

Predictive Modeling with Temporal Graphical Representation on Electronic Health Records

Jiayuan Chen, Changchang Yin, Yuanlong Wang and Ping Zhang*

The Ohio State University

{chen.12930, yin.731, wang.16050, zhang.10631}@osu.edu

Abstract

Deep learning-based predictive models, leveraging Electronic Health Records (EHR), are receiving increasing attention in healthcare. An effective representation of a patient’s EHR should hierarchically encompass both the temporal relationships between historical visits and medical events, and the inherent structural information within these elements. Existing patient representation methods can be roughly categorized into sequential representation and graphical representation. The sequential representation methods focus only on the temporal relationships among longitudinal visits. On the other hand, the graphical representation approaches, while adept at extracting the graph-structured relationships between various medical events, fall short in effectively integrate temporal information. To capture both types of information, we model a patient’s EHR as a novel temporal heterogeneous graph. This graph includes historical visits nodes and medical events nodes. It propagates structured information from medical event nodes to visit nodes and utilizes time-aware visit nodes to capture changes in the patient’s health status. Furthermore, we introduce a novel temporal graph transformer (TRANS) that integrates temporal edge features, global positional encoding, and local structural encoding into heterogeneous graph convolution, capturing both temporal and structural information. We validate the effectiveness of TRANS through extensive experiments on three real-world datasets. The results show that our proposed approach achieves state-of-the-art performance.

1 Introduction

Electronic Health Records (EHR) play a crucial role in healthcare by enhancing the quality of care and improving the efficiency and accuracy of medical decisions. In recent years, with the advancement of artificial intelligence technologies, many deep learning models based on EHR data have been

extensively explored and applied in predictive modeling, such as patient risk assessment [Bai *et al.*, 2018; Luo *et al.*, 2020; Jiang *et al.*, 2023; Wang *et al.*, 2024], prediction of drug interactions [Zhang *et al.*, 2017; Yang *et al.*, 2021], and support for clinical decision-making [Yin *et al.*, 2022; Yang *et al.*, 2023c; Chen *et al.*, 2024].

Predictive modeling on EHR data is proposed to utilize clinical events from patients’ historical visits, such as diagnoses, procedures and medications to forecast their future health status. These clinical events are typically represented as a set of medical codes. To enhance the quality and accuracy of medical predictive modeling, models need to learn three hierarchical representations in EHR: the **patient-level**, **visit-level**, and **(medical) code-level**. Specifically, the patient level representation should capture the temporal relationships between historical visits to build a comprehensive health profile of the patient. Then, at the visit level, the model analyzes the various medical events encountered in each visit to understand the specific circumstances of each visit. Finally, at the medical code level, the model needs to examine the characteristics of individual medical events and their higher-order patterns across historical visits, as well as the correlations between different medical events.

Patients’ longitudinal EHR, as a series of successive visits, can be naturally conceptualized as a sequence, as illustrated in Figure 1a. Current literature extensively employs sequence-based models, such as RNN-based [Choi *et al.*, 2016; Gao *et al.*, 2020] and Transformer-based [Luo *et al.*, 2020; Yang *et al.*, 2023c], to track state changes across a patient’s historical visits. These models excel in extracting temporal information at the visit level. However, they fall short in discerning structured relationships between medical events within visits due to weak temporal associations among these events. In contrast, graph-based models, another prominent category of EHR models, are adept at understanding relationships between medical codes in individual visits, especially regarding their structured connections. Some models [Choi *et al.*, 2017; Ma *et al.*, 2018; Lu *et al.*, 2021; Jiang *et al.*, 2023] enhance medical concept representations by integrating elements like medical ontology trees or sub-graphs from medical knowledge graphs (KGs). As shown in Figure 1b, other approaches conceptualize EHR as assorted graphs, applying graph-based models to extract structured information. For instance, [Choi *et al.*, 2020] has introduced

*The corresponding author.

a Graph Convolutional Transformer, aimed at deciphering the graphical structure of medical codes. While graph-based models are proficient in learning enhanced visit-level representations through structured features, they face challenges in integrating temporal data. DDHGNN [Wen *et al.*, 2022] tries to remedy this by embedding dynamic temporal edges in heterogeneous patient graphs, achieving improved predictions but lacking an explicit focus on patient state changes over historical visits. This leads to a pivotal question: Is there a model capable of concurrently capturing all hierarchical latent relationships within EHRs, thereby achieving superior patient, visit, and medical code level representations?

To address this challenge, in this paper, we model EHR as a novel temporal heterogeneous graph, as depicted in the Figure 1c. The visit-level and medical code-level representations are learned through dedicated medical event nodes and time-dependent visit nodes, respectively. The patient graph is designed to model intrinsic graphical relationships via edges connecting medical event nodes to visit nodes. Additionally, it facilitates the propagation of temporal relationships reflective of changes in patient state through temporal edges interlinking these visits. Building upon the patient graph representation, we introduce the Temporal gRAph traNSformer (**TRANS**), a novel model designed to deeply unravel these complex relationships. In particular, our temporal convolution merges temporal embedding with heterogeneous message passing. This fusion allows for the extraction of graph-structured features from medical events within individual visits, while also capturing the complex, higher-order temporal dependencies that span across historical visits. We enhance the learning of graphical features by integrating local structural encoding and global positional encoding into the medical event node representations, facilitated through meta-paths. Additionally, a patient explainer has been developed to ensure that our predictions are interpretable. We highlight the contributions of this work as follows:

- We propose a novel EHR representation method, transforming patient EHR into a heterogeneous graph with distinct sequential visit nodes and medical event nodes. This approach adeptly captures both structural relationships among medical events and temporal dependencies across historical visits.
- A novel model **TRANS** with temporal graph convolution and the structural and positional encoding of heterogeneous graphs is proposed to extract the temporal graphical representation of a patient’s graph, leading to interpretable predictions.
- We conduct extensive experiments on three real-world datasets to demonstrate that both the constructed graph and the **TRANS** model outperform baseline methods, demonstrating the effectiveness of our approach to healthcare predictive modeling. The implementation code can be found at Github¹.

