

MEDFuse: Multimodal EHR Data Fusion with Masked Lab-Test Modeling and Large Language Models

Phan Nguyen Minh Thao*

Cong-Tinh Dao*

National Yang Ming Chiao Tung University
Hsinchu, Taiwan
{pnmthaoct,dcongtingh}@gmail.com

Chenwei Wu

University of Michigan
Michigan, USA
chenweiw@umich.edu

Jian-Zhe Wang

National Yang Ming Chiao Tung University
Hsinchu, Taiwan
jzawang.cs09@nycu.edu.tw

Shun Liu

Shanghai University of Finance and Economics
Shanghai, China
kevinliuleo@gmail.com

Jun-En Ding

Stevens Institute of Technology
New Jersey, USA
jding17@stevens.edu

David Restrepo

Massachusetts Institute of Technology
Massachusetts, USA
davidres@mit.edu

Feng Liu

Stevens Institute of Technology
New Jersey, USA
fliu22@stevens.edu

Fang-Ming Hung

Far Eastern Memorial Hospital
New Taipei, Taiwan
philip@mail.femh.org.tw

Wen-Chih Peng

National Yang Ming Chiao Tung University
Hsinchu, Taiwan
wcpengcs@nycu.edu.tw

Abstract

Electronic health records (EHRs) are multimodal by nature, consisting of structured tabular features like lab tests and unstructured clinical notes. In real-life clinical practice, doctors use complementary multimodal EHR data sources to get a clearer picture of patients' health and support clinical decision-making. However, most EHR predictive models do not reflect these procedures, as they either focus on a single modality or overlook the inter-modality interactions/redundancy. In this work, we propose **MEDFuse**, a Multimodal EHR Data Fusion framework that incorporates masked lab-test modeling and large language models (LLMs) to effectively integrate structured and unstructured medical data. MEDFuse leverages multimodal embeddings extracted from two sources: LLMs fine-tuned on free clinical text and masked tabular transformers trained on structured lab test results. We design a disentangled transformer module, optimized by a mutual information loss to 1) decouple modality-specific and modality-shared information and 2) extract useful joint representation from the noise and redundancy present in clinical notes. Through comprehensive validation on the public MIMIC-III dataset and the in-house FEMH dataset, MEDFuse demonstrates great potential in advancing clinical predictions, achieving over 90% F1 score in the 10-disease multi-label classification task.

*Both authors contributed equally to this research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CIKM '24, October 21–25, 2024, Boise, ID, USA.

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0436-9/24/10
<https://doi.org/10.1145/3627673.3679962>

CCS Concepts

• **Applied computing** → **Health care information systems**; • **Computing methodologies** → **Artificial intelligence**; • **Information systems** → **Data mining**.

Keywords

Computer-aided Diagnosis; Large Language Model Fine-tuning; Electronic Health Records

ACM Reference Format:

Phan Nguyen Minh Thao, Cong-Tinh Dao, Chenwei Wu, Jian-Zhe Wang, Shun Liu, Jun-En Ding, David Restrepo, Feng Liu, Fang-Ming Hung, and Wen-Chih Peng. 2024. MEDFuse: Multimodal EHR Data Fusion with Masked Lab-Test Modeling and Large Language Models. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management (CIKM '24)*, October 21–25, 2024, Boise, ID, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3627673.3679962>

1 Introduction

Electronic Health Records (EHRs) widely document heterogeneous patient data, including tabular records and unstructured clinical notes. Tabular records encompass medical concepts like diagnoses, medications, and lab tests, providing a structured overview of a patient's health. In contrast, clinical notes are free-text documents by healthcare providers, offering detailed and nuanced accounts of patient history, clinical findings, and progress. The vast volume and diversity of multimodal data in EHRs present a unique opportunity for deep learning to enhance disease prediction [11, 27]. Nevertheless, the heterogeneity and redundancy in multimodal EHR inputs present significant challenges for effectively distill and fuse clinically meaningful information for disease prediction.

The primary question at hand is: *can we effectively obtain and integrate useful representations for different EHR modalities to improve clinical predictions?* Current research in deep EHR modeling [24, 25] often focuses on single data modalities, neglecting significant insights from unstructured medical notes and lab tests, which

limits the model’s ability to learn a comprehensive view of patient health conditions. Lab tests consist of high-dimensional, usually discrete tabular data; however, the conventional approach models structured EHR data as numerical vectors, overlooking complex interactions between individual variables and does not consider their interactions [9, 20, 24, 25, 33]. Recent work has shifted towards deep learning architectures like BERT and LLMs. Fine-tuned LLM on clinical data shows promise in understanding clinical notes, answering medical questions, and making few-shot predictions [30]. However, LLMs struggle with numerical lab test nuances and underperform on tabular prediction tasks. [3, 13, 15].

