

Next Visit Diagnosis Prediction via Medical Code-Centric Multimodal Contrastive EHR Modelling with Hierarchical Regularisation

Heejoon Koo

University College London
heejoon.koo.17@alumni.ucl.ac.uk

Abstract

Predicting next visit diagnosis using Electronic Health Records (EHR) is an essential task in healthcare, critical for devising proactive future plans for both healthcare providers and patients. Nonetheless, many preceding studies have not sufficiently addressed the heterogeneous and hierarchical characteristics inherent in EHR data, inevitably leading to sub-optimal performance. To this end, we propose NECHO, a novel medical code-centric multimodal contrastive EHR learning framework with hierarchical regularisation. First, we integrate multifaceted information encompassing medical codes, demographics, and clinical notes using a tailored network design and a pair of bimodal contrastive losses, all of which pivot around a medical codes representation. We also regularise modality-specific encoders using a parental level information in medical ontology to learn hierarchical structure of EHR data. A series of experiments on MIMIC-III data demonstrates effectiveness of our approach.

1 Introduction

Predicting a patient’s future diagnosis has been a longstanding objective in both academic and industrial healthcare sectors. Its significance is highlighted for healthcare providers with refining decision-making processes and resource allocation, and also for patients with effective future plans. By leveraging the extensive accumulation of EHR data, data-driven deep learning methodologies have achieved considerable advancements in the healthcare practices, particularly in next admissions diagnosis prediction (Choi et al., 2016a; Ma et al., 2018; Qiao et al., 2019; Zhang et al., 2020a).

However, most of previous studies have shown limited consideration into multifaceted and hierarchical properties inherent in EHR data. First, it is heterogeneous, encompassing a range of modalities including demographics (e.g. age), medical images (e.g., Computed Tomography), text (e.g.

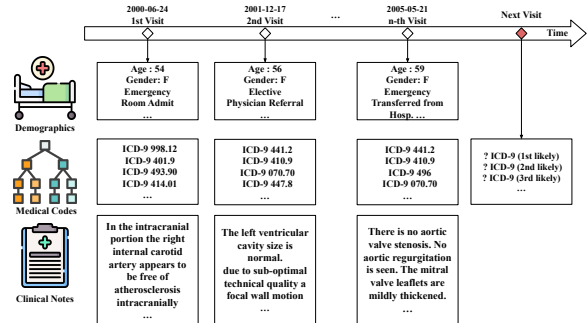


Figure 1: A Segment of Longitudinal EHR Data. It includes demographics, medical codes and clinical notes.

clinical notes), time series (e.g. laboratory tests), and medical codes (e.g. ICD-9). Each modality offers diverse and unique perspectives of a single observation and holds substantial potential to improve representative power if it is integrated seamlessly with other modalities. Nevertheless, the majority of previous works have solely focused on medical codes or shown limited exploration into effective multimodal fusion strategies (Choi et al., 2017; Zhang et al., 2020a; Yang and Wu, 2021).

Second, EHR data employs International Classification of Diseases (ICD) codes (Slee, 1978), an organised hierarchical medical concept ontology. It is used by domain experts to systematically categorise patient diagnoses into relevant medical concepts. For instance, in its ninth version (ICD-9), circulatory system diseases (ICD-9 code 390-459) are further categorised into 9 subcategories, each denoting specific conditions, such as hypertensive disease (ICD-9 code 401-405). Each is further divided into 10 subcategories (e.g. ICD-9 code 401.0 to 401.9). This shows a highly structured and hierarchical dependency amongst them. Despite the critical importance of these attributes, they have been largely overlooked in earlier studies.

To address the aforementioned characteristics of EHR data, we present a novel framework for Next Visit Diagnosis Prediction via Medical Code-

Centric Multimodal Contrastive EHR Modelling with Hierarchical Regularisation (NECHO). To the best of our knowledge, this framework is the first work designed in a medical code-centric fashion for diagnosis prediction. It tightly and seamlessly entangles three distinct modalities of medical codes, demographics, and clinical notes through a meticulously designed multimodal fusion network and two bimodal contrastive losses. Its goal is to boost representational power by positioning demographics and clinical notes as supplementary modalities. Furthermore, we harness an auxiliary loss to regularise each modality-specialised encoder based on the ancestral level of medical codes, thereby successfully injecting more general information from the ICD-9 medical ontology. Therefore, the main contributions of our work are threefold as follows:

- We effectively integrate three distinct modalities by developing a novel fusion network and a pair of bimodal contrastive losses, centralised around medical codes representation.
- We also propose to use auxiliary losses for each modality-specific model to regularise them using the parental level of medical codes to learn more general information, leveraging hierarchical nature of ICD-9 codes.
- Our proposed NECHO framework achieves superior performance over previous works on MIMIC-III (Johnson et al., 2016), a publicly available large-scale real-world healthcare data.

2 Related Works

2.1 Next Visit Diagnosis Prediction

AI research community has delved into future diagnosis predictions, employing various data modalities such as graph, text, or more than two. DoctorAI (Choi et al., 2016a) is the first work that predicts diagnoses utilising a simple recurrent neural networks (RNN). It is further refined to RETAIN (Choi et al., 2016b) and Dipole (Ma et al., 2017), which incorporate attention mechanisms.

Meanwhile, graph neural networks (GNN) have been influential, with models like GRAM (Choi et al., 2017) and KAME (Ma et al., 2018) constructing disease graphs from medical ontology, and others like MMORE (Song et al., 2019) and HAP (Zhang et al., 2020b) focusing on learning both ontology and diagnosis co-occurrence and leveraging

hierarchical attention, respectively. MIPO (Peng et al., 2021) predicts parental level medical codes based on the medical ontology additionally.

Biomedical domain specific pre-trained word2vec (Zhang et al., 2019) and language models have been introduced (Alsentzer et al., 2019) for clinical text understanding. The importance of them is particularly underscored in multimodal EHR learning (Husmann et al., 2022), often supplementing diverse prediction tasks. MNN (Qiao et al., 2019) and CGL (Lu et al., 2021) fuse medical codes and clinical notes. MAIN (An et al., 2021) further integrates demographics to learn more comprehensive information of patients. (Yang and Wu, 2021) explore multiple fusion strategies for clinical event prediction.

2.2 Multimodal Learning

Beyond EHR, multimodality learning has been explored to various domains, particularly in multimodal sentiment analysis (MSA) (Gandhi et al., 2022). We introduce a few works that have somewhat influenced our work.

First, Tensor Fusion Network (TFN) (Zadeh et al., 2017; Liu et al., 2018) and Multimodal Adaptation Gate (MAG) (Rahman et al., 2020) perform an outer product and attentional gate on representations from varying modalities, respectively. (Tsai et al., 2019) use cross-modal and self-attention transformers (Vaswani et al., 2017). (Yu et al., 2021) introduce Unimodal Label Generation Module (ULGM) to boost modality-wise representations. However, the above literature do not consider the modality imbalance, such as the superiority of text-based models. Based on such findings, text-centred multimodal fusion strategies have been developed (Qiu et al., 2022; Huang et al., 2023).

2.3 Contrastive Learning

Contrastive Learning has emerged as a predominant paradigm, showing its superior performance in many research areas recently. Originally, it aims to learn features from different views of a single sample and discriminate samples from different classes (Oord et al., 2018; Chen et al., 2020). Next, it is extended to multimodality. CLIP (Radford et al., 2021) is a seminal work on multimodal contrastive learning, employing InfoNCE loss (Oord et al., 2018) to learn transferable features between images and texts. (Zhang et al., 2022) apply this strategy to medical domain, whilst (Mai et al., 2022) exploit trimodal contrastive learning in MSA.