¹<https://github.com/The-Real-JerryChen/TRANS>

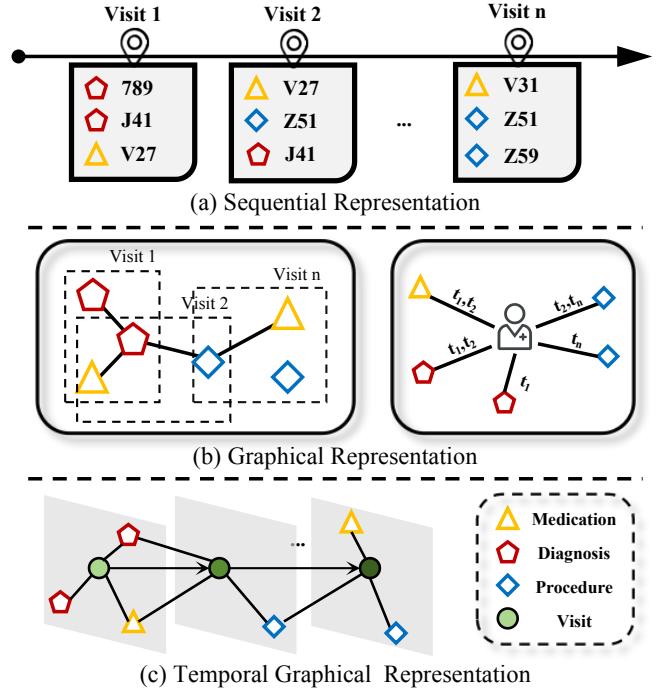


Figure 1: Different EHR representation methods. (a) An example of sequential representation. (b) Examples of graphical representation, which include graphs at the visit level and the patient level. (c) The temporal graphical representation proposed in this paper.

2 Preliminary

In this section, we provide some background knowledge about EHR data and then formulate the diagnosis prediction task.

2.1 Medical Codes

Each medical event can be represented as a set of unique medical codes, denoted as $\mathcal{C} = \{c_1, c_2, \dots, c_{|\mathcal{C}|}\}$, where $|\mathcal{C}|$ is the size of unique medical codes. Each code may represent a diagnosis, procedure, or medication. In this paper, we abstractly refer to these events as e .

2.2 EHR Data

The EHR of each patient is represented as a time-sequenced array of medical events, encompassing diagnoses, procedures, and medications. Specifically, the historical visit information of patient p can be represented as $\mathbf{X}^p = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T]$, where the t -th visit is denoted by a multi-hot vector $\mathbf{x}_t \in \{0, 1\}^{|\mathcal{C}|}$, $t = 1, \dots, T^p$ indicates the index of visit, and T^p denotes the total visit times of patient p . The i -th element of one visit vector is set to 1 if it contains the medical code c_i . To simplify, we explain the algorithms for one patient, omitting the superscript p when clear.

2.3 Sequential Diagnosis Prediction

The goal of predictive modeling is to predict the label at each time stamp or at the end of the visit sequence. In this study, we focus on the task of next-visit diagnosis prediction. It involves predicting diagnosis codes $\mathcal{C}_{diag} \in \{0, 1\}^{|\mathcal{C}|}$ based on

the medical events from historical visits as well as the medication codes and procedure codes for the current visit.

3 Methodology

In this section, we first introduce the method for temporal graph construction, followed by details of the proposed model TRANS. Figure 2 shows the overall architecture of TRANS, which is composed of three main modules: (1) heterogeneous temporal message passing; (2) spatial encoder; and (3) patient graph explainer.

3.1 Graph Construction

To capture the structured information of potential disease development trends, treatment pathways, and other temporal graphical relationships within the patient EHR, we construct a heterogeneous patient graph $G = (\mathcal{N}, \mathcal{E})$ for each patient, as illustrated in Figure 2, where \mathcal{N} and \mathcal{E} are sets of nodes and edges, respectively. The heterogeneous graph comprises four types of nodes: diagnoses, procedures, medications, and timestamped visit nodes: $\mathcal{N} = \mathcal{N}_e \cup \mathcal{N}_v$, where $|\mathcal{N}_v| = T$, $|\mathcal{N}_e| = |\mathbf{e}|$. Here, \mathbf{e} represents medical events that include diagnosis, medication, and procedure. The edge set contains two types of edges: (1) undirected edges between medical codes and visits when a visit includes medical events $\mathcal{E}_{ev} = \{(u_1, u_2) : \forall u_1 \in \mathcal{N}_e, \forall u_2 \in \mathcal{N}_v\}$; (2) directed edges between visits in chronological order $\mathcal{E}_{vv} = \{(u_t, u_{t+1})\} : \forall u \in \mathcal{N}_v, 0 \leq t < T$.

To the best of our knowledge, our graph construction differs from existing methods [Choi *et al.*, 2020; Wen *et al.*, 2022; Jiang *et al.*, 2023] of EHR representation learning. We explicitly represent each visit with time-dependent visit node. This approach decouples the representations of medical codes and the historical visits, allowing the embedding of clinical event nodes to focus on the characteristics of the codes themselves, while using the embedding of visit nodes to capture changes in the patient’s state. Additionally, we incorporate temporal embedding on the edges between medical codes and visits, as well as between successive visits with different time indexes.

3.2 Temporal Heterogeneous Message Passing

To effectively gather temporal information and advanced connections between medical events across historical visits, we propose a tailored temporal-aware message-passing mechanism for heterogeneous graphs. This system not only propagates medical code information into patient visit representations but also subtly connects various medical events through the visit nodes. We represent the features of medical code node and visit node outputted by the (l) -th layer of the network as $\mathbf{h}_e^{(l)} \in \mathbb{R}^d$ and $\mathbf{h}_v^{(l)} \in \mathbb{R}^d$, respectively. They also serve as the input for the subsequent $(l+1)$ -th layer. We denote linear mappings as $\mathbf{L}(\cdot)$, with each instance representing a distinct mapping.