Another significant challenge in fusing information from different types of EHR data is *How do we distill the overlapping clinically important features from both modalities?* The information contained within different modalities can be categorized as either modality-specific or modality-shared [21]. For example, a patient’s dietary habits would be considered information specific to the clinical notes modality; hypertension record and lab test value would be regarded as modality-shared information. Cai et al. [5] focused on integrating the modality-shared information by emphasizing the inherent consistency through alignment techniques. However, this approach often leads to the common information dominating the alignment and integration process, resulting in the distinctive perspectives offered by each modality being disregarded. Lab tests and clinical notes have significantly different noise-to-information ratios, making it hard to distill useful joint representation from the noises and redundancy present in EHR. Therefore, there is an urgent need for methods to extract the diverse yet collaborative perspectives both modalities offer for informing therapeutic decision-making.

In this work, we propose MEDFuse, a novel Multimodal EHR Data Fusion diagnostic model consisting of modality-specific embedding extractors followed by a disentangled transformer for multimodal fusion. Our model integrates embeddings between fine-tuned LLMs on unstructured clinical text and masked lab-test modeling models pre-trained on structured lab results. We further utilize a disentangled transformer optimized by mutual information loss to decouple modality-specific and modality-common information and learn meaningful joint representations for downstream prediction tasks. The key contributions of our work are as follows:

- We proposed a novel diagnostic model integrating structured lab test data and unstructured clinical notes, utilizing embeddings from fine-tuned LLMs and Masked Lab-Test Modeling, enhancing understanding of diverse clinical information.
- We improved joint patient representation by incorporating a disentangled transformer module to effectively separate and integrate modality-specific and shared information, leading to better prediction outcomes across multiple diseases.
- We conducted empirical evaluations to illustrate our model’s effectiveness through EHR datasets on various metrics.

2 Related work

2.1 EHR For Multi-Label Disease Prediction

Most recent works in medical Multi-Label Text Classification (MLTC) entirely rely on medical texts. For instance, Kim et al. [7] introduced a convolutional attention network designed to extract meaningful

document representations across varying text lengths. Recent developments in LLMs, such as those discussed by Luo et al. (2022) and Elliot et al. [4], utilize extensive data from medical literature for domain-specific tasks such as natural language inferencing. Additionally, some studies have employed graph neural networks (GNNs) to organize sequences from Electronic Medical Records (EMR) into hierarchical graphs [32], or to integrate entity relationships from text using attention mechanisms in neural networks [6, 12]. Nevertheless, many of these studies overlook the potential advantages of integrating medical expert knowledge from official guidelines and critical blood tests. A combined approach that harnesses both unstructured and structured data could offer extra help to offset issues like label and data scarcity in the medical domain.

2.2 Extraction of Clinical Relevant Information from Multimodal EHR

Recent work has leveraged self-supervised learning methods, like contrastive pretraining of clinical notes [5] and prompt-based large language modeling [11, 15], to facilitate multimodal learning of EHR data. The former encourages the alignment between paired patient data via contrastive loss, and the latter directly converts the structured data into text by prompt templates and feeds it into LLMs. However, if the data fusion process focuses solely on aligning the common information, such as diabetes history (text) and blood glucose levels (lab), the rich, modality-specific insights like exercise habits may be overlooked. This limits understanding of the patient’s health and impact predictive models. Therefore, developing multimodal EHR data fusion techniques to integrate both modality-specific and modality-shared information is essential.

3 Method

3.1 Overview

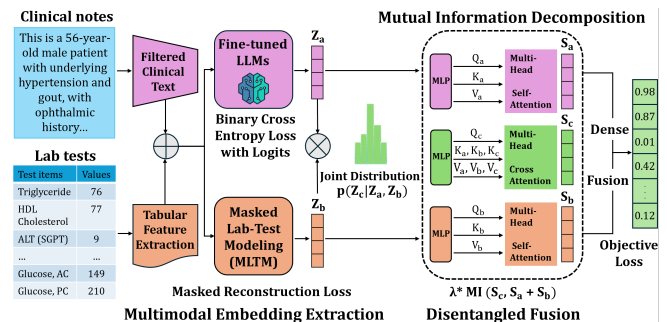


Figure 1: The proposed model architecture.

