

KG-TREAT: Pre-training for Treatment Effect Estimation by Synergizing Patient Data with Knowledge Graphs

Ruoqi Liu¹, Lingfei Wu², Ping Zhang¹

¹The Ohio State University

²Anytime.AI

liu.7324@osu.edu, lwu@anytime-ai.com, zhang.10631@osu.edu

Abstract

Treatment effect estimation (TEE) is the task of determining the impact of various treatments on patient outcomes. Current TEE methods fall short due to reliance on limited labeled data and challenges posed by sparse and high-dimensional observational patient data. To address the challenges, we introduce a novel pre-training and fine-tuning framework, KG-TREAT, which synergizes large-scale observational patient data with biomedical knowledge graphs (KGs) to enhance TEE. Unlike previous approaches, KG-TREAT constructs *dual-focus* KGs and integrates a deep bi-level attention synergy method for in-depth information fusion, enabling distinct encoding of treatment-covariate and outcome-covariate relationships. KG-TREAT also incorporates two pre-training tasks to ensure a thorough grounding and contextualization of patient data and KGs. Evaluation on four downstream TEE tasks shows KG-TREAT’s superiority over existing methods, with an average improvement of 7% in Area under the ROC Curve (AUC) and 9% in Influence Function-based Precision of Estimating Heterogeneous Effects (IF-PEHE). The effectiveness of our estimated treatment effects is further affirmed by alignment with established randomized trial findings.

Introduction

Treatment effect estimation (TEE), which identifies the causal effects of treatment options on patient outcomes given observational covariates, is a pivotal task in healthcare (Glass et al. 2013). Yet, existing TEE methods (Shalit, Johansson, and Sontag 2017; Shi, Blei, and Veitch 2019; Zhang et al. 2022) are limited in both generalizability and accuracy due to their dependence on small, task-specific datasets that might not fully encompass the complex relationships among covariates, treatments, and outcomes.

To address this, one might consider deploying foundation models (Devlin et al. 2019; Brown et al. 2020; Bommasani et al. 2021) trained on large datasets, to improve generalizability. However, the application of foundation models to TEE is not straightforward. Medical data, often characterized by high-dimensional and sparse medical concepts, continue to pose challenges to these models (Huang, Altosaar, and Ranganath 2019; Rasmy et al. 2021). Even with large-scale datasets, developing a domain-specific understanding

of these medical concepts and identifying potential confounders to reduce estimation bias remains difficult. The sheer volume of data does not necessarily equate to rich, specific instances that the model needs to learn effectively.

Therefore, we turn to biomedical knowledge graphs (KGs) - structured representations of diverse medical concepts and their relations. By synergizing potentially sparse patient data with domain-specific knowledge from KGs, we can derive meaningful insights and identify key confounders for adjustment in TEE.

Despite its potential, synergizing patient data with KGs poses several challenges. Firstly, utilization of the entire KG can introduce noise that is irrelevant to patient data. While recent works have suggested constructing personalized knowledge graphs (PKG) by retrieving relevant medical information for each patient from the entire KG to mitigate noise (Ye et al. 2021; Xu et al. 2023; Yang et al. 2023), these methods fail to distinguish among various types of medical codes, such as treatments, covariates, and outcomes. This issue can lead to bias and spurious correlations in TEE. Secondly, the primary application scenario of existing works is clinical risk prediction (Choi et al. 2017; Ma et al. 2018a), neglecting the encoding of vital causal relationships among covariates, treatments, and outcomes unique to TEE. Additionally, existing methods (Ye et al. 2021; Xu et al. 2023) often incorporate patient data and KGs only at the final prediction stage, leading to superficial combination and inefficient information utilization.