Temporal Embedding

The temporal embedding in TRANS consists of two parts: the time encoding of the visit nodes; and the temporal encoding of edge features. We employ Time2Vec [Kazemi *et al.*, 2019]

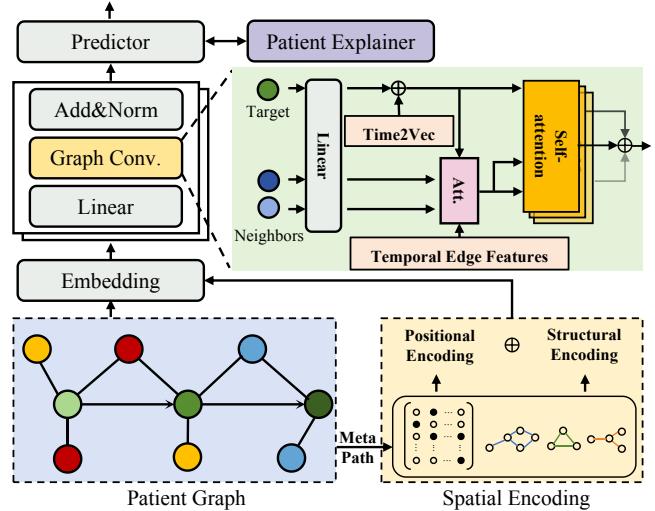


Figure 2: The overall framework of TRANS. The input EHR data is first constructed into a heterogeneous graph containing visit nodes and medical event nodes. Then the node features are mapped to an embedding space and combined with structural and sequential information. Subsequently, a temporal graph transformer is used to aggregate the node information. Finally, the features of the last visit node are input into a predictor.

to encode the absolute times of the visits, normalized for each patient, into vectors. This time-aware vector is then concatenated with the initial features of the visit node $\mathbf{h}_v^{(0)}$, thereby incorporating information about the time intervals between different visits.

Time2Vec is suitable for handling real-world timestamps, such as specific visit times, while for the indexed time of edges between medical events and visits, we use functional time encoding. Inspired by [Xu *et al.*, 2020], we use $\mathcal{T}(t)$ to encode visit index t into a continuous vector, which is then regulated by a linear layer to obtain the time factor α_t .

$$\mathcal{T}(t) = \sqrt{\frac{1}{d}} [\cos(\omega_1 t), \sin(\omega_1 t), \dots, \cos(\omega_d t), \sin(\omega_d t)],$$

where $\omega_1, \dots, \omega_d$ are learnable parameters. And $\alpha_t = \mathbf{L}(\mathcal{T}(t))$. We separate the temporal vector from node features, treating it as an edge type, to maintain focus on the nodes’ semantic information while infusing temporal data during message propagation into visit node features. Details can be found in the appendix.

Heterogeneous Message Passing

The entire message-passing process can be formalized into two stages: aggregation and updating. The first step is to aggregate the information of neighboring nodes. Specifically, we use an attention mechanism to weigh and integrate the features of neighboring nodes and concatenate the output from multi-head attention to obtain the final message. Taking as an example the process of propagating the features of neighbor nodes $e_t \in N(v_t)$ to the visit node v_t , the i -th head attention

is as follows:

$$\mathbf{q}_v = \mathbf{L}(\mathbf{h}_v^{(l)}), \quad \mathbf{k}_{e_t} = \mathbf{L}(\mathbf{h}_{e_t}^{(l)})$$

$$\text{Att}^i(e_t, v_t) = \frac{\alpha_t \cdot (\mathbf{q}_v \mathbf{W}^{\mu(e)} \mathbf{k}_e^\top)}{\sqrt{d}}, \quad (1)$$

where e_t represents the medical event node such as diagnosis, medication, and procedure contained in t -th visit as well as previous visit node v_{t-1} . The target node for message propagation is t -th visit node v_t . We use weight matrix $\mathbf{W}^{\mu(e)} \in \mathbb{R}^{d \times d}$ to encode different types of medical events, where $\mu(e)$ represents the type of the node e_t , such as diagnosis, procedure, or medication. After obtaining the attention scores for different neighboring nodes, we combine them with the mapped neighbor features to complete the entire aggregation process. We formalize it as follows:

$$\tilde{\mathbf{h}}_v^i = \text{Softmax}_{\forall e_t \in N(v_t)} (\text{Att}^i(e_t, v_t)) \cdot \mathbf{L}(\mathbf{h}_{e_t}^{(l)}), \quad (2)$$

where $N(v_t)$ denotes neighbors of v_t , and $\tilde{\mathbf{h}}_v^i$ represents the message passed to the visit node by the i -th head attention. Following the acquisition of aggregated information, the next stage is to combine this aggregated information with the target node's ego information to update the representation of the target node, as follows:

$$\mathbf{h}_v^{(l+1),i} = \gamma \cdot \mathbf{L}(\sigma(\tilde{\mathbf{h}}_v^i)) + (1 - \gamma) \cdot \mathbf{h}_v^{(l)},$$

$$\mathbf{h}_v^{(l+1)} = \big\|_{i \in [1, h]} \mathbf{h}_v^{(l+1),i} \quad (3)$$

where $\big\|_{i \in [1, h]}$ denotes concatenating the outputs of multiple heads, γ is the coefficient for the skip connection and σ represents activation function. Similar to the network structure in [Chen *et al.*, 2020], it can be either a learnable parameter or a hyper-parameter. Leveraging heterogeneous message passing, we integrate the features of medical events in the visit, along with the graphical structure information and historical visit information, into the current visit node, thereby updating the patient's representation.

3.3 Spatial Encoding

We integrate global positional encoding and local structural encoding into the initial feature of medical event nodes by concatenating them. This approach effectively captures structured features from patients' historical visits, which is crucial for illuminating key patterns like diagnosis sequences and treatment regimes, essential for accurately modeling disease stages. Such an integration enables us to discern intricate relationships in medical event co-occurrences and disease trajectories, thereby enhancing our predictive accuracy for a patient's future health status.