In Figure 1, MEDFuse integrates clinical notes and lab tests to obtain patient representation for accurate multi-disease prediction. In the Multimodal Embedding Extraction stage, textual data from clinical notes, are filtered and structured. Simultaneously, abnormal lab tests like triglyceride levels, HDL cholesterol, are extracted and formatted into textual data by prompt templates. The filtered clinical text is then processed by fine-tuned LLMs to generate embeddings capturing its semantic meaning. Meanwhile, raw structured tabular data are processed with a domain-specific masked lab-test model to create embeddings for the quantitative lab data. These embeddings are then merged through the disentangled transformer module for multimodal fusion and final disease prediction.

3.2 Multimodal Embedding Extraction

3.2.1 Fine-tuning LLMs on Unstructured Text. Clinical notes, comprising diverse fields derived from physicians’ diagnoses, form the textual component of the dataset. We filtered the text by Chief Complaint, Present Illness, Medical History, and Medication on Admission. These specific fields are crucial for accurately predicting a patient’s disease. To integrate this tabular data with the textual clinical notes, we converted the tabular data into a textual format, a process referred to as tabular feature extraction. This method involves extracting abnormal lab test results and formatting them into a text template – “These are abnormal results recorded: ITEMID <ITEMID>: <VALUE> <VALUEUOM>; ITEMID <ITEMID>: <VALUE> <VALUEUOM>; ...”. Here, <ITEMID> refers to the specific lab test names, <VALUE> indicates the test values, and <VALUEUOM> denotes the units of measure for the test values.

Inspired by fine-tuning the language models [10, 23] for classification purposes, we fine-tuned various LLMs for disease prediction. Our best-performing backbone is the publicly accessible Medical-Llama3-8B model [28], which is fine-tuned from Meta-Llama-3-8B [2]. It is trained on a comprehensive medical chatbot dataset and optimized for addressing health-related inquiries. We extracted latent vector representations from the final layer of the Llama decoder, which was originally engineered for autoregressive prediction of subsequent tokens. These extracted vectors were subsequently processed through feed-forward neural layers, effectively transforming them into a label space. The output from these transformations, in the form of logits, was then utilized to perform discriminative classification based on labels. This method aims to harness the latent embedding of LLMs to achieve targeted, efficient task adaptation.

3.2.2 Masked Lab-Test Modeling. The Masked Lab-Test Modeling (MLTM) module extends the Masked Autoencoders (MAE) [8, 14, 29] framework to reconstruct masked components based on observed components in EHR data. MLTM consists of an encoder that maps observed values to their representations and a decoder that reconstructs the masked values from the latent representations. To account for the inherent incompleteness in the imputation task, MLTM employs an additional masking approach in the training to make sure a uniform 75% value is masked out. The encoder applies a learnable linear encoding function $wx + b$ to each unmasked x and passes through a transformer architecture, while the decoder operates on the embeddings of both observed and masked values. Positional encoding is added to the embeddings to preserve the lab test positions. The reconstruction loss is defined as the mean square error between the reconstructed and original values on the re-masked and unmasked sets. MLTM is designed with an asymmetric architecture, using a deep encoder and a shallow decoder to extract useful lab-test representations.

3.3 Disentangled Transformer Module

Initially, features from each modality are multiplied by the Kronecker product to approximate a joint distribution, $C = A \otimes B \in \mathbb{R}^{(m \times (a) \times (b))}$, effectively capturing the pairwise interactions. Self-attention is applied to Z_a and Z_b to obtain S_a and S_b , controlling the expressivity of each modality and preventing noisy features. Subsequently, the common information of the joint distribution is extracted via cross attention of $Q_c, K_c + K_a + K_b$, and $V_c + V_a + V_b$ to

model modality-common features S_c . To preserve modality-specific information, we minimize the Mutual Information (MI) loss between concatenated $S_a + S_b$ and S_c . As the computation of mutual information is intractable, we calculate a variational upper bound called contrastive log-ratio upper bound (vCLUB) as an MI estimator to achieve MI minimization. Given two variables a and b , the $L_v^{CLUB}(a, b)$ is calculated as follows [34]:

$$\begin{aligned} L_v^{CLUB}(a, b) &= \mathbb{E}_p(a, b) [\log q_\theta(b|a)] - \mathbb{E}_p(a) \mathbb{E}_p(b) [\log q_\theta(b|a)] \\ &= \frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N [\log q_\theta(b_i|a_i) - \log q_\theta(b_j|a_i)] \end{aligned} \quad (1)$$

