complex relationships, but RNNs like LSTMs and GRUs continue to be relevant, particularly for modeling temporal dependencies in time series data.

The T-MAG model demonstrates its strength in effectively integrating four different modalities and handling the complexities of temporal data and long clinical texts, achieving superior predictive performance on myocardial infarction and stroke prediction tasks(Ma et al., 2024). The novel attention-backtracking module appears to be particularly effective in capturing long-term temporal dependencies(Ma et al., 2024). However, its reliance on fixed-length 3-day subsequences might limit its adaptability to varying temporal patterns in patient data, and the selection of a primary modality in the MAG framework could be challenging in the medical domain where the importance of different modalities can vary(Ma et al., 2024).

The intermediate fusion model excels in jointly modeling time series and clinical records for mortality prediction in ICU patients, effectively capturing cross-modal correlations through its specialized fusion module and achieving improved accuracy(Niu et al., 2023). A potential limitation is its specific focus on ICU mortality, which might restrict its generalizability to other clinical prediction tasks or different patient populations. Additionally, the preprocessing of time series data to a fixed hourly interval might still lead to some loss of information or introduce artificial regularity(Niu et al., 2023).

The multimodal Transformer model stands out as the first Transformer-based approach to integrate physiological time series and clinical notes for early sepsis prediction, demonstrating superior performance on two large datasets and providing enhanced interpretability through attention visualization(Wang et al., 2022). While interpretability is a significant strength, the inherent complexity of Transformer models can still make it difficult to fully understand the reasoning behind their predictions, and the model's sensitivity for sepsis prediction was observed to decrease with longer prediction windows before onset(Wang et al., 2022).

Collectively, these three papers highlight the effectiveness of multimodal fusion across different critical clinical prediction tasks, including mortality prediction for various patient populations and the early detection of a severe condition like sepsis. The diverse approaches and strong performance metrics reported underscore the broad applicability and potential of these techniques in advancing healthcare.

5 Challenges in Multimodal EHR Analysis

Looking at multimodal EHR data brings some big problems that researchers are still working on. One key issue is missing and uneven data(Niu et al., 2023). EHR data, especially measurements taken over time, often has gaps with missing numbers and different recording times across patients or even for the same patient(Niu et al., 2023). These differences make it tough to use normal time series tools. The paper called "Self-supervised transformer for sparse and irregularly sampled multivariate clinical time-series" (Knez and Žitnik, 2024) shows this is a big focus. Self-supervised learning can help by training models to guess missing parts, finding patterns even with holes in the data. This could make time-based EHR studies more solid and less affected by missing pieces.

Another big problem is how to mix different types of data well. Figuring out the best way to combine things like time series numbers and written doctor notes is still unclear. There are different mixing methods—like doing it early, late, or in the middle—and each has good and bad points. A review paper, "Research on Multimodal Fusion..."(Ma et al., 2024), says mixing time series and fixed data directly might hide important time order details. Also, in medicine, unlike feelings analysis where text often has the most emotion, it's hard to pick one main data type because each kind can give special and important clues (Ma et al., 2024). So, making smart mixing methods that fit the situation and handle the mixed EHR data types, plus catch tricky links between them, is very important for progress. Papers show many mixing ideas, like attention tools or middle-step mixing, which shows this work is still growing.

Also, EHR data has lots of details and differences. The huge amount of features and items in EHRs, along with different forms, setups, and scales of data types, makes analysis hard. Cutting down the number of details is often needed to handle this, and good fixing methods are key to make sure different data types can work together in models.

Lastly, understanding and explaining results matter a lot in healthcare (Wang et al., 2022). Deep learning models can predict well, but their hidden way of working can stop them from being used in real life. Doctors need to know why a model says something to trust it and use it in their choices. So, building models that show what affects their predictions—like using attention pictures or showing important features—is a key area researchers are still studying.

6 Future Research Directions

The world of multimodal fusion and temporal EHR analysis is packed with chances for new studies. Tapping into self-supervised learning tricks—like those hinted at in a paper about spotty and uneven time series (Tipirneni and Reddy, 2022)—looks like a bright path forward. It's all about teaching models to pull strong, flexible patterns from the huge piles of unlabeled EHR data. By getting models ready with tasks like guessing missing bits or what comes next in time, researchers might boost how well they tackle prediction jobs later on, especially when there's not much labeled data to work with.

Another neat idea to dig into is prompt learning. This method(Wang and Sun, 2022) uses plain language prompts to steer LLMs toward specific EHR tasks. Instead of the usual tweaking, prompt learning could make pre-trained models more bendy and quick to adapt—think creating virtual EHR data to keep things private or helping doctors decide by pulling up useful info based on their questions.