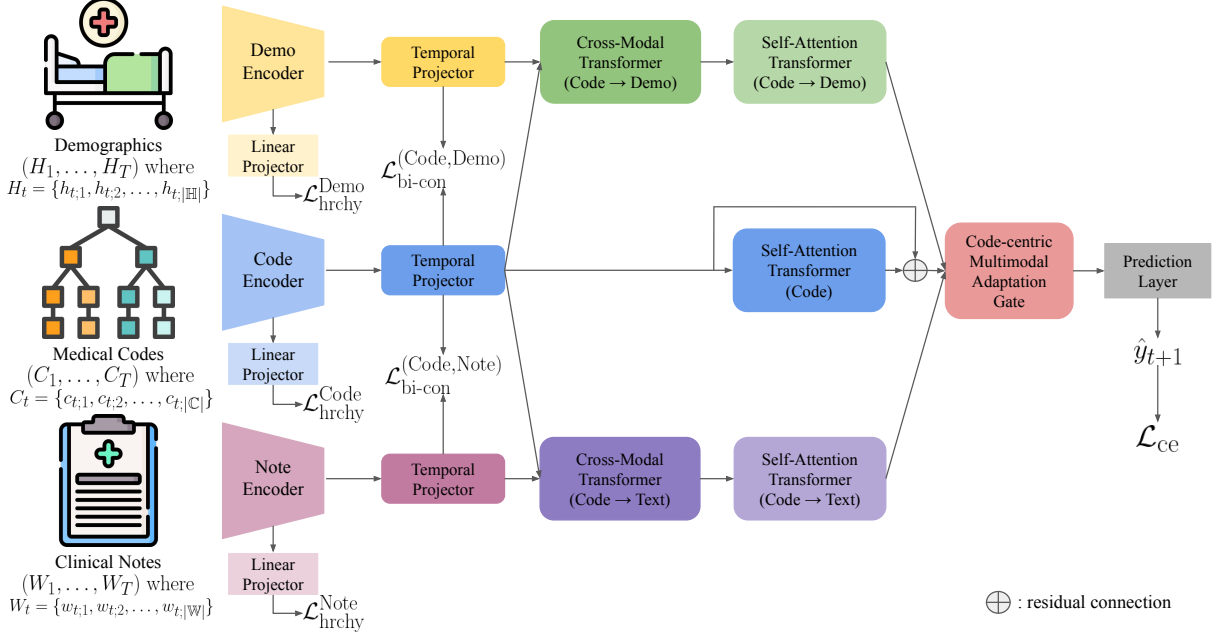


Figure 2: The Overall Framework of Our Proposed NECHO.

3 Methodology

In this section, we firstly introduce notations and problem formulation on next visit diagnosis prediction. Thereafter, we describe an overview and details of our proposed framework, NECHO.

3.1 Problem Formulation

Multimodal EHR Data A clinical record can be represented as a time-ordered sequence of visits V_1, \dots, V_T , where T is the total number of visits of any patient \mathcal{P} . Each visit V_t is denoted as (C_t, A_t, H_t, W_t) , where C_t is a set of diagnosis codes, A_t is a set of diagnosis codes at their ancestral level, H_t is demographics, W_t is a clinical note at t -th admission, respectively.

We denote a set of medical codes from EHR data as $c_1, c_2, \dots, c_{|\mathbb{C}|} \in \mathbb{C}$, where $|\mathbb{C}|$ is the number of unique medical codes at a level in ICD-9 code hierarchy \mathcal{G} . Similarly, a set of medical codes at their direct ancestral level is denoted as $a_1, a_2, \dots, a_{|\mathbb{A}|} \in \mathbb{A}$. The total number of unique medical codes in parental level is $|\mathbb{A}|$. Note that, $|\mathbb{A}| \ll |\mathbb{C}|$.

Diagnosis code at t -th visit is represented by $C_t = \{c_{t;1}, c_{t;2}, \dots, c_{t;|\mathbb{C}|}\}$, where $|\mathbb{C}|$ represents the number of diagnosis codes. Its ancestral level code is denoted by $A_t = \{a_{t;1}, a_{t;2}, \dots, a_{t;|\mathbb{A}|}\}$ with of the number of parental level diagnosis codes $|\mathbb{A}|$. Demographics is represented as $H_t =$

$\{h_{t;1}, h_{t;2}, \dots, h_{t;|\mathbb{H}|}\}$, where $|\mathbb{H}|$ is the total number of demographics features. Clinical note is represented as $W_t = \{w_{t;1}, w_{t;2}, \dots, w_{t;|\mathbb{W}|}\}$, where $|\mathbb{W}|$ is the maximum number of words to process.

Next Visit Diagnosis Prediction Task Based on the above notations, next visit diagnosis prediction is defined as follows. Given the patient’s multifaceted clinical records for the previous T visits, the objective is to predict a $(T + 1)$ -th visit’s diagnosis codes, denoted as \hat{y}_{T+1} .

3.2 Medical Code Information Centred Multimodal Fusion

One of the major challenges in the realm of AI healthcare is how to integrate the multifaceted data effectively. This has catalysed a surge of research on multimodal EHR learning (Zhang et al., 2020a; Yang and Wu, 2021). Nonetheless, a notable limitation in prior studies is the oversight of modality imbalance and the adoption of a modality-symmetric strategy, resulting in an unsatisfactory performance. We empirically observe that the medical code representations show the best performance. Also, previous works on MSA prioritise text representations at the core (Qiu et al., 2022; Huang et al., 2023) due to their superiority. Based on these findings, we introduce a novel medical code-centric multimodal fusion training scheme, which encompasses a tailored multimodal fusion network and a couple of bimodal contrastive losses.

3.2.1 Modality-Specific Feature Extraction

Before introducing our novel fusion strategies, we first explain modality-specific encoders that extract features from each modality. We design them as simple as possible to highlight the efficacy of our proposed fusion strategies. In other words, our framework is modular, with the potential for performance enhancement if the encoders are switched to more representative ones.

We employ a simple embedding layer for both medical codes and demographics, and a combination of BioWord2Vec (Zhang et al., 2019) and 1D CNN (Kim, 2014) to process clinical notes. Subsequently, the feature vector is passed to a fully connected layer (Linear) connected with ReLU activation function (Nair and Hinton, 2010).

$$\begin{aligned} M_t &= \text{Encoder}_m(m_t), \\ \bar{M}_t &= \text{ReLU}(\text{Linear}(M_t)) \end{aligned} \quad (1)$$

where m_t is a data of modality $m \in (C, H, W)$ at t -th visit and Encoder_m is a modality-specialised encoder, passing the feature vector M_t to MLP. Finally, a modality-specific feature \bar{M}_t is yielded. Appendix A provides a detailed information on how each modality-specific encoder operates.

3.2.2 Multimodal Fusion Network

Cross-Modal Transformer After acquiring representations from all modalities, we entangle them using two cross-modal transformers (CMTs), introduced by MulT (Tsai et al., 2019). It has verified its effectiveness in integrating meaningful information across different modalities. Initially, we put the each distinct representation to a temporal non-linear projector, 1D CNN:

$$\hat{H}_t^m = \text{Conv1D}(\bar{M}_t) \quad (2)$$

where \bar{M}_t is a representation from any modality m and \hat{H}_t^m is a resultant representation. Conv1D is equivalent to 1D CNN. Next, we introduce cross-modal attention, which facilitates the information transfer from the source modality to the target modality, e.g. medical codes \rightarrow clinical notes.