We employ an MLP $q_\theta(b|a)$ to provide a variational approximation of $q_\theta(b|a)$, which can be optimized by maximizing the log-likelihood [34]: $L_{\text{estimator}}(a, b) = \frac{1}{N} \sum_{i=1}^N \log q_\theta(b_i|a_i)$. The MI Loss is then defined as: $\text{MI Loss} = L_v^{CLUB}(S_a + S_b) + L_{\text{estimator}}(S_a + S_b, S_c)$.

After optimizing the mutual information between the modality-specific information and the modality-common information, we utilize dense fusion [16] to enable denser interaction between modalities. Instead of directly connecting a prediction classifier on top of the fused representation S_c , we learn deeper representations of the clinical notes and lab test features and add skip connections to concatenate with the fused representation, forming a final fused embedding: $h_a = f_a(S_a)$ and $h_b = f_b(S_b)$ where f_a and f_b are fully-connected layers. This final representation not only aggregates the modality-specific features but also incorporates the modality-common representation from the previous stage of the network: $h_{\text{final}} = \text{concat}(h_a, S_c, h_b)$. Finally, a dense block g is used to generate $y = g(h_{\text{final}})$, and the model is trained by optimizing the prediction loss (focal loss for multilabel prediction). This allows for dense interaction of features from each modality, aggregating information across different stages of the network. The final loss optimizes a combination of the prediction objective and the mutual information loss, controlled by a hyperparameter λ with a value range of $[0, 1]$. In our experiment, we choose a λ value of 0.1. $\text{Loss}_{\text{final}} = L_{\text{objective}}(g(h_{\text{final}})) + \lambda * \text{MI}(\text{concat}(S_a, S_b), S_c)$.

4 EXPERIMENTS

4.1 Datasets and Metrics

We used two real-world EHR datasets: MIMIC-III [19] and FEMH, which collected five years of EHRs from the Far Eastern Memorial Hospital (FEMH) in Taiwan between 2017 and 2021. FEMH includes 1,420,596 clinical notes, 387,392 lab results, and over 1,505 lab test items. The FEMH Research Ethics Review Committee¹ approved the study, and all data were de-identified. We selected patients with at least two recorded visits from each dataset.

For the multi-label classification task in MIMIC-III, we identified the top 10 most prevalent conditions: “Hypertension, uncomplicated”, “Cardiac arrhythmias”, “Fluid and electrolyte disorders”, “Congestive heart failure”, “Diabetes w/o chronic complications”, “Chronic pulmonary disease”, “Valvular disease”, “Renal failure”, “Hypertension, complicated”, and “Other neurological disorders”. In the FEMH dataset, the top 10 most common diseases include “Hypertension”, “Diabetes”, “Heart disease”, “Cancer”, “Cerebrovascular

¹<https://www.femhirb.org/>

Table 1: Training and Validation Performance Comparison of Various Models on the MIMIC-III Dataset.

Model	Precision	Recall	F1 macro	F1 weighted	Accuracy
Bert	0.8333 / 0.6790	0.2000 / 0.2000	0.1818 / 0.1618	0.3686 / 0.3162	0.6515 / 0.2692
LoRA Mistral-7B-v0.1	0.8759 / 0.8616	0.8459 / 0.8289	0.8449 / 0.8274	0.9007 / 0.8886	0.9089 / 0.8974
LoRA Llama-2-7B-hf	0.8828 / 0.8585	0.8592 / 0.8364	0.8559 / 0.8301	0.9097 / 0.8924	0.9168 / 0.9004
LoRA Meta-Llama2-13B	0.9153 / 0.8732	0.8852 / 0.8430	0.8874 / 0.8414	0.9297 / 0.8990	0.9363 / 0.9071
LoRA Meta-Llama3-8B	0.8899 / 0.8667	0.8569 / 0.8306	0.8579 / 0.8305	0.9121 / 0.8935	0.9211 / 0.9040
LoRA Medical-Llama3-8B	0.9283 / 0.8807	0.9008 / 0.8474	0.9026 / 0.8466	0.9367 / 0.9003	0.9417 / 0.9068
MEDFuse	0.9375 / 0.9025	0.9217 / 0.8534	0.9216 / 0.8615	0.9462 / 0.9103	0.9535 / 0.9122

Table 2: Training and Validation Performance on the FEMH Dataset.