To address these challenges, we propose a novel pre-training and fine-tuning framework for TEE, named KG-TREAT, by synergizing patient data and KGs. Firstly, we address the bias issue by constructing *dual-focus* PKGs, one focusing on the relationship between treatment and covariates (treatment-covariate PKG) and the other focusing on the relationship between outcome and covariates (outcome-covariate PKG). These PKGs explicitly capture and represent the key relationships and dependencies among treatments, outcomes, and covariates, thereby mitigating the risk of spurious correlations. Secondly, we present a deep *bi-level* attention synergy method, namely DIVE. The first level of attention applies a treatment(outcome) attention mechanism to patient data and specific treatment(outcome) information from PKGs, explicitly encoding the complex relationships among covariates, treatments, and outcomes. The

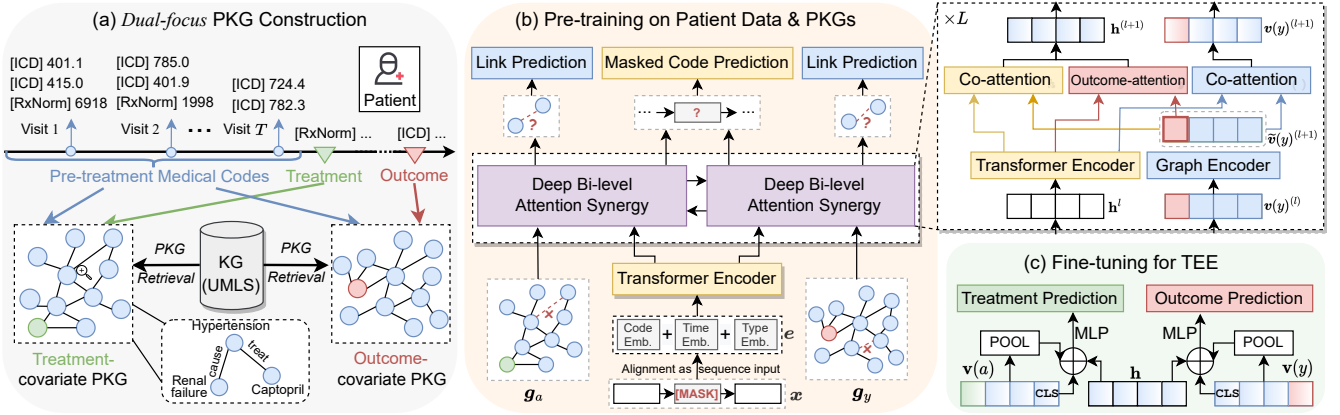


Figure 1: A detailed illustration of KG-TREAT. (a) *Dual-focus* PKGs are constructed by extracting relevant treatment-covariate and outcome-covariate information for each individual patient from KG. (b) The model is pre-trained by synergizing patient data with corresponding PKGs through the proposed deep bi-level attention synergy method. Two pre-training tasks are unified to learn contextualized representations. (c) The pre-trained model is fine-tuned on downstream data for TEE.

second level employs a multi-layer co-attention mechanism to patient data and its corresponding PKGs, ensuring deep information synergy between these modalities.

KG-TREAT is first pre-trained by combining two self-supervised tasks: masked code prediction and link prediction. These tasks ensure that the patient data and KGs are fully grounded and contextualized to each other. The model is then fine-tuned on downstream data for TEE.

Our contributions include:

- We propose KG-TREAT, a novel pre-training and fine-tuning framework that integrates observational patient data with KGs for TEE. This includes the construction of *dual-focus* PKGs and the introduction of a deep bi-level attention synergy method, DIVE.
- We compile a large dataset of 3M patient records from MarketScan Research Databases¹ and KG (300K nodes, 1M edges) from Unified Medical Language System (UMLS) (Bodenreider 2004) for pre-training, and 4 TEE fine-tuning datasets for assessing treatment effectiveness in coronary artery disease (CAD).
- Comprehensive experiments demonstrate the superior performance of KG-TREAT over existing TEE methods. It shows an average improvement of 7% in AUC for outcome prediction and a 9% improvement in IF-PEHE for TEE compared to the best baseline across four tasks.
- Case study shows that the estimated treatment effects align with established randomized controlled trials (RCTs) findings, further demonstrating the effectiveness of our approach in real-world use.