Let two modalities as m_1 and m_2 . Then, using trainable weights $W^{(\cdot)}$ with a dimension of d_k , we define the query, key and values as $Q^{m_1} = H^{m_1}W^{Q^{m_1}}$, $K^{m_2} = H^{m_2}W^{K^{m_2}}$, and $V^{m_2} = H^{m_2}W^{V^{m_2}}$, respectively. The cross-modal atten-

tion, denoted as CA, from m_1 to m_2 is then:

$$\begin{aligned} Z^{m_1 \rightarrow m_2} &= \text{CA}^{m_1 \rightarrow m_2}(\hat{H}^{m_1}, \hat{H}^{m_2}) \\ &= \text{Softmax}\left(\frac{Q^{m_1}(K^{m_2})^T}{\sqrt{d_k}}\right)V^{m_2}. \end{aligned} \quad (3)$$

We omit t for brevity. CMT is an extension of the CA. It is composed of a multi-head cross-modal attention block (MHA) and a Layer Normalisation layer (LM) (Ba et al., 2016). It is computed feed-forwardly for $i = 1, \dots, D$ layers as follows:

$$Z_{(0)}^{m_1 \rightarrow m_2} = H_{(0)}^{m_2}, \quad (4)$$

$$\begin{aligned} \hat{Z}_{(i)}^{m_1 \rightarrow m_2} &= \text{MHA}_{(i)}^{m_1 \rightarrow m_2}(\text{LM}(Z_{(i-1)}^{m_1 \rightarrow m_2}) \\ &\quad \text{LM}(H_{(0)}^T)) + \text{LM}(Z_{(i-1)}^{m_1 \rightarrow m_2}), \end{aligned} \quad (5)$$

$$\begin{aligned} Z_{(i)}^{m_1 \rightarrow m_2} &= f_{\theta_{(i)}^{m_1 \rightarrow m_2}}(\text{LM}(\hat{Z}_{(i)}^{m_1 \rightarrow m_2})) + \\ &\quad \text{LM}(\hat{Z}_{(i)}^{m_1 \rightarrow m_2}). \end{aligned} \quad (6)$$

During the process at MHA, the representations from the source modality are correlated with the target modality, enhancing the representational power across different modalities. As presented in Fig. 2, the fusion is performed in a medical code-centric fashion, thus we set m_1 as medical code C and m_2 as either demographics H or clinical notes W . Thus, we acquire two representations of $Z_t^{C \rightarrow H}$ and $Z_t^{C \rightarrow W}$ from the two CMTs.

Self-Attention Transformer To extract sequential feature representations effectively and boost dependencies from the above two cross-modal and medical code representations, a self-attention transformer (SA) is employed. It processes across the single-patient visits:

$$\begin{aligned} \hat{y}^C &= \text{SA}^C(\hat{H}^C), \\ \hat{y}^{C \rightarrow H} &= \text{SA}^{C \rightarrow H}(Z^{C \rightarrow H}), \\ \hat{y}^{C \rightarrow W} &= \text{SA}^{C \rightarrow W}(Z^{C \rightarrow W}). \end{aligned} \quad (7)$$

Additionally, we perform a residual connection (He et al., 2016) between the code representation before and after SA^C to enhance the influence of the medical code modality representation.

$$\hat{y}_t^C = \hat{y}_t^C + \hat{H}_t^C. \quad (8)$$

Multimodal Adaptation Gate Rather than performing a simple concatenation of the three distinct representations, we modify and adopt previous multimodal adaptation gate (MAG) (Rahman et al., 2020; Yang and Wu, 2021) in the medical code-centric manner. First, we calculate the trimodal

gating value $g \in \mathbb{R}$ and the displacement vector \mathbf{H} by concatenating (\oplus) meaningful representations in the previous stage as:

$$g = \text{Linear}([\hat{y}_t^C \oplus \hat{y}_t^{C \rightarrow H} \oplus \hat{y}_t^{C \rightarrow W}]), \quad (9)$$

$$\mathbf{H} = \text{Linear}(g[\hat{y}_t^{C \rightarrow H} \oplus \hat{y}_t^{C \rightarrow W}]). \quad (10)$$

This modification maximises the influence of medical code representation during the multimodal fusion process. Then, a weighted summation is performed between the medical code representation \hat{y}_t^C and the displacement vector \mathbf{H} to derive the multimodal representation \mathbf{M} :

$$\begin{aligned} \mathbf{M} &= \hat{y}_t^C + \alpha \mathbf{H}, \\ \text{where } \alpha &= \min\left(\frac{\|\hat{y}_t^C\|_2}{\|\mathbf{H}\|_2} \beta, 1\right). \end{aligned} \quad (11)$$

Here, α is a scaling factor, modulating the influence of the displacement vector \mathbf{H} and β is a trainable parameter that is randomly initialised. Both $\|\hat{y}_t^C\|_2$ and $\|\mathbf{H}\|_2$ are the L_2 norm of their respective entities. Finally, we apply a layer normalisation and dropout to \mathbf{M} .

Prediction To predict next visit diagnosis, we feed the representation \mathbf{M} in the previous stage into a single linear layer with a Sigmoid activation function to calculate the predicted probability \hat{y}_{t+1} .

$$\hat{y}_{t+1} = \text{Sigmoid}(\text{Linear}(\mathbf{M})), \quad (12)$$

$$\begin{aligned} \mathcal{L}_{ce} &= \frac{1}{T} \sum_{t=1}^T - (y_{t+1}^T \log \hat{y}_{t+1} + \\ &\quad (1 - y_{t+1})^T \log(1 - \hat{y}_{t+1})) \end{aligned} \quad (13)$$

where cross-entropy loss \mathcal{L}_{ce} is applied as the loss function. y_{t+1} is a ground truth with elements $|\mathbb{C}|$, which takes a value of 1 if the i -th code exists in V_{t+1} , otherwise 0.

3.2.3 Bimodal Contrastive Losses

Contrastive learning has been leveraged in multimodal pre-training literature (Radford et al., 2021; Zhang et al., 2022) to align diverse modalities effectively. Inspired by prior works, we apply two bimodal contrastive losses to further intricately entangle the different modalities by anchoring on the medical code representations.

Again, let two distinct modalities of m_1 and m_2 , where representation vectors derived from each modality be $\hat{H}_i^{m_1}$ and $\hat{H}_i^{m_2}$. Given a i -th pair of $(\hat{H}_i^{m_1}, \hat{H}_i^{m_2})$, our bimodal contrastive loss

scheme incorporates two asymmetric losses, m_1 -to- m_2 contrastive loss for the i -th pair and its inverse.

$$l_i^{(m_1 \rightarrow m_2)} = -\log \frac{\exp(\langle \hat{H}_i^{m_1}, \hat{H}_i^{m_2} \rangle / \tau)}{\sum_{k=1}^N \exp(\langle \hat{H}_i^{m_1}, \hat{H}_k^{m_2} \rangle / \tau)}, \quad (14)$$

$$l_i^{(m_2 \rightarrow m_1)} = -\log \frac{\exp(\langle \hat{H}_i^{m_2}, \hat{H}_i^{m_1} \rangle / \tau)}{\sum_{k=1}^N \exp(\langle \hat{H}_i^{m_2}, \hat{H}_k^{m_1} \rangle / \tau)} \quad (15)$$

where \langle, \rangle is cosine similarity and temperature $\tau \in \mathbb{R}^+$ is a parameter modulating distribution's concentration and Softmax function's gradient. Subsequently, a bimodal contrastive loss is determined by a weighted combination of $l_i^{(m_1 \rightarrow m_2)}$ and $l_i^{(m_2 \rightarrow m_1)}$ using a weighting parameter $\alpha \in [0, 1]$ and averaging over the mini-batch N as:

$$\begin{aligned} \mathcal{L}_{\text{bi-con}}^{(m_1, m_2)} &= \frac{1}{N} \sum_{i=1}^N (\alpha l_i^{(m_1 \rightarrow m_2)} + \\ &\quad (1 - \alpha) l_i^{(m_2 \rightarrow m_1)}). \end{aligned} \quad (16)$$

We apply this to two pairs, one between medical codes and demographics, and the other between medical codes and clinical notes.

$$\mathcal{L}_{\text{bi-con}} = \mathcal{L}_{\text{bi-con}}^{(C, H)} + \mathcal{L}_{\text{bi-con}}^{(C, W)}. \quad (17)$$

Note that, our multimodal contrastive loss is applied inter-modally, in line with the CLIP (Radford et al., 2021), rather than intra-modally. Moreover, we consider at the patient level rather than at the visit level. This is because patient level representations share similar patterns between their visits.