Local structural encoding $\mathcal{S}(e)$ is designed to capture local structural information in the graph based on random walks. Positional encoding $\mathcal{P}(e)$ adds a global positional encoding based on the Laplacian eigenvectors of the patient graph. However, in heterogeneous graphs, unlike homogeneous graphs [Rampášek *et al.*, 2022], direct computation of

certain properties is not straightforward, owing to the complexity introduced by diverse types of nodes and edges. To navigate this complexity, we employ medical event meta-paths, offering a structured and meaningful method to traverse the heterogeneous graph. For instance, in constructing meta-paths for medication nodes, visit nodes are used as intermediaries to connect multi-hop neighbors, forming a sequence represented as:

$$(N_m, N_v) - (N_v, N_d) - (N_d, N_v) - (N_v, N_m),$$

where N_m , N_v , and N_d represent medication, visit, and diagnosis nodes, respectively, and $-$ denotes the pseudo edge. This meta-path construction results in an adjacency matrix $\tilde{\mathbf{A}} \in \mathbb{R}^{|N_m| \times |N_m|}$, where $|N_m|$ is the number of unique medication nodes within graph.

Once the meta-paths are established, we can compute structural embedding based on the corresponding adjacency matrix $\tilde{\mathbf{A}}$. And for $\mathcal{P}(e)$, we construct a Laplacian matrix and compute its top k eigenvectors to represent the spatial information of nodes within the graph. $\mathcal{P}(e) \in \mathbb{R}^k$ for node e is expressed:

$$\mathcal{P}(e) = [\mathbf{v}_e^{(1)}, \mathbf{v}_e^{(2)}, \dots, \mathbf{v}_e^{(i)}, \dots, \mathbf{v}_e^{(k)}], \quad (4)$$

where each $\mathbf{v}_e^{(i)}$ is the element corresponding to node e in the i -th eigenvector. For structural encoding $\mathcal{S}(e) \in \mathbb{R}^k$, we calculate k -step random walk matrices $\tilde{\mathbf{W}}^k$ based on the adjacency matrix $\tilde{\mathbf{A}}$, and extract the cumulative probabilities of node e connecting to itself at each step to form $\mathcal{S}(e)$ as :

$$\mathcal{S}(e) = [\mathbf{d}_e^{(1)}, \mathbf{d}_e^{(2)}, \dots, \mathbf{d}_e^{(i)}, \dots, \mathbf{d}_e^{(k)}], \quad (5)$$

where $\mathbf{d}_e^{(i)}$ represents the diagonal element of node e after the i -th step of random walk. This approach not only enhances the representational ability of the graph but also provides a richer context for the deep analysis of complex interactions between different types of medical events.

3.4 Patient Explainer

Based on the patient graph, we are able to predict the future development of a patient's status using their historical diagnosis and medical information. However, the decision-making process of GNNs also requires intuitive and interpretable analysis to increase the credibility of decisions and to assist doctors in better understanding them. Therefore, we have incorporated a heterogeneous graph explainer module into our model. By masking certain medical event nodes, the explainer aims to identify a significant subgraph G_s crucial for the model's predictive results. Utilizing mutual information (MI) as the criterion, this process can be expressed as follows:

$$\max_{G_s} \text{MI}(\hat{Y}, G_s) \quad (6)$$

where \hat{Y} is the prediction of the TRANS model. The obtained subgraph contains the most influential medical events in the EHR for predicting the patient's status. We will provide a detailed introduction to this module in the appendix

Dataset	MIMIC-III	MIMIC-IV	CCAE
# of patients	5,449	14,155	565,897
# of visits	14,141	42,053	5,355,668
Avg. # of visits per patient	2.60	2.97	9.46
Max # of visits per patient	39	70	257
# of unique diagnoses	3,874	11,225	8,751
# of unique procedures	1,412	8,352	11,358
# of unique medicines	193	196	200
# of CCS codes	264	275	280

Table 1: Statistics of processed datasets.

4 Experiments

4.1 Experimental Setting

Datasets

We conduct experiments on the MIMIC-III [Johnson *et al.*, 2016], MIMIC-IV [Johnson *et al.*, 2023], and MarketScan Commercial Claims and Encounters (CCAE) [Butler *et al.*, 2021] datasets. MIMIC-III and MIMIC-IV are two publicly available medical databases released on PhysioNet. Following [Ma *et al.*, 2018], we filtered the original datasets to include patients who made at least two visits. The CCAE² dataset is a healthcare database, which contains patients with longer historical visit information. The goal of the diagnosis prediction task is to predict the diagnostic labels for the next visit. Similar to [Ma *et al.*, 2018], we chose to predict the category of the diagnosis, which is the CCS (Clinical Classifications Software) code corresponding to the ICD code. We use PyHealth [Yang *et al.*, 2023a] for parsing and preprocessing the dataset. Table 1 presents some statistics of the processed datasets. Details on **dataset processing** can be found in the Appendix.

Baselines

We evaluate our model against the following baselines:

Sequence-based Models. We use four sequence models based on different backbones. The Transformer [Vaswani *et al.*, 2017] consists of encoders stacked based on the self-attention mechanism. RETAIN [Choi *et al.*, 2016] is an RNN-based model that integrates a Reverse Time Attention mechanism. StageNet [Gao *et al.*, 2020] is enabled by a stage-aware LSTM module and a stage-adaptive convolutional module. HiTANet [Luo *et al.*, 2020] proposes a hierarchical time-aware Transformer model.

Graph-based Models. We compared our model with three graph-based models: KAME [Ma *et al.*, 2018], GCT [Choi *et al.*, 2020], and DDHGNN [Wen *et al.*, 2022]. KAME incorporates medical ontology knowledge of ICD codes into the sequence model. GCT utilizes graph convolutional networks to learn the graphical structure in EHR. We extend it by concatenating a transformer, making it suitable for sequential input tasks. DDHGNN constructs EHR as heterogeneous graphs of patients and employs dynamic graph convolution to learn patient’s features. Additionally, the results of some models [Hu *et al.*, 2020; Anonymous, 2023] can also be regarded as ablation studies for our proposed model.