Preliminary

Patient Data. A patient record is a collection of multiple visits, denoted as $\bar{x} = \{x_1, \dots, x_T\}$. Each visit is characterized by a series of medications $m_1, \dots, m_{|\mathcal{M}|} \in \mathcal{M}$ (with $|\mathcal{M}|$ total medication codes), and diagnosis codes

¹<https://www.merative.com/real-world-evidence>

$d_1, \dots, d_{|\mathcal{D}|} \in \mathcal{D}$, (with $|\mathcal{D}|$ total diagnosis codes). A patient’s demographics include age and gender, encoded as categorical and binary values respectively, and are denoted as $c \in \mathcal{C}$. We create a comprehensive medical vocabulary, $\mathcal{W} = \{\mathcal{M}, \mathcal{D}, \mathcal{C}\}$ that includes all these patient attributes. We denote the pre-train data as \mathcal{X} and downstream data as \mathcal{Z} , where $\mathcal{X} \cap \mathcal{Z} = \emptyset$.

Personalized Knowledge Graph. A biomedical KG, which contains extensive relationships among various medical codes (e.g., medications and diagnoses), can be represented as a multi-relational graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} is the set of entity nodes and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{R} \times \mathcal{V}$ is the set of edges. These edges connect nodes in \mathcal{V} via triplets, and \mathcal{R} is the set of relation types. A triplet is denoted as (h, r, t) , $h, t \in \mathcal{V}$, $r \in \mathcal{R}$, represents a relationship within the KG. As an entire KG can be large and contain noises, a personalized KG (PKG), $g = (v, e)$ is considered by extracting a subgraph from the KG, where $v \in \mathcal{V}$ and $e \in \mathcal{E}$.

Treatment Effect Estimation. Given a patient’s visit sequence $\bar{x} = \{x_1, \dots, x_T\}$, demographics c , binary treatment $a \in \{0, 1\}$ (where 1 indicates treated and 0 indicates control status), disease outcome $y \in \{0, 1\}$ (where 1 indicates its presence and 0 indicates its absence), and PKGs g , we aim to estimate treatment effect as $\mathbb{E}[Y(1) - Y(0) | X = \bar{x}, C = c, G = g]$, where $Y(A)$ is the potential outcome if the patient receives treatment A (Rubin 2005). We make three standard assumptions in our TEE analysis: consistency, positivity, and ignorability (see details in Appendix A). These assumptions ensure that the treatment effects estimated as $\mathbb{E}[Y | A = 1, X = \bar{x}, C = c, G = g] - \mathbb{E}[Y | A = 0, X = \bar{x}, C = c, G = g]$ are identifiable.

Method

In this section, we introduce our model (Fig. 1), including three main modules: 1) *Dual-focus* PKG construction, 2) Pre-training on patient data & PKGs, and 3) Fine-tuning for TEE. Algorithm 1 shows the model training procedure.

Data Encoding

As the raw patient data and KGs are not applicable for direct modeling, we need to encode the patient data into dense embeddings and construct *dual-focus* PKGs given the relevant information from individual patients.

Patient Data Encoding. Compared to natural text, patient data presents unique challenges owing to its irregular temporality (i.e., variability in time intervals between patient visits) and hierarchical structures (i.e., a patient record includes multiple visits and each visit includes different types of medical codes). To address these, we propose a comprehensive embedding approach that extends the original BERT (Devlin et al. 2019) embedding by incorporating both the code type and temporal information. For every input medical code, the patient embedding e is obtained as:

$$e = \mathbf{w}_{\text{code}} + \mathbf{t}_{\text{type}} + \mathbf{v}_{\text{visit}} + \mathbf{p}_{\text{physical}} \quad (1)$$

where $\mathbf{w}_{\text{code}} \in \mathbb{R}^{d_{\text{emb}}}$ is the medical code embedding, and $\mathbf{t}_{\text{type}} \in \mathbb{R}^{d_{\text{emb}}}$ is the type embedding. Our input data includes three types: demographics, medication, and diagnosis. The visit time embedding $\mathbf{v}_{\text{visit}} \in \mathbb{R}^{d_{\text{emb}}}$ is the actual time of a visit. The physical time embedding $\mathbf{p}_{\text{physical}} \in \mathbb{R}^{d_{\text{emb}}}$ measures the physical time over a fixed time interval. The code embedding, time embedding, and type embedding are integrated as the input to the patient sequence encoder.