We also need to tackle the shaky ground under Large Language Models (LLMs) for EHRs, like the paper points out (Wornow et al., 2023). LLMs have tons of promise for sorting through and making sense of the rich text in EHRs, but we've got to watch out. Stuff like keeping data private, hidden biases in the models, needing solid medical knowledge, and the risk of spitting out wrong or confusing info all need serious study and fixes.

Keeping a strong focus on interpretability and explainability is a must if we want these fancy analysis tools to catch on in clinics. Future studies should search for fresh ways to make multimodal EHR models clearer and easier for doctors to get. This might mean trying out new ways to show the data, figuring out how to tie predictions to specific bits or data types, or even suplying clear explanations for what the model spits out.

Lastly, the hunt for new fusion setups that can really grab the twisty links between different EHR data types should keep rolling. This could mean diving into slick deep learning ideas like graph neural networks to map out the messy web of medical events and items, or memory networks that pick out the key stuff across long time stretches and mixed data.

7 Conclusion

In conclusion, combining different types of EHR data through multimodal fusion has great potential to change how we predict medical outcomes and help patients get better. Although researchers have made big steps forward, we still need more studies to tackle the current problems and unlock the full strength of mixing various data sources in healthcare. Exploring new ways to blend data, building stronger and clearer models, and staying serious about solving ethical issues will open the door for these tools to be widely used in clinics in a responsible way. Looking at the chosen research papers shows how much progress is happening in mixing electronic medical records for healthcare predictions. Studies have tried different methods to combine time series data and clinical notes, using deep learning tools like LSTMs and Transformers, and testing blending tricks such as attention mechanisms and intermediate feature fusion. These results keep showing that putting together multiple data types boosts the accuracy of predicting important things like death rates or sepsis risks.

Even with this progress, some tough issues stick around. EHR data is often uneven and missing parts, models need to be easier to understand, and there aren't enough shared rules for data terms or ways to measure success. Future work should focus on fixing these gaps by trying fresh ideas like self-supervised learning, creating fake data with generative models, and improving how we mix data. Plus, we need to set up standard ways to share data and check models so that research can move forward smoothly and turn into real help for clinics.

References

Knez, Timotej and Slavko Žitnik. 2024. "Multimodal learning for temporal relation extraction in clinical texts." *Journal of the American Medical Informatics Association* 31(6):1380–1387.

Ma, Moxuan, Muyu Wang, Binyu Gao, Yichen Li, Jun Huang and Hui Chen. 2024. "Research on Multimodal Fusion of Temporal Electronic Medical Records." *Bioengineering* 11(1):94.

Niu, Ke, Ke Zhang, Xueping Peng, Yijie Pan and Naian Xiao. 2023. "Deep multi-modal intermediate fusion of clinical record and time series data in mortality prediction." *Frontiers in Molecular Biosciences* 10:1136071.

Tipirneni, Sindhu and Chandan K Reddy. 2022. "Self-supervised transformer for sparse and irregularly sampled multivariate clinical time-series." *ACM Transactions on Knowledge Discovery from Data (TKDD)* 16(6):1–17.

Wang, Yuqing, Yun Zhao, Rachael Callcut and Linda Petzold. 2022. "Integrating physiological time series and clinical notes with transformer for early prediction of sepsis." *arXiv preprint arXiv:2203.14469* .

Wang, Zifeng and Jimeng Sun. 2022. PromptEHR: Conditional electronic healthcare records generation with prompt learning. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing. Conference on Empirical Methods in Natural Language Processing*. Vol. 2022 p. 2873.

Wornow, Michael, Yizhe Xu, Rahul Thapa, Birju Patel, Ethan Steinberg, Scott Fleming, Michael A Pfeffer, Jason Fries and Nigam H Shah. 2023. "The shaky foundations of large language models and foundation models for electronic health records." *npj digital medicine* 6(1):135.

Appendix: Use of AI-Based Tools

This appendix documents the use of artificial intelligence (AI)-based tools in the preparation of this academic work.

List of Steps Involving AI-Based Tools

- **DeepSeek**: I consulted DeepSeek models to learn more formal organization of an conclusion chapter. The suggested frameworks were adapted and rewritten entirely in my own words.
- **QuillBot**: QuillBot was used sparingly to rephrase sentences for improved readability and flow. All suggestions were manually reviewed and edited to ensure alignment with my original intent and academic style.
- **DeepL and Youdao Translation**: DeepL and Youdao Translation assisted in translating a small number of technical terms and short phrases from Chinese to English to clarify meaning during drafting. These translations were verified and incorporated into my own text.