²Available at <https://www.merative.com/real-world-evidence>.

The **experimental setup** and **implementation details** can be found in the Appendix.

Evaluation Metrics

We use two evaluation metrics [Ma *et al.*, 2018] for the diagnosis prediction task: visit-level precision@k and code-level accuracy@k. Visit-level precision@k is defined as the number of correct medical codes within the top k predictions, normalized by the minimum of k and the actual number of categories present in a visit. Code-level accuracy@k is the ratio of correct predictions to the total number of predictions made for individual codes. The value of k varies from 10 to 30, with higher values indicating better model performance, where visit-level precision assesses broader performance and code-level accuracy reflects more detailed accuracy.

4.2 Experimental Results

Table 2 demonstrates the performance of the proposed TRANS model and all the other baselines on three datasets, including MIMIC-III, MIMIC-IV, and CCAE. We ran 10 times for each model under random dataset splits and reported the mean and standard deviation of the results, where the best results are highlighted in **bold** and the second best scores are underlined.

Overall results. Compared to baseline models, our proposed model TRANS demonstrated SOTA performance across all three datasets. Notably, the improvements on the MIMIC-IV and CCAE datasets were more pronounced than on MIMIC-III. On the other hand, compared to MIMIC-III, almost all models performed better on the MIMIC-IV and CCAE datasets. This outcome suggests that more comprehensive datasets can enhance model training outcomes.

Sequence-based models. As mentioned before, sequence models treat patients’ historical visits as text, where visits and medical events correspond to sentences and words, respectively. Table 2 shows that for the MIMIC-III and MIMIC-IV datasets, it’s evident that the overall performance of sequence models is not as good as that of graph-based models. This may be attributed to the fact that in contrast to graph-based models like GCT and DDHGNN, sequence models are less effective at extracting structured information within a visit. Our proposed TRANS, which constructs a patient graph with temporally connected visit nodes, therefore achieves significant improvements. For the CCAE dataset, our model also outperforms those sequence model-based models, which further highlights the critical role of the graph module.

Graph-based models. In baseline models like KAME, tree-structured information of ICD codes is incorporated into the sequence structure. However, this does not deeply relate to the patient’s state, resulting in KAME’s suboptimal performance on several datasets. On the other hand, GCT, DDHGNN, and our proposed TRANS model integrate a graph feature extraction module for visit data into the sequential network, leading to notable improvements on datasets with shorter historical sequences like MIMIC-III and MIMIC-IV. Yet, as Table 2 indicates, these graph-based models are outperformed by ours on the CCAE dataset. This disparity may stem from their limited focus on visit characteristics, essential for capturing state changes in longer sequences

Dataset	Model	Visit-Level Precision@ k			Code-Level Accuracy@ k		
		10	20	30	10	20	30
MIMIC-III	Transformer	55.01 \pm 0.47	62.51 \pm 0.37	72.30 \pm 0.37	39.99 \pm 0.37	58.40 \pm 0.44	69.73 \pm 0.39
	RETAIN	55.92 \pm 0.31	63.51 \pm 0.29	73.15 \pm 0.25	40.56 \pm 0.37	59.27 \pm 0.35	70.55 \pm 0.29
	KAME	52.94 \pm 0.41	60.76 \pm 0.39	70.74 \pm 0.31	39.56 \pm 0.42	57.48 \pm 0.28	68.66 \pm 0.20
	StageNet	51.49 \pm 0.43	59.57 \pm 0.40	69.33 \pm 0.30	38.08 \pm 0.33	55.72 \pm 0.32	66.61 \pm 0.25
	HiTANet	56.33 \pm 0.38	63.66 \pm 0.35	73.38 \pm 0.29	40.95 \pm 0.30	59.87 \pm 0.27	70.68 \pm 0.21
	GCT	56.95 \pm 0.28	63.78 \pm 0.26	73.16 \pm 0.20	41.01 \pm 0.29	59.94 \pm 0.25	70.66 \pm 0.17
	DDHGNN	56.80 \pm 0.35	63.73 \pm 0.23	72.99 \pm 0.19	40.59 \pm 0.36	59.90 \pm 0.27	70.59 \pm 0.20
	TRANS (Ours)	57.65 \pm 0.25	64.18 \pm 0.17	73.69 \pm 0.16	41.30 \pm 0.28	60.19 \pm 0.19	71.17 \pm 0.12
MIMIC-IV	Transformer	62.27 \pm 0.39	65.98 \pm 0.25	74.66 \pm 0.15	40.32 \pm 0.33	60.07 \pm 0.26	71.62 \pm 0.17
	RETAIN	60.79 \pm 0.33	65.20 \pm 0.20	74.10 \pm 0.16	39.18 \pm 0.32	59.20 \pm 0.26	71.04 \pm 0.20
	KAME	61.16 \pm 0.40	64.55 \pm 0.33	73.03 \pm 0.20	39.42 \pm 0.41	58.72 \pm 0.29	70.01 \pm 0.19
	StageNet	57.89 \pm 0.44	61.97 \pm 0.31	71.21 \pm 0.22	37.15 \pm 0.46	55.87 \pm 0.30	67.80 \pm 0.21
	HiTANet	63.01 \pm 0.30	66.24 \pm 0.24	75.01 \pm 0.19	40.51 \pm 0.34	60.04 \pm 0.25	72.19 \pm 0.15
	GCT	63.15 \pm 0.27	66.39 \pm 0.21	75.05 \pm 0.20	41.00 \pm 0.30	60.30 \pm 0.24	72.23 \pm 0.21
	DDHGNN	59.67 \pm 0.33	64.61 \pm 0.26	73.32 \pm 0.18	39.11 \pm 0.19	59.54 \pm 0.15	70.16 \pm 0.13
	TRANS (Ours)	65.68 \pm 0.22	68.99 \pm 0.20	77.12 \pm 0.15	42.83 \pm 0.20	63.04 \pm 0.18	74.05 \pm 0.13
CCAE	Transformer	68.97 \pm 0.09	81.47 \pm 0.08	87.56 \pm 0.08	65.86 \pm 0.10	78.68 \pm 0.08	86.55 \pm 0.08
	RETAIN	68.14 \pm 0.10	77.12 \pm 0.09	87.14 \pm 0.10	65.60 \pm 0.08	76.98 \pm 0.07	85.92 \pm 0.07
	KAME	71.69 \pm 0.12	81.72 \pm 0.10	86.87 \pm 0.08	66.59 \pm 0.11	78.50 \pm 0.12	84.63 \pm 0.10
	StageNet	74.72 \pm 0.12	84.44 \pm 0.12	88.86 \pm 0.08	69.65 \pm 0.11	81.22 \pm 0.10	86.61 \pm 0.09
	HiTANet	<u>75.68 \pm 0.11</u>	<u>84.63 \pm 0.13</u>	<u>89.29 \pm 0.10</u>	<u>70.35 \pm 0.15</u>	<u>81.21 \pm 0.11</u>	<u>86.99 \pm 0.11</u>
	GCT	75.64 \pm 0.14	84.57 \pm 0.12	89.09 \pm 0.09	70.32 \pm 0.14	81.22 \pm 0.13	86.75 \pm 0.09
	DDHGNN	74.49 \pm 0.15	84.13 \pm 0.13	88.39 \pm 0.13	69.20 \pm 0.16	81.07 \pm 0.15	85.89 \pm 0.15
	TRANS (Ours)	81.81 \pm 0.07	89.29 \pm 0.05	92.82 \pm 0.06	76.47 \pm 0.07	86.15 \pm 0.06	90.70 \pm 0.04