3.3 Hierarchical Regularisation

Medical ontologies organise diseases in a hierarchical manner. By effectively leveraging this, models are capable of acquiring knowledge at both general and specific levels of medical codes. This approach also mitigates the risk of error propagation and minimises the loss of pertinent information throughout the intricate multimodal fusion processes.

In ULGM (Yu et al., 2021), modality-tailored encoders are also tasked with predicting ground truths. Meanwhile, MIPO (Peng et al., 2021) introduces an auxiliary loss to learn parental level ICD-9 code prediction. Inspired by them, we introduce a regularisation strategy for each modality-specialised encoder to learn parental level of ICD-9 codes.

Specifically, the modality-specific features \bar{M}_t are passed to fully connected layers and Sigmoid activation function, yielding modality-specific

parental level prediction \hat{o}_t^m . Subsequently, we employ three cross-entropy losses, denoted as L_{hrchy}^m , to each modality m for this auxiliary task:

$$\hat{o}_{t+1}^m = \text{Sigmoid}(\text{Linear}(\bar{M}_t)), \quad (18)$$

$$\mathcal{L}_{\text{hrchy}}^m = \frac{1}{T} \sum_{t=1}^T - (o_{t+1}^T \log \hat{o}_{t+1}^m + (1 - o_{t+1}^T) \log(1 - \hat{o}_{t+1}^m)) \quad (19)$$

o_{t+1} is a ground truth with elements $|\mathbb{A}|$, where 1 is assigned if the i -th code presents in V_{t+1} and 0 if absent. This is re-written to encompass three distinct modalities as:

$$\mathcal{L}_{\text{hrchy}} = \mathcal{L}_{\text{hrchy}}^C + \mathcal{L}_{\text{hrchy}}^H + \mathcal{L}_{\text{hrchy}}^W. \quad (20)$$

3.4 Model Optimisation

The final objective function $\mathcal{L}_{\text{total}}$ is a weighted sum of three loss terms: the cross-entropy loss \mathcal{L}_{ce} between ground truth diagnosis and prediction, the medical code-centric two bimodal contrastive losses $\mathcal{L}_{\text{cont}}$, and the three modality-specific direct ancestral level hierarchical losses $\mathcal{L}_{\text{hrchy}}$. It is formulated as:

$$\mathcal{L}_{\text{total}} = \lambda_{\text{ce}} \mathcal{L}_{\text{ce}} + \lambda_{\text{bi-con}} \mathcal{L}_{\text{bi-con}} + \lambda_{\text{hrchy}} \mathcal{L}_{\text{hrchy}} \quad (21)$$

where λ_{ce} , λ_{cont} , and λ_{hrchy} are parameters that balance the different loss terms. The parameters of the model are updated via stochastic gradient descent (SGD) technique with respect to the calculated loss.

4 Experiments

4.1 Experimental Setup

4.1.1 Dataset

We conduct experiments on a publicly available large-scale, deidentified real-world EHR data, MIMIC-III (Johnson et al., 2016). It is acquired from intensive care units (ICU) patients at Beth Israel Deaconess Medical Center between 2001 and 2012. It contains multifaceted data, including ICD-9 medical codes, demographics, clinical notes, and so on. We provide descriptions on data pre-processing and the corresponding statistics to Appendix B.

4.1.2 Implementation Details

We describe the details for implementation. First, we set 256 and 0.1 as a hidden dimension and a dropout rate across the entirety of the model (e.g.

medical code and demographics feature extraction modules, Transformers including CMT and SA, and MAG), respectively. In the clinical note extraction module, filter sizes are set to [2, 3, 4], and the hidden dimension is 512. For the CMTs and SAs, we set the number of heads and encoder layers to be 4 and 3, respectively.

Also, following the previous work (Radford et al., 2021), the temperature τ and alpha α are 0.1 and 0.25 for the contrastive loss. The coefficients of loss terms, λ_{ce} , λ_{con} , and λ_{hrchy} are set to 1, 1, and 0.1, respectively. Especially, the λ_{hrchy} is set relatively small to weakly regularise each modality-specific encoder to learn the parental levels of ICD-9 codes, without overly constraining them. We provide the experimental results on the different λ_{hrchy} to Appendix C.

4.1.3 Training Details

We train models using Adam optimiser (Kingma and Ba, 2014) with a constant learning rate of 1e-4 and mini-batch size of 4, for a maximum of 50 epochs. The training is stopped if there is no gain for consecutive 5 epochs on validation data. Also, following the previous work (Choi et al., 2017), our proposed framework is evaluated using top- k accuracy, ranging k from 5, 10, 20 to 30. This is consistent with how physicians consider a comprehensive set of potential diagnoses, and is suitable for multi-label classification scenarios where multiple diseases often co-occur. Details on other baselines are provided to Appendix D.

Our proposed framework is implemented using PyTorch (Paszke et al., 2019) and accelerated via a single NVIDIA GeForce RTX 3090 GPU.

4.2 Experimental Results

4.2.1 Next Visit Diagnosis Prediction Results

Table 1 provides quantitative results of the proposed NECHO in comparison to the baselines on the MIMIC-III data for the diagnosis prediction task. NECHO notably excels over all existing baselines in EHR modelling and multimodal fusion strategies. Its effectiveness is attributed to its ability to leverage unique and complementary information from other modalities, which especially improves top-30 accuracy ranging from 0.5% to 10.7% over modality-specific encoders that constitute NECHO.

As shown in Table 1, the multimodal fusion is imperative. It’s noteworthy that whilst MAIN (An et al., 2021) employs a trimodal representation learning, its performance falls short compared to

Criteria	Modalities	Models	Acc@k			
			5	10	20	30
EHR Modelling	Code	GRAM (Choi et al., 2017)	24.16	36.47	52.48	62.76
		KAME (Ma et al., 2018)	25.34	36.93	54.25	64.54
		MMORE (Song et al., 2019)	25.97	38.58	57.05	68.23
		MIPO (Peng et al., 2021)	28.70	43.98	60.85	71.07
		Code Extractor (Ours)	28.16	41.83	57.99	68.31
	Demo	Demo Extractor (Ours)	17.96	29.58	47.13	58.94
	Note	BioWord2Vec _{10k} (Zhang et al., 2019)	27.31	41.14	58.53	69.21
		BioWord2Vec ₅₁₂ (Zhang et al., 2019)	23.05	35.74	52.76	63.20
		Clinical BERT ₅₁₂ (Alsentzer et al., 2019)	24.63	37.21	54.96	66.37
	Code + Note	MNN (Qiao et al., 2019)	28.16	41.83	59.75	69.44
	Code + Demo + Note	MAIN (An et al., 2021)	27.25	41.07	57.37	67.69
		NECHO _{w/o code centring} (Ours)	28.10	42.13	59.32	70.01
NECHO _{w/o \mathcal{L}_{hrcy}} (Ours)		28.71	43.14	59.83	70.22	
NECHO (Ours)		28.66	43.55	60.77	71.45	
NECHO _{w/MIPO} (Ours)		29.05	43.80	61.33	72.08	
Fusion Strategies	Code + Demo + Note	Concat	28.38	42.39	58.63	68.89
		TFN (Zadeh et al., 2017)	24.66	36.80	52.93	63.85
		MuT (Tsai et al., 2019)	28.27	41.87	58.12	68.50
		MAG (Rahman et al., 2020)	28.26	42.36	58.40	69.16
		ULGM (Yu et al., 2021)	28.58	42.09	58.70	68.53
		TeFNA (Huang et al., 2023)	28.12	41.78	59.11	69.21