Model	Precision	Recall	F1 macro	F1 weighted	Accuracy
LoRA Medical-Llama3-8B	0.8702 / 0.8691	0.8496 / 0.8478	0.8453 / 0.8435	0.9182 / 0.9167	0.9267 / 0.9252
MEDFuse	0.8839 / 0.8823	0.8707 / 0.8670	0.8637 / 0.8607	0.9260 / 0.9243	0.9311 / 0.9296

Table 3: Ablation Study on Training and Validation Performance on the MIMIC-III Dataset.

Model	Precision	Recall	F1 macro	F1 weighted	Accuracy
MEDFuse w/o (MLTM & LABTEXT)	0.8882 / 0.8580	0.8663 / 0.8406	0.8620 / 0.8321	0.9100 / 0.8901	0.9148 / 0.8955
MEDFuse w/o (MLTM & TEXT)	0.6553 / 0.6224	0.6461 / 0.6203	0.6282 / 0.5980	0.7869 / 0.7627	0.8239 / 0.8008
MEDFuse w/o TEXT	0.7730 / 0.7666	0.7912 / 0.7923	0.7600 / 0.7573	0.8331 / 0.8230	0.8331 / 0.8271
MEDFuse w/o Disentangled Transformer	0.9330 / 0.8974	0.9164 / 0.8483	0.9162 / 0.8564	0.9417 / 0.9074	0.9489 / 0.9082
MEDFuse	0.9375 / 0.9025	0.9217 / 0.8534	0.9216 / 0.8615	0.9462 / 0.9103	0.9535 / 0.9122

Disease”, “Kidney Disease”, “Liver Disease”, “Asthma”, “Hyperlipidemia”, and “Lung Disease”. We applied several metrics to assess model performance such as Macro-average and weighted-average F1-Scores, precision, recall, and accuracy [17, 22, 26].

4.2 Experimental Results

Table 1 and Table 2 illustrate the training and validation performance of various models, highlighting the effectiveness of our proposed method on the MIMIC-III and FEMH datasets, respectively. In Table 1, our approach outperforms baseline models such as Bert [10], Mistral-7B-v0.1 [18], Llama-2-7B-hf [31], Meta-Llama2-13B [1], Meta-Llama3-8B [2], and Medical-Llama3-8B [28] across all key metrics. Cai et al. [5] pretrained multimodal models, consuming significant time and resources, while our approach uses finetuning and test-time embedding alignment, which is much more efficient. Specifically, MEDFuse shows significant improvements over the best-performing LoRA fine-tuned LLM, Medical-Llama3-8B. On the test set, our model performs 1.49% better in macro F1 score, and similar trends are observed in other metrics. Table 2 shows MEDFuse consistently outperforms Medical-Llama3-8B on the FEMH dataset. For example, training and validation in precision is a 1.53% increase, in the recall is a 2.07% increase, the training accuracy is 0.9311 (0.55% increase), and validation accuracy is 0.9296 (0.41% increase). These results validate the robustness and generalizability of our approach, underscoring its potential for accurate and reliable clinical predictions across diverse datasets.

4.3 Ablation Study

We conducted an ablation study to examine the contributions of various components in our proposed method, which integrates Medical-Llama3 with a transformer module, utilizing lab tests (LABTEXT) and clinical notes (TEXT). The results show a clear performance

drop when any component is omitted. Removing both the transformer and LABTEXT results in a 4.81% drop in training precision and a 4.40% decrease in validation precision. The most substantial performance reduction occurs when both the transformer and TEXT are excluded, leading to a 29.76% decrease in training precision and a 30.66% decrease in validation precision. This underscores the indispensable role of TEXT and the transformer in our method. Even when only TEXT is removed, performance significantly deteriorates, with a 17.14% decline in training precision and a 14.60% decline in validation precision. These findings illustrate that each component contributes significantly to the model’s overall efficacy. Our full model, combining LLMs and MLTM, demonstrates the highest performance, with a training accuracy of 0.9535 and a validation accuracy of 0.9122.