Table 2: Results for the diagnosis prediction task on MIMIC-III, MIMIC-IV and CCAE dataset, where the best results are highlighted in **bold** and the second best scores are underlined.

for precise diagnosis. HiTANe’s good performance on longer sequence datasets also reflects its hierarchical approach to analyzing a patient’s historical visit stages.

4.3 Model Analyses

Ablation Study

In this section, we conduct ablation studies to investigate the contributions of different components. Specifically, we removed three key parts of TRANS, namely Spatial Encoding (SE), Temporal Embedding (TE), and Sequence information in node feature initialization (Seq). Spatial Encoding includes both structural encoding and positional encoding. Temporal Embedding comprises the Time2Vec features of visit nodes and the temporal embedding of edge features in the graph. As shown in Table 3, the experimental results first reveal that each of these components contributes to improving the model’s performance to some extent. Second, we observe that removing TE causes more performance degradation than removing SE, especially in the CCAE dataset where both Seq and TE play a more significant role. This suggests that graph convolution layers in TRANS are capable of capturing structured information to a certain degree, which overlaps with some features in SE. However, the inclusion of SE also results in better performance, indicating its important role.

Graph Construction

In this section, we delve into the distinctions among various graph construction methods. There are three notable approaches: (1) The first method, exemplified by GCT [Choi

SE	TE	Seq	MIMIC-III	MIMIC-IV	CCAE
\times	\times	\times	55.03 \pm 0.35	62.23 \pm 0.33	68.80 \pm 0.14
\times	\times	\checkmark	55.92 \pm 0.36	62.74 \pm 0.27	74.14 \pm 0.12
\checkmark	\times	\checkmark	56.73 \pm 0.32	63.97 \pm 0.27	76.75 \pm 0.10
\times	\checkmark	\checkmark	57.16 \pm 0.28	65.48 \pm 0.24	81.63 \pm 0.08
\checkmark	\checkmark	\checkmark	57.65 \pm 0.25	65.68 \pm 0.22	81.81 \pm 0.07

Table 3: Variants of TRANS. We only report average precision of the visit-level Top@10 due to the space limitations.

et al., 2020], prioritizes the structural relationships between medical events within individual visits, omitting temporal data. This method aggregates medical event nodes to create a patient representation. (2) The second approach, as introduced by DDHGNN [Wen *et al.*, 2022], conceptualizes EHR as dynamic graphs with patient nodes at their core. It infuses temporal information into the dynamic edges linked to medical events. (3) Our proposed method, the third approach, models EHRs utilizing both time-dependent visit nodes and medical event nodes. The comparative performance of these methods is showcased in Table 2, where our method demonstrates superior efficacy.

DDHGNN’s graph construction, while acknowledging historical visit temporal information, incorporates it solely into dynamic edge features. This leads to a conflation of medical event node representations with patient state data, diluting the focus on specific ICD code attributes. In our analysis of the CCAE dataset, we selected 4,000 codes that reflect a balanced

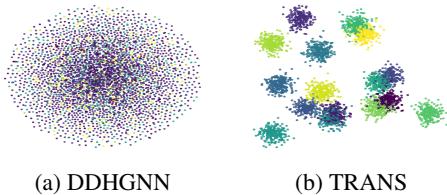


Figure 3: t-SNE Scatterplots of medical codes Learned by DDHGNN and TRANS on the CCAE dataset.

distribution across disease categories. We visualized these codes in a 2D space using t-SNE (Figure 3), where different colors represent different categories. The clustering patterns of ICD code representations in our model (Figure 3b) are notably distinct and category-specific, unlike the more chaotic representations observed in DDHGNN (Figure 3a). This observation underscores the efficacy and superiority of our approach in graph construction.

4.4 Graph Explainer

We employ a patient graph explainer to enhance the interpretability of our model’s predictions, which is crucial for dissecting the predictive role of each node, including medical event nodes like medications and diagnoses, and visit nodes. Leveraging our capability to discern the significance of various nodes, we have performed a detailed statistical analysis of key historical diagnoses within the CCAE dataset, which specifically pertains to patients with Chronic Obstructive Pulmonary Disease (COPD). For each COPD patient (labeled with CCS code ‘127’), we calculate the importance of each diagnosis occurrence and then average these values across the entire dataset. Our analysis of the diagnostic nodes related to COPD allows us to pinpoint several potential comorbidities. As depicted in Table 4, the explainer has identified a range of likely comorbidities, systematically categorized by disease type. Notably, these findings, encompassing prevalent complications such as Asthma, Bronchitis, and Hypertension, resonate strongly with established clinical knowledge³, demonstrating the efficacy of the TRANS on predictive modeling.