Table 1: Experimental Results on MIMIC-III Data for Next Visit Diagnosis Prediction. Code, Demo, and Note are short for Medical Codes, Demographics, Clinical Notes, respectively. Best results are in boldface. 10k and 512 indicates the number of words. Unless specified otherwise, 10k words are processed for multimodal models with with clinical notes.

the bimodal MNN (Qiao et al., 2019). This discrepancy might arise from the harmful effects of improperly fusing demographic data lately. Especially, bimodal MNN shows comparable performance to trimodal fusion strategies baselines. This confirms the limitations of the tertiary symmetric multimodal fusion methodologies and raises the need for a medical code-centric approach, taking into account the modality imbalance.

To validate the efficacy of our fusion strategy, we compare NECHO that excludes the hierarchical regularisation (NECHO_{w/o \mathcal{L}_{hrcy}}) amongst multimodal EHR modelling and fusion baselines. Our method demonstrates superior performance over them, including NECHO_{w/o code centring}. These findings highlight the significance of designing multimodal fusion framework by centring medical codes representation that ensures a seamless aggregation of diverse data modalities. Furthermore, we also provide a comparative study on our novel code-centric MAG with others (Rahman et al., 2020; Yang and Wu, 2021) to Appendix E.

Next, we delve into the significance of regularising modality-specific encoders using parental level of medical codes. We juxtapose NECHO with

NECHO_{w/o \mathcal{L}_{hrcy}} and ULGM (Yu et al., 2021), at which modality-specific encoders learn the same level of medical ontology as the final prediction. They two show inferior performance, emphasising the importance of our novel strategy. It is discussed further in Ablation Studies (Section 4.2.2).

Furthermore, whilst NECHO does not completely surpass MIPO, replacing its simple medical code encoder with MIPO (NECHO_{w/MIPO}) outperforms MIPO. It especially achieves a 1.01% increase in top-30 accuracy, indicating that 1) our framework is modular, and 2) NECHO can predict additional accurate diseases than MIPO by leveraging complementary information from various modalities, emphasising its significance in real clinical settings. We provide a regarding case study to Section 4.2.3.

Another noteworthy point outside the multimodal strategies is that, amongst the clinical note baselines, Clinical BERT (Alsentzer et al., 2019) that is trained with a maximum of 512 tokens surpasses the combination model of BioWord2Vec (Zhang et al., 2019) and 1D CNN (Kim, 2014) with equivalent number of tokens but is inferior to that model trained with 10k tokens. This suggests that

enhancing performance is more about processing a large number of tokens than increasing model complexity in EHR learning. This also justifies our preference for BioWord2Vec over Clinical BERT within the realm of Pretrained Language Models.

4.2.2 Ablation Studies

We conduct ablation studies to discern influence of each module on the overall performance as: 1) individual modalities, 2) the multimodal fusion strategies (including Transformers, MAG, and bimodal contrastive losses), and 3) the hierarchical regularisation. The results are reported in Table 2.

Firstly, we assess the contribution of each modality within our proposed framework. The results demonstrate a clear superiority of the trimodal approach over its unimodal and bimodal ones. This underscores the unique representations from each modality are complementary to one another. Also, the significant performance degradation is observed upon the exclusion of medical code representation (w/o code), highlighting its pivotal role and rationalising our medical code-centred strategy. Additionally, whilst the exclusion of either notes or demographics similarly harms the performance, the note contains more meaningful information necessary than demographics, as shown in Table 1.

Secondly, we evaluate the impact of our medical code-centred strategies by removing each component. The resultant performance decline highlights their importance. Intriguingly, the performance disparities between models lacking transformers (w/o Transformers), lacking MAG (w/o MAG), and the full model (NECHO) widen as the value of k increases, suggesting an amplified effect in scenarios involving a broader range of disease sampling. Conversely, the influence of contrastive losses (w/o $\mathcal{L}_{\text{bi-con}}$) remains relatively stable across different top- k accuracies, indicating that they effectively align the distinct modalities in a semantically consistent fashion. These observations show that the adaptation of the proposed modules simultaneously is essential for effective inter-modality interaction and integration, thereby yielding significant performance enhancements.

Finally, the effectiveness of our novel parental level hierarchical regularisation is investigated. Its omission (w/o $\mathcal{L}_{\text{hrchy}}$) affects adversely model’s accuracy across various top- k accuracies. This suggests that enforcing the encoders for three distinct modalities, guided by the parental levels of medical codes using an ICD-9 hierarchy, is essential for en-

Criteria	Components	Acc@k	
		10	30
Modalities	w/o Code	36.78	65.54
	w/o Demo	42.56	70.12
	w/o Note	41.94	69.00
Multimodal Fusion	w/o Transformers	42.93	69.68
	w/o MAG	42.77	69.48
	w/o $\mathcal{L}_{\text{bi-con}}$	42.69	70.84
Hierarchical Regularisation	w/o $\mathcal{L}_{\text{hrchy}}$	43.14	70.22
NECHO	Full	43.55	71.45

Table 2: Ablation Studies on MIMIC-III Data.

hancing performance as it injects the general information and thus prevents the possible transmission of erroneous information when combining representations from distinct data modalities, thereby encouraging effective and accurate training.

4.2.3 Case Study

To qualitatively evaluate the predictive performance between MIPO (Peng et al., 2021) and our NECHO, we present a case study (Table 3) using a patient whose medical history shows a progression from a mitral valve issue to complications after surgery and cardiac rhythm disturbances. In the study, codes are formatted according to the Clinical Classifications Software (CCS) and are sequenced based on their priority, significantly influencing the reimbursement for treatment. We prefix them with "D" to make them appear akin to diagnosis codes.

Notably, our NECHO model accurately predicts 6 out of the top-10 diagnosis, outperforming MIPO, which predicts only 3. Firstly, both successfully identify D53 (Disorders of lipid metabolism), D106 (Cardiac dysrhythmias) and D101 (Coronary atherosclerosis and other heart disease), likely due to these diagnoses being part of the patient’s prior medical codes. However, NECHO uniquely predicts D238 (Complications of surgical procedures or medical care), D49 (Diabetes mellitus without complication), D2616 (E Codes: Adverse effects of medical care) and D96 (heart valve disorder) which MIPO fails to identify.