Dual-focus PKG Construction. To facilitate a nuanced understanding of patient data and obtain personalized estimation for TEE, we propose to construct *dual-focus* PKGs that capture diverse medical contexts of treatment-covariate and outcome-covariate relationship, respectively. We first map the medical codes from a patient’s record to their corresponding medical concepts in KG, resulting in an initial set of graph nodes v' . Then we augment the graph with treatment and outcome data. The treatment-covariate PKG includes the mapped treatment concepts added to v' as $v'(a)$, while the outcome-covariate PKG incorporates mapped outcome concepts, added to v' as $v'(y)$. To leverage implicit contextual knowledge, we include k -hop bridge nodes in the final graph node set $v(a)$ and $v(y)$. A k -hop bridge node denotes the entity node positioned within a k -hop path between any pair of linked entities in the node set $v'(a)$ or $v'(y)$. Finally, we establish a link between any pair of entities in $v(a)$ and $v(y)$ if there exists an edge between them. This procedure results in *dual-focus* PKGs as $g_a = (v(a), e(a))$ and $g_y = (v(y), e(y))$.

Pre-training KG-TREAT

The patient embeddings and PKGs are first encoded and then synergized through the proposed deep bi-level attention synergy method. KG-TREAT is pre-trained by unifying two tasks: masked code prediction and link prediction.

Patient Sequence Encoder. Patient visit sequences are encoded using an N-stacked Transformer (Vaswani et al. 2017). Each Transformer encoder has a multi-head self-attention layer followed by a fully-connected feed-forward layer. The patient sequence representations are computed as:

$$\tilde{\mathbf{h}}_{\text{CLS}}^{(l+1)}, \tilde{\mathbf{h}}_1^{(l+1)}, \dots, \tilde{\mathbf{h}}_T^{(l+1)} = f_{\text{seq}}(\mathbf{h}_{\text{CLS}}^{(l)}, \mathbf{h}_1^{(l)}, \dots, \mathbf{h}_T^{(l)}), \quad (2)$$

Algorithm 1: KG-TREAT Pre-training and Fine-tuning

Input: Pre-train data \mathcal{X} , KG \mathcal{G} , downstream data \mathcal{Z}

Output: Pre-trained model f_{θ^*} , treatment effects δ

- 1: Obtain patient data encoding e by Eq. (1);
- 2: Extract *dual-focus* PKGs $g(a), g(y)$ from entire KG;
- 3: Obtain patient representations $\{\tilde{\mathbf{h}}_{\text{CLS}}, \tilde{\mathbf{h}}_1, \dots, \tilde{\mathbf{h}}_T\}$ by Eq. (2);
- 4: Obtain PKG representations $\{\tilde{\mathbf{v}}_{\text{CLS}}, \tilde{\mathbf{v}}_1, \dots, \tilde{\mathbf{v}}_T\}$ by Eq. (3);
- 5: Synergize patient and PKG representations by Eq. (8);
- 6: Pre-train the model by unifying MCP Eq. (10) and LP Eq. (11);
- 7: Initialize the model with parameters θ^* from pre-training;
- 8: Obtain patient representations $\{\mathbf{h}_{\text{CLS}}, \mathbf{h}_1, \dots, \mathbf{h}_T\}$ and PKG representations $\{\mathbf{v}_{\text{CLS}}, \mathbf{v}_1, \dots, \mathbf{v}_T\}$ by Eq. (2), (3);
- 9: Fine-tune the model and estimate effects by Eq. (14), (15);

where $l = 1, \dots, N$ denotes the Transformer layer and the representations in layer $l = 0$ are initialized with the patient embedding e . The term \mathbf{h}_{CLS} is the encoding of a special code that is appended to the patient sequence and acts as the pooling point for prediction. More details of the Transformer architecture are provided in Appendix B.