5 Related Work

In this section, We review two main EHR representation methods: sequence-based and graph-based, along with their predictive modeling algorithms.

Sequence-based models. Sequence-based models utilize sequential models to capture temporal information in EHR. In sequence-based studies, patients’ EHRs are represented a sequence of visit. Each visit contains a list of medical events (e.g., diagnosis codes and medications). These studies focus on extracting temporal or contextual features between visits. For example, the classic model RETAIN [Choi *et al.*, 2016] introduces a attention-based combination of forward and backward visit-level representations to obtain more accurate predictions. Dipole [Ma *et al.*, 2017] employs bidirectional recurrent neural networks combined with three dis-

³<https://www.mayoclinic.org/diseases-conditions/copd/symptoms-causes/syc-20353679>

Respiratory System Diseases

- J06 - Acute upper respiratory infections
- J20 - Acute bronchitis
- J45 - Asthma

Mental and behavioral disorders

- F32 - Major depressive disorder
- F17 - Nicotine dependence

Circulatory System Diseases

- I10 - Essential hypertension
- E78 - Disorders of lipoprotein metabolism

Table 4: Potential COPD comorbidities grouped by disease categories, as identified by the Patient Explainer.

tinct attention mechanisms for patient visit information prediction. More recently, many works [Gao *et al.*, 2020; Luo *et al.*, 2020; Chen *et al.*, 2022; Yang *et al.*, 2023c] have applied the powerful Transformer-based model adapted from natural language processing like BERT [Devlin *et al.*, 2018] to EHR, treating each visit or medical code as a token.

Graph-based models. Different from sequence-based models, graph-based models represent EHR in the form of graphs, which naturally model the inherent structural information among medical events. For instance, many models leverage external medical KGs to improve medical code representation learning. To name a few, GRAM [Choi *et al.*, 2017] utilizes medical ontologies and graph-based attention to develop strong medical code representations, while CGL [Lu *et al.*, 2021] utilizes disease domain knowledge and unstructured text. KerPrint [Yang *et al.*, 2023b] proposed a time-aware KG attention method to understand knowledge decay over time. There is also considerable research focusing on the extraction of graph structural features within EHR. [Choi *et al.*, 2018; Choi *et al.*, 2020] explored the prior and hidden structured relationships between medical events. [Wen *et al.*, 2022] expanded the graph from visit-level to patient-level, constructing a dynamic graph to represent EHR. However, existing graph-based models fail to simultaneously extract the temporal and structured relationships between visits.

6 Conclusion

In this paper, we investigate the three levels of relationships with EHR and the limitations of existing patient representation methods for predictive modeling. To capture both temporal and structural information about EHR, we propose a novel temporal heterogeneous graph to represent each patient’s EHR. Moreover, we propose a TRANS model to learn the three-level relationships, which use temporal graph transformer to extract temporal and structured information from the graph. We validated our methods and the TRANS model through experiments on three EHR datasets, demonstrating superior performance in diagnosis prediction. Utilizing KGs in healthcare AI is also a crucial research area, particularly for improving the effectiveness of predictive models in the field. We will explore the expansion of our models through the incorporation of KGs in the future work.

Ethical Statement

The datasets used in the paper are from PhysioNet (MIMIC-III and MIMIC-IV) and Mерative (CCAE), which are fully HIPAA-compliant de-identified, and have very minimal risk of the potential for loss of privacy. Moreover, Per the DUAs, all users to access the data will need to take full research, ethics, and compliance training courses. Thus, potential privacy and security risks would be eliminated and/or mitigated.

Acknowledgments

This work was funded in part by the National Science Foundation under award number IIS-2145625, by the National Institutes of Health under award number R01GM141279, and by The Ohio State University President's Research Excellence Accelerator Grant.

References

[Anonymous, 2023] Anonymous. Graph transformers on EHRs: Better representation improves downstream performance. In *Submitted to The Twelfth International Conference on Learning Representations*, 2023. under review.

[Bai *et al.*, 2018] Tian Bai, Shanshan Zhang, Brian L. Egleston, and Slobodan Vucetic. Interpretable representation learning for healthcare via capturing disease progression through time. In *KDD*, 2018.

[Butler *et al.*, 2021] Anne M. Butler, Katelin B. Nickel, Robert A. Overman, and M. Alan Brookhart. *IBM MarketScan Research Databases*, pages 243–251. 2021.

[Chen *et al.*, 2020] Ming Chen, Zhewei Wei, Zengfeng Huang, Bolin Ding, and Yaliang Li. Simple and deep graph convolutional networks. In *ICML*, 2020.

[Chen *et al.*, 2022] Xianlai Chen, Jiamiao Lin, and Ying An. Dl-bert: a time-aware double-level bert-style model with pre-training for disease prediction. In *2022 IEEE International Conference on Big Data (Big Data)*, pages 1801–1808. 2022.

[Chen *et al.*, 2024] Jiayuan Chen, Yu Wang, Ruining Deng, Quan Liu, Can Cui, Tianyuan Yao, Yilin Liu, Jianyong Zhong, Agnes B Fogo, Haichun Yang, et al. Spatial pathomics toolkit for quantitative analysis of podocyte nuclei with histology and spatial transcriptomics data in renal pathology. In *Medical Imaging 2024: Digital and Computational Pathology*. SPIE, 2024.

[Choi *et al.*, 2016] Edward Choi, Mohammad Taha Bahadori, Joshua A. Kulas, Andy Schuetz, Walter F. Stewart, and Jimeng Sun. Retain: An interpretable predictive model for healthcare using reverse time attention mechanism. In *NeurIPS*, page 3512–3520, 2016.

[Choi *et al.*, 2017] Edward Choi, Mohammad Taha Bahadori, Le Song, Walter F. Stewart, and Jimeng Sun. Gram: Graph-based attention model for healthcare representation learning. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, page 787–795, 2017.