Additionally, our model predicts D238 and D2616 using multifaceted information of both demographics and notes. D238 should be predicted for two points: 1) the patient was initially hospitalised due to emergency health problem according to demographics, and 2) his notes states visual hallucinations, monitoring for pericardial and pleural effusions. The prediction of D2616 aligns with

Visit	Modalities / Models	Contents
Preceding	Demo	Age: 67, Gender: Male, Admission Type: Emergency, Admission Location: Transfer from hospital ...
	Codes	D96, D109, D97, D131, D101, D49, D110, D53, D138, D257
	Notes	... he was taken to the Operating Room where mitral valve replacement was performed ... Discharge Diagnosis: mitral valve mass ... He experienced some visual hallucinations ... IMPRESSION: 1. Enlarging bilateral pleural effusions. 2. Enlarging cardiac silhouette suspicious for a pericardial effusion, echocardiographic confirmation is suggested.
	Codes	D238, D53 , D130, D106, D101, D49 , D2, D3, D2616, D96
Subsequent	MIPO	D101 , D128, D53 , D108, D95, D259, D106 , D131, D98, D55
	NECHO	D96 , D98, D101, D53 , D138, D238, D49, D106, D2616 , D663

Table 3: Case Study of Next Visit Diagnosis Prediction for a Subject ID of 42129 in MIMIC-III Data. The preceding visit part provides a comprehensive information of a patient on demographics, medical codes, and clinical notes whilst the subsequent visit provides the patient’s real medical codes along with predicted ones by MIPO and NECHO. The accurately predicted codes and their matching ground truths are both in boldface.

potential risks associated with mitral valve replacement. On the contrary, MIPO’s prediction of D259 (Residual codes; unclassified) and D131 (Respiratory failure; insufficiency; arrest (adult)), which is considered less informative and a simple repetition from previous patient visits. D2 (Septicemia) and D3 (Bacterial infection) are not explicitly mentioned in the patient’s history, thus extremely challenging to predict. Hence, this demonstrates the necessity of the effective multimodal fusion strategy for its capability of capturing complementary and unique information in other modalities, verifying the effectiveness of the NECHO.

Apart from multimodal EHR learning, the content following the "Impression" in the preceding notes is only explicitly found in radiology reports. This indicates the importance of considering all available clinical note types to acquire a thorough understanding of a patient’s information. This contrasts with previous findings (Hsu et al., 2020; Husmann et al., 2022) suggesting that certain specific note types are representative in EHR learning.

5 Conclusion

Next visit diagnosis prediction is beneficial in AI-driven healthcare applications and has shown remarkable progress. However, the multifaceted and hierarchical properties of EHR data are beyond the consideration for the most of existing studies. To address these limitations, we introduce the novel multimodal EHR modelling framework, NECHO. It effectively aggregates representations from three heterogeneous modalities through meticulously de-

signed multimodal fusion network and the pair of two bimodal contrastive losses in a medical code-centric manner. It also uses parental level information of ICD-9 codes to regularise each modality-specialised encoder to learn more general information. Experimental results including the ablation studies and case study on MIMIC-III data highlight the NECHO’s efficacy and superiority.

6 Limitations

Whilst our proposed framework demonstrates promising advancements in multimodal EHR modelling for next visit diagnosis prediction, it is not without its limitations.

From a data perspective, the model’s predictions are heavily biased to the training data. This means there’s a potential risk that the model might underperform when encountering patterns that is non-existent in the dataset or originating from the different healthcare settings. Additionally, from a model perspective, firstly, the framework’s applicability is confined and has not been extended to a variety of clinical event prediction tasks, such as mortality, re-admissions, and length of stay, where different modalities might take main status. Secondly, it operates under the assumption that all data modalities are readily and consistently available for every patient. However, this assumption is impractical in that the availability of data can be compromised due to device malfunctions or human errors.

We hope to mitigate aforementioned challenges in the near future, enhancing NECHO’s adaptability in real-world clinical scenarios.

Acknowledgement

We highly appreciate anonymous reviewers for their valuable comments that helped us to enhance quality and completeness of this manuscript.

References

- Emily Alsentzer, John R Murphy, Willie Boag, Wei-Hung Weng, Di Jin, Tristan Naumann, and Matthew McDermott. 2019. Publicly available clinical bert embeddings. *arXiv preprint arXiv:1904.03323*.
- Ying An, Haojia Zhang, Yu Sheng, Jianxin Wang, and Xianlai Chen. 2021. Main: Multimodal attention-based fusion networks for diagnosis prediction. In *2021 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 809–816. IEEE.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR.
- Edward Choi, Mohammad Taha Bahadori, Andy Schuetz, Walter F Stewart, and Jimeng Sun. 2016a. Doctor ai: Predicting clinical events via recurrent neural networks. In *Machine learning for healthcare conference*, pages 301–318. PMLR.
- Edward Choi, Mohammad Taha Bahadori, Le Song, Walter F Stewart, and Jimeng Sun. 2017. Gram: graph-based attention model for healthcare representation learning. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 787–795.
- Edward Choi, Mohammad Taha Bahadori, Jimeng Sun, Joshua Kulas, Andy Schuetz, and Walter Stewart. 2016b. Retain: An interpretable predictive model for healthcare using reverse time attention mechanism. *Advances in neural information processing systems*, 29.
- 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*.
- Ankita Gandhi, Kinjal Adhvaryu, Soujanya Poria, Erik Cambria, and Amir Hussain. 2022. Multimodal sentiment analysis: A systematic review of history, datasets, multimodal fusion methods, applications, challenges and future directions. *Information Fusion*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Chao-Chun Hsu, Shantanu Karnwal, Sendhil Mulinathan, Ziad Obermeyer, and Chenhao Tan. 2020. Characterizing the value of information in medical notes. *arXiv preprint arXiv:2010.03574*.
- Changqin Huang, Junling Zhang, Xuemei Wu, Yi Wang, Ming Li, and Xiaodi Huang. 2023. Tefna: Text-centered fusion network with crossmodal attention for multimodal sentiment analysis. *Knowledge-Based Systems*, 269:110502.
- Severin Husmann, Hugo Yèche, Gunnar Rätsch, and Rita Kuznetsova. 2022. On the importance of clinical notes in multi-modal learning for ehr data. *arXiv preprint arXiv:2212.03044*.
- 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):1–9.
- Swaraj Khadanga, Karan Aggarwal, Shafiq Joty, and Jaideep Srivastava. 2019. Using clinical notes with time series data for ICU management. *arXiv preprint arXiv:1909.09702*.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Jianxun Lian, Xiaohuan Zhou, Fuzheng Zhang, Zhongxia Chen, Xing Xie, and Guangzhong Sun. 2018. xdeepfm: Combining explicit and implicit feature interactions for recommender systems. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 1754–1763.
- Zhun Liu, Ying Shen, Varun Bharadhwaj Lakshminarasimhan, Paul Pu Liang, Amir Zadeh, and Louis-Philippe Morency. 2018. Efficient low-rank multimodal fusion with modality-specific factors. *arXiv preprint arXiv:1806.00064*.
- Chang Lu, Chandan K Reddy, Prithwish Chakraborty, Samantha Kleinberg, and Yue Ning. 2021. Collaborative graph learning with auxiliary text for temporal event prediction in healthcare. *arXiv preprint arXiv:2105.07542*.
- 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*, pages 1903–1911.
- Fenglong Ma, Quanzeng You, Houping Xiao, Radha Chitta, Jing Zhou, and Jing Gao. 2018. Kame:

- Knowledge-based attention model for diagnosis prediction in healthcare. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 743–752.
- Sijie Mai, Ying Zeng, Shuangjia Zheng, and Haifeng Hu. 2022. Hybrid contrastive learning of tri-modal representation for multimodal sentiment analysis. *IEEE Transactions on Affective Computing*.
- Vinod Nair and Geoffrey E Hinton. 2010. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pages 807–814.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Xueping Peng, Guodong Long, Sen Wang, Jing Jiang, Allison Clarke, Clement Schlegel, and Chengqi Zhang. 2021. Mipo: Mutual integration of patient journey and medical ontology for healthcare representation learning. *arXiv preprint arXiv:2107.09288*.
- Zhi Qiao, Xian Wu, Shen Ge, and Wei Fan. 2019. Mnn: multimodal attentional neural networks for diagnosis prediction. *Extraction*, 1(2019):A1.
- Feng Qiu, Wanzeng Kong, and Yu Ding. 2022. Intermulti: Multi-view multimodal interactions with text-dominated hierarchical high-order fusion for emotion analysis. *arXiv preprint arXiv:2212.10030*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.
- Wasifur Rahman, Md Kamrul Hasan, Sangwu Lee, Amir Zadeh, Chengfeng Mao, Louis-Philippe Morency, and Ehsan Hoque. 2020. Integrating multimodal information in large pretrained transformers. In *Proceedings of the conference. Association for Computational Linguistics. Meeting*, volume 2020, page 2359. NIH Public Access.
- Vergil N Slee. 1978. The international classification of diseases: ninth revision (icd-9).
- Lihong Song, Chin Wang Cheong, Kejing Yin, William K Cheung, Benjamin CM Fung, and Jonathan Poon. 2019. Medical concept embedding with multiple ontological representations. In *IJCAI*, volume 19, pages 4613–4619.
- Yao-Hung Hubert Tsai, Shaojie Bai, Paul Pu Liang, J Zico Kolter, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2019. Multimodal transformer for unaligned multimodal language sequences. In *Proceedings of the conference. Association for Computational Linguistics. Meeting*, volume 2019, page 6558. NIH Public Access.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Bo Yang and Lijun Wu. 2021. How to leverage the multimodal ehr data for better medical prediction? In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4029–4038.
- Wenmeng Yu, Hua Xu, Ziqi Yuan, and Jiele Wu. 2021. Learning modality-specific representations with self-supervised multi-task learning for multimodal sentiment analysis. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 10790–10797.
- Amir Zadeh, Minghai Chen, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. 2017. Tensor fusion network for multimodal sentiment analysis. *arXiv preprint arXiv:1707.07250*.
- Dongdong Zhang, Changchang Yin, Jucheng Zeng, Xi-aohui Yuan, and Ping Zhang. 2020a. Combining structured and unstructured data for predictive models: a deep learning approach. *BMC medical informatics and decision making*, 20(1):1–11.
- Muhan Zhang, Christopher R King, Michael Avidan, and Yixin Chen. 2020b. Hierarchical attention propagation for healthcare representation learning. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 249–256.
- Yijia Zhang, Qingyu Chen, Zhihao Yang, Hongfei Lin, and Zhiyong Lu. 2019. Biowordvec, improving biomedical word embeddings with subword information and mesh. *Scientific data*, 6(1):52.
- Yuhao Zhang, Hang Jiang, Yasuhide Miura, Christopher D Manning, and Curtis P Langlotz. 2022. Contrastive learning of medical visual representations from paired images and text. In *Machine Learning for Healthcare Conference*, pages 2–25. PMLR.

A Modality-Specific Feature Extraction Modules

A.1 Feature Extraction Module for Medical Codes

Medical codes, particularly those from ICD-9 codes, play a vital role in that they directly indicate a patient’s status. They are highly specific, unambiguous and succinct, thus they have acted as a primary modality for next admission diagnosis prediction and shown better performance than models leveraging other modalities. Hence, here in this task, we consider them as a main modality.

We employ a single embedding layer E_C to process a set of diagnosis codes at t -th patient record, c_t . The features are passed to a single linear layer followed by a ReLU activation function. It is formulated as:

$$\bar{c}_t = E_C(c_t), \quad (22)$$

$$\bar{C}_t = \text{ReLU}(\text{Linear}(\bar{c}_t)) \quad (23)$$

where \bar{C}_t represents a feature vector from medical code information of each patient \mathcal{P} at t -th visit.

A.2 Feature Extraction Module for Demographics

Each patient has unique demographics, such as gender, age, admission and discharge location, to just name a few. Those provide the supplementary but highly personalised information, allowing an improvement in predictive performance.

We capture the non-stationary nature of the aforementioned attributes across clinical records at the individual level. For example, variables such as age and insurance type may change over time. Thus, we employ a single embedding layer E_H^n to n -th attribute at t -th patient record, h_t^n . The features from each embedding layer are then concatenated (\oplus) and fed into a single linear layer paired with a ReLU activation function. It can be represented as:

$$\bar{h}_t = E_H^1(h_t^1) \oplus E_H^2(h_t^2) \oplus \dots \oplus E_H^n(h_t^n), \quad (24)$$

$$\bar{H}_t = \text{ReLU}(\text{Linear}(\bar{h}_t)) \quad (25)$$

where \bar{H}_t represents a feature vector from demographics of each patient \mathcal{P} at t -th visit.

A.3 Feature Extraction Module for Clinical Notes

Clinical notes inherently possess a free, unstructured format but carry a comprehensive insight into

a patient’s condition from the perspective of health-care provider. They offer potential diagnoses and planned procedures, providing complementary and supplementary information not explicitly specified in medical codes.

We leverage a combination of pre-trained BioWord2Vec (Zhang et al., 2019) (frozen during both training and inference) and 1D CNN (Kim, 2014), which is capable of processing more tokens with computational efficiency. Although many preceding studies utilise PLMs like Clinical BERT (Alsentzer et al., 2019), they are still limited by a 512-token maximum, preventing themselves from processing an entire note in a single visit. Thus, we do not utilise them here.

First, we combine all notes $W_t^1, W_t^2, \dots, W_t^K$ in a single patient visit V_t to generate a single note W_t . Then, using the pre-trained BioWord2Vec (Zhang et al., 2019) E_W , each discrete word w_t^n in the note W_t is mapped to a low-dimensional embedding space, generating e_t^n . With the maximum number of words $|\mathbb{W}|$, the word embeddings $e_t = (e_t^1, e_t^2, \dots, e_t^{|\mathbb{W}|})$ from the combined note W_t are then fed into the 1D CNN (Conv1D) with multiple filters with a subsequent max-pooling layer (Max) to generate the most salient features \bar{w}_t using a filter (equivalent to window size) f . The outputs from each filter are concatenated (\oplus) and passed to a linear layer with ReLU activation function. It yields the note representation \bar{W}_t at t -th visit of each patient \mathcal{P} . The aforementioned processes are mathematically described as follows:

$$W_t = W_t^1 \oplus W_t^2 \oplus \dots \oplus W_t^K, \quad (26)$$

$$e_t^n = E_W(w_t^n), \quad (27)$$

$$\bar{e}_t^f = \text{ReLU}(\text{Conv1D}^f(e_t)) \quad (28)$$

where $f \in [2, 3, 4]$,

$$\bar{w}_t^f = \text{Max}(\bar{e}_t^f), \quad (29)$$

$$\bar{w}_t = \bar{w}_t^2 \oplus \bar{w}_t^3 \oplus \bar{w}_t^4, \quad (30)$$

$$\bar{W}_t = \text{ReLU}(\text{Linear}(\bar{w}_t)). \quad (31)$$

B Data Pre-processing

Patient Selection Criteria We follow the previous work of GRAM (Choi et al., 2017). First, we select patients with minimum two visits. Also, we truncate visits beyond the 21st visit.

Demographics Processing Attributes such as age, gender, admission type, admission and discharge locations, and insurance type are considered. Patients with ages 0 or above 120 are excluded. The admission types encompass categories such as emergency, elective, and urgent whilst the insurance types include medicare, private, medicaid, government and self pay. The dataset also offers a diverse range of features for both admission and discharge locations.