Graph Encoder. We utilize graph neural networks (GNNs) to encode the PKGs. We initialize node embeddings following existing work (Feng et al. 2020). We transform KG triplets into textual data and feed these sentences into a pre-trained language model, BioLinkBERT (Yasunaga, Leskovec, and Liang 2022), to obtain sentence embeddings. We compute node embeddings by pooling all token outputs of the entity nodes. We encode the treatment-covariate PKG $g(a)$ and outcome-covariate PKG $g(y)$ as follows:

$$\begin{aligned} \tilde{\mathbf{v}}(a)_{\text{CLS}}^{(l+1)}, \dots, \tilde{\mathbf{v}}(a)_I^{(l+1)} &= f_{\text{gnn}}(\mathbf{v}(a)_{\text{CLS}}^{(l)}, \dots, \mathbf{v}(a)_I^{(l)}), \\ \tilde{\mathbf{v}}(y)_{\text{CLS}}^{(l+1)}, \dots, \tilde{\mathbf{v}}(y)_J^{(l+1)} &= f_{\text{gnn}}(\mathbf{v}(y)_{\text{CLS}}^{(l)}, \dots, \mathbf{v}(y)_J^{(l)}), \end{aligned} \quad (3)$$

where $l = 1, \dots, M$ denotes the GNN encoder layer, and \mathbf{v}_{CLS} is the encoding of a special node added to the PKG (with edges connecting to all other nodes) to serve as the pooling point for prediction. The node representations are updated via iterative message passing between neighbors as:

$$\mathbf{v}(a)_i^{(l+1)} = f_v\left(\sum_{s \in \mathcal{N}_i \cup \{i\}} \alpha_{s,i} \mathbf{m}_{si}\right) + \mathbf{v}(a)_i^{(l)}, \quad (4)$$

where \mathcal{N}_i denotes the neighbors of entity node i , $\alpha_{s,i}$ is the attention weight for scaling the message pass, and f_v is a multi-layer perceptron with batch normalization. The message \mathbf{m}_{si} from s to i is computed as follows:

$$\mathbf{m}_{si} = f_m(\mathbf{v}(a)_s^{(l)}, \mathbf{r}_{si}), \quad (5)$$

where \mathbf{r}_{si} is the relation embedding and f_m is a linear transformation. The attention weight $\alpha_{s,i}$, which controls the impact of each neighbor on the current node, is computed as:

$$\begin{aligned} \mathbf{q}_s &= f_q(\mathbf{v}(a)_s^{(l)}), \quad \mathbf{k}_i = f_k(\mathbf{v}(a)_i^{(l)}, \mathbf{r}_{si}), \\ \alpha_{s,i} &= \text{Softmax}(\mathbf{q}_s \mathbf{k}_i^\top / \sqrt{d}), \end{aligned} \quad (6)$$

where f_q and f_k are linear transformations. The outcome-covariate PKG can be encoded similarly as above.

Deep Bi-level Attention Synergy. To address the challenges of shallow synergy and inefficient information utilization in

existing work (Ye et al. 2021), we propose a deep bi-level attention synergy method, DIVE.

The first level of attention handles the complex treatment-covariate and outcome-covariate relationships for bias adjustment and accurate estimation. Given the patient sequence representation $\tilde{\mathbf{h}}_t^{(l)}$ and treatment node representation $\tilde{\mathbf{v}}(a)_a^{(l)}$, we compute the treatment attention weight $\alpha_{a,h,t}$ and treatment-related attention-pooling of patient sequence representation $\hat{\mathbf{h}}^{(l)}$ as follows:

$$\begin{aligned} \alpha_{a,h,t} &= \text{Softmax}(\tilde{\mathbf{v}}(a)_a^{(l)} \tilde{\mathbf{h}}_t^{(l)\top} / \sqrt{d}), \\ \hat{\mathbf{h}}^{(l)} &= \sum_{t=1}^T \alpha_{a,h,t} \tilde{\mathbf{h}}_t^{(l)}. \end{aligned} \quad (7)$$