[Choi *et al.*, 2018] Edward Choi, Cao Xiao, Walter F. Stewart, and Jimeng Sun. Mime: multilevel medical embedding of electronic health records for predictive healthcare. In *NIPS*, page 4552–4562, 2018.

[Choi *et al.*, 2020] Edward Choi, Zhen Xu, Yujia Li, Michael Dusenberry, Gerardo Flores, Emily Xue, and Andrew Dai. Learning the graphical structure of electronic health records with graph convolutional transformer. *AAAI*, 34:606–613, 04 2020.

[Devlin *et al.*, 2018] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

[Gao *et al.*, 2020] Junyi Gao, Cao Xiao, Yasha Wang, Wen Tang, Lucas M. Glass, and Jimeng Sun. Stagenet: Stage-aware neural networks for health risk prediction. In *Proceedings of The Web Conference 2020*, page 530–540, 2020.

[Hu *et al.*, 2020] Ziniu Hu, Yuxiao Dong, Kuansan Wang, and Yizhou Sun. Heterogeneous graph transformer. In *Proceedings of The Web Conference 2020*, page 2704–2710, 2020.

[Jiang *et al.*, 2023] Pengcheng Jiang, Cao Xiao, Adam Cross, and Jimeng Sun. Graphcare: Enhancing healthcare predictions with personalized knowledge graphs, 2023.

[Johnson *et al.*, 2016] Alistair Johnson, Tom Pollard, Lu Shen, Li-wei Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Celi, and Roger Mark. Mimic-iii, a freely accessible critical care database. *Scientific Data*, 3, 2016.

[Johnson *et al.*, 2023] Alistair Johnson, Lucas Bulgarelli, Lu Shen, Alvin Gayles, Ayad Shammout, Steven Horng, Tom Pollard, Sicheng Hao, Benjamin Moody, Brian Gow, Li-wei Lehman, Leo Celi, and Roger Mark. Mimic-iv, a freely accessible electronic health record dataset. *Scientific Data*, 2023.

[Kazemi *et al.*, 2019] Seyed Mehran Kazemi, Rishab Goel, Sepehr Eghbali, Janahan Ramanan, Jaspreet Sahota, Sanjay Thakur, Stella Wu, Cathal Smyth, Pascal Poupart, and Marcus Brubaker. Time2vec: Learning a vector representation of time. *arXiv preprint arXiv:1907.05321*, 2019.

[Lu *et al.*, 2021] Chang Lu, Chandan Reddy, Prithwish Chakraborty, Samantha Kleinberg, and Yue Ning. Collaborative graph learning with auxiliary text for temporal event prediction in healthcare. In *IJCAI*, pages 3529–3535, 08 2021.

[Luo *et al.*, 2020] Junyu Luo, Muchao Ye, Cao Xiao, and Fenglong Ma. 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*, 2020.

[Ma *et al.*, 2017] Fenglong Ma, Radha Chitta, Jing Zhou, Quanzeng You, Tong Sun, and Jing Gao. 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, 2017.

[Ma *et al.*, 2018] Fenglong Ma, Quanzeng You, Houping Xiao, Radha Chitta, Jing Zhou, and Jing Gao. Kame: Knowledge-based attention model for diagnosis prediction in healthcare. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, page 743–752, 2018.

[Rampášek *et al.*, 2022] Ladislav Rampášek, Mikhail Galkin, Vijay Prakash Dwivedi, Anh Tuan Luu, Guy Wolf, and Dominique Beaini. Recipe for a General, Powerful, Scalable Graph Transformer. *Advances in Neural Information Processing Systems*, 35, 2022.

[Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 2017.

[Wang *et al.*, 2024] Yuanlong Wang, Changchang Yin, and Ping Zhang. Multimodal risk prediction with physiological signals, medical images and clinical notes. *Heliyon*, 10(5), 2024.

[Wen *et al.*, 2022] Qianlong Wen, Zhongyu Ouyang, Jianfei Zhang, Yiyue Qian, Yanfang Ye, and Chuxu Zhang. Disentangled dynamic heterogeneous graph learning for opioid overdose prediction. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, page 2009–2019, 2022.

[Xu *et al.*, 2020] Da Xu, Chuanwei Ruan, Evren Korpeoglu, Sushant Kumar, and Kannan Achan. Inductive representation learning on temporal graphs. In *ICLR*, 2020.

[Yang *et al.*, 2021] Chaoqi Yang, Cao Xiao, Fenglong Ma, Lucas Glass, and Jimeng Sun. Safedrug: Dual molecular graph encoders for recommending effective and safe drug combinations. In *IJCAI*, pages 3735–3741, 2021.

[Yang *et al.*, 2023a] Chaoqi Yang, Zhenbang Wu, Patrick Jiang, Zhen Lin, Junyi Gao, Benjamin Danek, and Jimeng Sun. PyHealth: A deep learning toolkit for healthcare predictive modeling. In *Proceedings of the 27th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD) 2023*, 2023.

[Yang *et al.*, 2023b] Kai Yang, Yongxin Xu, Peinie Zou, Hongxin Ding, Junfeng Zhao, Yasha Wang, and Bing Xie. Kerprint: Local-global knowledge graph enhanced diagnosis prediction for retrospective and prospective interpretations. *AAAI*, 37(4), 2023.

[Yang *et al.*, 2023c] Zhichao Yang, Avijit Mitra, Weisong Liu, Dan Berlowitz, and Hong Yu. Transformehr: transformer-based encoder-decoder generative model to enhance prediction of disease outcomes using electronic health records. *Nature Communications*, 14:7857, 2023.

[Yin *et al.*, 2022] Changchang Yin, Ruoqi Liu, Jeffrey Caterino, and Ping Zhang. Deconfounding actor-critic network with policy adaptation for dynamic treatment regimes. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, page 2316–2326, 2022.

[Zhang *et al.*, 2017] Yutao Zhang, Robert Chen, Jie Tang, Walter F. Stewart, and Jimeng Sun. Leap: Learning to prescribe effective and safe treatment combinations for multimorbidity. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, page 1315–1324, 2017.