Clinical Note Processing Even though some prior works (Hsu et al., 2020; Husmann et al., 2022) emphasise the significance of specific note types for EHR representation learning, we consider all available note types (e.g. radiology, discharge summary, and nursing) for universality.

We first pre-process the notes, following the previous work (Khadanga et al., 2019). It involves a removal of non-alphabetical characters, stopwords and conversion of uppercase to lowercase letters. Then, we add two special tokens to BioWord2Vec (Zhang et al., 2019), <UNK> and <PAD>, the same as those used in BERT (Devlin et al., 2018). They are initialised using matrices filled with zeros and uniform distribution, respectively. Any visit records lacking note information are excluded. Next, each note is tokenised with maximum 10k words using BioWord2Vec. This approach effectively captures the entirety of note information for approximately 85% of all the visits.

Medical Ontology & Label Construction Following the GRAM (Choi et al., 2017), a medical ontology is constructed based on ICD-9 codes using the Clinical Classifications Software (CCS) from the Healthcare Cost and Utilization Project¹. The labels are derived from nodes present in the primary² and secondary³ hierarchy of the ICD-9 codes. This renders the next visit diagnosis prediction task as a hierarchical multi-label multi-class classification.

Summary A comprehensive statistical summary of the pre-processed dataset is provided in Table 4.

Dataset	MIMIC-III
# of patients	6,812
# of visits	18,256
Avg. # of visits per patient	2.68
# of Training Data	5449
# of Validation Data	681
# of Test Data	682
# of unique ICD9 codes	4,138
Avg. # of ICD9 codes per visit	13.27
Max # of ICD9 codes per visit	39
# of category codes	265
Avg. # of category codes per visit	11.40
Max # of category codes per visit	34
# of disease typing code	17
Avg. # of disease typing codes per visit	6.68
Max # of disease typing codes per visit	15
# of Age	73
# of Gender	2
# of Admission Type	3
# of Admission Location	8
# of Discharge Location	16
# of Insurance Type	5
Avg. # of words per visit	6743
Max # of words per visit	239,102

Table 4: Statistics of the Pre-processed MIMIC-III Data.

¹<https://hcup-us.ahrq.gov/toolssoftware/ccs/ccs.jsp>

²<https://hcup-us.ahrq.gov/toolssoftware/ccs/AppendixCMultiDX.txt>

³<https://hcup-us.ahrq.gov/toolssoftware/ccs/AppendixASingleDX.txt>

C Experiments on the Coefficient for Hierarchical Regularisation

We assume that modality-specific encoders necessitate soft regularisation for two reasons: firstly, their representations are relatively incomplete in comparison to the full framework (NECHO); secondly, since the general information embodies a broader scope, it should not impose excessive constraints on these encoders during training.

The empirical results on Table 5, delineated on a logarithmic scale for λ_{hrchy} values ranging from 0.01, 0.1, to 1, substantiate our hypothesis. Notably, setting it as 0.1 enhances the overall model performance the most, thereby verifying its optimal effectiveness.

Coefficients	Values	Acc@k	
		10	30
λ_{hrchy}	0.01	42.24	70.09
	0.1	43.55	71.45
	1	43.02	70.82

Table 5: Experimental Results on MIMIC-III Data of the Coefficient for Hierarchical Regularisation, λ_{hrchy} .

D Baselines

D.1 Unimodal EHR Modelling Baselines

- GRAM (Choi et al., 2017) considers medical ontology with an attention mechanism.
- KAME (Ma et al., 2018) employs an attention mechanism at the knowledge level, specifically tailored for medical ontology.
- MMORE (Song et al., 2019) attentively learns both the multiple ontological representation and the co-occurrence statistics.
- MIPO (Peng et al., 2021) utilises an auxiliary task of disease typing task. In other words, it learns parental level ICD-9 codes additionally.
- Medical Code Encoder (Ours) employs a simple combination of embedding layers and a couple of linear layers, which are followed by ReLU and Sigmoid activation function. It is utilised in our pipeline. Refer to Appendix A.1 for details.
- Demographics Encoder (Ours) utilises a simple combination of attribute-specific embedding layers and two linear layers, whose subsequent layers are ReLU and Sigmoid activation function, respectively. It is employed in our pipeline. Refer to Appendix A.2 for details.
- BioWord2Vec (Zhang et al., 2019) model is combined with 1D CNN (Kim, 2014). For brevity, we simplify it as BioWord2Vec. It uses pre-trained embedding with 16,545,454 words (with an arbitrary addition of two special tokens), which are subsequently processed by 1D CNN. In our framework, this serves as the notes feature extraction module. Refer to Appendix A.3 for details.
- Bio-Clinical BERT (Alsentzer et al., 2019) is a derivative of the original BERT (Devlin et al., 2018) on bio-medical domain. It is trained on MIMIC-III dataset (Johnson et al., 2016) and has a maximum input sequence length of 512.

D.2 Multimodal EHR Modelling Baselines

Both MNN and MAIN process 10k words from a clinical note within a single visit. The parameters (e.g. hidden dimension, the number of heads) are set in accordance with the specifications detailed in their original paper.

- MNN (Qiao et al., 2019) is trained using both medical codes and clinical notes. It employs a single embedding layer for the former and a combination of BioWord2Vec 1D CNN for the latter. The fusion of representations from these two modalities is achieved through deep feature mixture (Lian et al., 2018) and bi-directional RNN with attention.
- MAIN (An et al., 2021) is a trimodal model, integrating medical codes, clinical notes, and demographics, which is akin to our approach. First, medical codes and clinical notes are fused using a combination of low-rank fusion (Liu et al., 2018) and cross-modal attention. Next, demographics is merged using low-rank fusion subsequently.

D.3 Multimodal Fusion Strategies Baselines

We employ the same feature extraction module as used in our approach for the subsequent baselines, and fuse different modalities using their proposed mechanisms. For fairness, we set the parameters as the same as ours.

- Concat, an abbreviation for concatenation, is a straightforward method that merges distinct modalities without any computations, ensuring a raw and unaltered integration.
- TFN (Tensor fusion Network) (Zadeh et al., 2017) executes an outer product on the representations of different modalities.
- MulT (Multimodal Transformer) (Tsai et al., 2019) utilises both cross-modal and self-attention transformers to integrate distinct modalities.
- MAG (Multimodal Adaptation Gate) (Rahman et al., 2020) refines the representation of one modality by adjusting it with a displacement vector, which is derived from the other modalities.
- ULGM (Unimodal Label Generation Module) (Yu et al., 2021) uses modality-specific encoders to predict the ground truths as well.
- TeFNA (Text Enhanced Transformer Fusion Network) (Huang et al., 2023) learns text-centric pairwise cross-modal representations.

E A Comparative Study on Different MAGs

We present a comparative analysis of various MAGs, including our newly developed code-centric MAG and others (Rahman et al., 2020; Yang and Wu, 2021). (Rahman et al., 2020) introduce MAG initially while MAG from (Yang and Wu, 2021) combines representations from different modalities at the sample level dynamically with an attention gate. They are replaced with our MAG in the framework for a comparison.

From the Table 6, it demonstrates the superiority of our method over preceding approaches. It can be attributed to the meticulous consideration of the modality imbalance, one of factors not adequately addressed by previous methodologies. This validates that considering the dominance of main modality is essential in multimodal modelling.

Criteria	Methodologies	Acc@k	
		10	30
MAG	(Rahman et al., 2020)	42.36	69.16
	(Yang and Wu, 2021)	42.24	70.22
	NECHO (Ours)	43.55	71.45

Table 6: Experimental Results on MIMIC-III Data on Different MAGs.