5 Conclusion

In conclusion, we have presented a novel multi-disease diagnostic model that integrates multimodal data, closely mirroring real-life clinical decision-making. By combining fine-tuned LLMs with domain-specific transformers, we achieved enhanced synthesis of structured and unstructured medical data. Using a disentangled transformer further refined this integration, significantly improving disease prediction accuracy. Our experimental results across two practical EHR datasets demonstrated the proposed model’s robustness and effectiveness. Despite achieving high performance on two datasets, our approach may face overfitting and generalization challenges when applied to specific data types or healthcare environments. Utilizing embeddings and fine-tuning LLMs proved more efficient than pretraining a new model from scratch. However, computational challenges may persist. In future work, we will enhance model transparency, expand coverage to include complex and rare diseases, assess more diverse datasets, and combine real-time and other data modalities [35] in dynamic clinical environments.

References

- [1] AI@Meta. 2024. Llama 2 Model Card. (2024). <https://huggingface.co/meta-llama/Llama-2-13b>
- [2] AI@Meta. 2024. Llama 3 Model Card. (2024). https://github.com/meta-llama/llama3/blob/main/MODEL_CARD.md
- [3] David R Bellamy, Bhawesh Kumar, Cindy Wang, and Andrew Beam. 2023. Labrador: Exploring the Limits of Masked Language Modeling for Laboratory Data. *arXiv preprint arXiv:2312.11502* (2023).
- [4] Elliot Bolton, Abhinav Venigalla, Michihiro Yasunaga, David Hall, Betty Xiong, Tony Lee, Roxana Daneshjou, Jonathan Frankle, Percy Liang, Michael Carbin, et al. 2024. BioMedLM: A 2.7 B Parameter Language Model Trained On Biomedical Text. *arXiv preprint arXiv:2403.18421* (2024).
- [5] Tianxi Cai, Feiqing Huang, Ryumei Nakada, Linjun Zhang, and Doudou Zhou. 2024. Contrastive Learning on Multimodal Analysis of Electronic Health Records. *arXiv preprint arXiv:2403.14926* (2024).
- [6] Jindong Chen, Yizhou Hu, Jingping Liu, Yanghua Xiao, and Haiyun Jiang. 2019. Deep short text classification with knowledge powered attention. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 6252–6259.
- [7] Yahui Chen. 2015. *Convolutional neural network for sentence classification*. Master's thesis. University of Waterloo.
- [8] Yinda Chen, Haoyuan Shi, Xiaoyu Liu, Te Shi, Ruobing Zhang, Dong Liu, Zhiwei Xiong, and Feng Wu. 2024. TokenUnify: Scalable Autoregressive Visual Pre-training with Mixture Token Prediction. *arXiv preprint arXiv:2405.16847* (2024).
- [9] Edward Choi, Mohammad Taha Bahadori, Jimeng Sun, Joshua Kulas, Andy Schuetz, and Walter Stewart. 2016. Retain: An interpretable predictive model for healthcare using reverse time attention mechanism. *Advances in neural information processing systems* 29 (2016).
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [11] Jun-En Ding, Phan Nguyen Minh Thao, Wen-Chih Peng, Jian-Zhe Wang, Chun-Cheng Chug, Min-Chen Hsieh, Yun-Chien Tseng, Ling Chen, Dongsheng Luo, Chi-Te Wang, et al. 2024. Large Language Multimodal Models for 5-Year Chronic Disease Cohort Prediction Using EHR Data. *arXiv preprint arXiv:2403.04785* (2024).
- [12] Yaqian Dun, Kefei Tu, Chen Chen, Chunyan Hou, and Xiaojie Yuan. 2021. Kan: Knowledge-aware attention network for fake news detection. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 35. 81–89.
- [13] Léo Grinsztajn, Edouard Oyallon, and Gaël Varoquaux. 2022. Why do tree-based models still outperform deep learning on typical tabular data? *Advances in neural information processing systems* 35 (2022), 507–520.
- [14] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. 2022. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 16000–16009.
- [15] Stefan Hegselmann, Alejandro Buendia, Hunter Lang, Monica Agrawal, Xiaoyi Jiang, and David Sontag. 2023. Tabllm: Few-shot classification of tabular data with large language models. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 5549–5581.
- [16] Gregory Holste, Douwe van der Wal, Hans Pinckaers, Rikiya Yamashita, Akinori Mitani, and Andre Esteva. 2023. Improved Multimodal Fusion for Small Datasets with Auxiliary Supervision. In *2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI)*. IEEE, 1–5.
- [17] Mohammad Hossin and Md Nasir Sulaiman. 2015. A review on evaluation metrics for data classification evaluations. *International journal of data mining & knowledge management process* 5, 2 (2015), 1.
- [18] Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7B. *arXiv preprint arXiv:2310.06825* (2023).
- [19] Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. 2016. MIMIC-III, a freely accessible critical care database. *Scientific data* 3, 1 (2016), 1–9.
- [20] Yikuan Li, Shishir Rao, José Roberto Ayala Solares, Abdelaali Hassaine, Rema Ramakrishnan, Dexter Canoy, Yajie Zhu, Kazem Rahimi, and Gholamreza Salimi-Khorshidi. 2020. BEHRT: transformer for electronic health records. *Scientific reports* 10, 1 (2020), 7155.
- [21] Paul Pu Liang, Yun Cheng, Xiang Fan, Chun Kai Ling, Suzanne Nie, Richard Chen, Zihao Deng, Faisal Mahmood, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2023. Quantifying & modeling feature interactions: An information decomposition framework. *arXiv e-prints* (2023), arXiv-2302.
- [22] Zachary C Lipton, Charles Elkan, and Balakrishnan Naryanaswamy. 2014. Optimal thresholding of classifiers to maximize F1 measure. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2014, Nancy, France, September 15-19, 2014. Proceedings, Part II 14*. Springer, 225–239.
- [23] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* (2019).
- [24] Junyu Luo, Muchao Ye, Cao Xiao, and Fenglong Ma. 2020. Hitanet: Hierarchical time-aware attention networks for risk prediction on electronic health records. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 647–656.
- [25] Fenglong Ma, Radha Chitta, Jing Zhou, Quanzeng You, Tong Sun, and Jing Gao. 2017. Dipole: Diagnosis prediction in healthcare via attention-based bidirectional recurrent neural networks. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*. 1903–1911.
- [26] Julio-Omar Palacio-Niño and Fernando Berzal. 2019. Evaluation metrics for unsupervised learning algorithms. *arXiv preprint arXiv:1905.05667* (2019).
- [27] David Restrepo, Chenwei Wu, Constanza Vásquez-Venegas, Luis Filipe Nakayama, Leo Anthony Celi, and Diego M López. 2024. DF-DM: A foundational process model for multimodal data fusion in the artificial intelligence era. *arXiv preprint arXiv:2404.12278* (2024).
- [28] ruslanmv. 2024. Medical-Llama3-8B-16bit: Fine-Tuned Llama3 for Medical Q&A. (2024). <https://huggingface.co/ruslanmv/Medical-Llama3-8B>
- [29] Xingchen Song, Di Wu, Binbin Zhang, Zhendong Peng, Bo Dang, Fuping Pan, and Zhiyong Wu. 2023. Zeroprompt: Streaming acoustic encoders are zero-shot masked lms. *arXiv preprint arXiv:2305.10649* (2023).
- [30] Arun James Thirunavukarasu, Darren Shu Jeng Ting, Kabilan Elangovan, Laura Gutierrez, Ting Fang Tan, and Daniel Shu Wei Ting. 2023. Large language models in medicine. *Nature medicine* 29, 8 (2023), 1930–1940.
- [31] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruiti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288* (2023).
- [32] Haoran Wu, Wei Chen, Shuang Xu, and Bo Xu. 2021. Counterfactual supporting facts extraction for explainable medical record based diagnosis with graph network. In *Proceedings of the 2021 conference of the north American chapter of the association for computational linguistics: human language technologies. 1942–1955*.
- [33] Xianli Zhang, Buyue Qian, Shilei Cao, Yang Li, Hang Chen, Yefeng Zheng, and Ian Davidson. 2020. INPREM: An interpretable and trustworthy predictive model for healthcare. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 450–460.
- [34] Yilan Zhang, Yingxue Xu, Jianqi Chen, Fengying Xie, and Hao Chen. 2024. Prototypical Information Bottlenecking and Disentangling for Multimodal Cancer Survival Prediction. *arXiv preprint arXiv:2401.01646* (2024).
- [35] Lijuan Zheng, Zihan Wang, Junqiang Liang, Shifan Luo, and Senping Tian. 2021. Effective compression and classification of ECG arrhythmia by singular value decomposition. *Biomedical Engineering Advances* 2 (2021), 100013.