The second level of attention enables deep synergy of patient data with KGs. We apply a multi-head co-attention (Murahari et al. 2020) across patient sequence and graph representations in multiple hidden layers. Concretely, we obtain synergized patient sequence representations by transforming $\tilde{\mathbf{h}}^{(l)}$ to queries, $\tilde{\mathbf{v}}^{(l)}(a)$ to keys and values. These synergized representations are then concatenated with the treatment-related patient representations derived from Eq. (7), and passed through a multi-layer perceptron f_c to yield bi-level attention synergized patient sequence representations $\mathbf{h}^{(l)}$. This process is formally denoted as:

$$\begin{aligned} \mathbf{h}^{(l)'} &= \text{MHCA}_{h,v(a)}(\tilde{\mathbf{h}}^{(l)}, \tilde{\mathbf{v}}(a)^{(l)}, \tilde{\mathbf{v}}(a)^{(l)}), \\ \mathbf{h}^{(l)} &= f_c[\mathbf{h}^{(l)'}; \hat{\mathbf{h}}^{(l)}], \end{aligned} \quad (8)$$

where $\text{MHCA}_{Q,K}(Q, K, K)$ is the multi-head co-attention applied to Q, K , with Q as query, K as both key and value. We compute the synergized node representations $\mathbf{v}(a)^{(l)}$ by transforming $\tilde{\mathbf{v}}(a)^{(l)}$ to queries, $\tilde{\mathbf{h}}^{(l)}$ to keys and values as:

$$\mathbf{v}(a) = \text{MHCA}_{v(a),h}(\tilde{\mathbf{v}}(a)^{(l)}, \tilde{\mathbf{h}}^{(l)}, \tilde{\mathbf{h}}^{(l)}). \quad (9)$$

The synergized outcome node representations can be obtained similarly. In essence, DIVE introduces a bi-level attention synergy system: one level that adjusts bias by handling relationships among covariates, treatments, and outcomes, and a second level that focuses on deep synergy. These levels work together to facilitate efficient information utilization and overcome the limitations of shallow synergy methods.

Pre-training Tasks. The goal of pre-training is to encourage a thorough grounding and contextualization of patient data and KGs. To approach this, two self-supervised pre-training tasks are adopted: masked code prediction (MCP) and KG link prediction (LP).

MCP predicts the masked medical code with position $j \in \mathcal{J}$ using the representation \mathbf{h}_j . The loss of MCP, $\mathcal{L}_{\text{MCP}}(\theta_m)$, with optimization parameters θ_m , is formulated as:

$$\mathcal{L}_{\text{MCP}}(\theta^m) = - \sum_{j \in \mathcal{J}} \log(P(\mathbf{w}_j | \mathbf{h}_j)), \quad (10)$$

where $P(\mathbf{w}_j | \mathbf{h}_j)$ is the softmax probability of the masked code over all codes in the vocabulary. By using the synergized representations, the patient data are enhanced with external knowledge to predict the masked codes.

LP is widely used in KG representation learning, which aims to distinguish existing (positive) triplets from corrupted (negative) triplets using the representations of entities and relations. Formally, given the representations of a triplet obtained from the graph encoder as $(\mathbf{v}_h, \mathbf{r}, \mathbf{v}_t)$, the pre-training loss for LP is optimized over parameters θ^l as:

$$\begin{aligned} \mathcal{L}_{\text{LP}}(\theta^l) &= \sum_{(h,r,t) \in \mathcal{S}} \left(-\sigma(d(\mathbf{v}_h, \mathbf{r}, \mathbf{v}_t)) + \right. \\ &\quad \left. \sum_{(h',r',t') \in \mathcal{S}'} \sigma(d(\mathbf{v}_{h'}, \mathbf{r}, \mathbf{v}_{t'})) \right) \end{aligned} \quad (11)$$

where $(h', r', t') \in \mathcal{S}'_{(h,r,t)}$ are corrupted triplets in which either the head or tail entity is replaced by a random entity (but not both simultaneously), σ denotes logarithm sigmoid function, and d is the score function such as TransE (Bordes et al. 2013) and DistMult (Yang et al. 2015).

The final pre-training loss is optimized over parameters $\theta = \{\theta^m, \theta^l\}$ by integrating MLP and LP as $\mathcal{L}(\theta) = \mathcal{L}_{\text{MCP}}(\theta^m) + \mathcal{L}_{\text{LP}}(\theta^l)$.

Fine-tuning KG-TREAT for TEE

After pre-training, we fine-tune the model on downstream data for TEE. To mitigate confounding bias, the model is fine-tuned to simultaneously predict the treatment and outcome using shared representations. This strategy discourages reliance on unrelated features and prioritizes confounders for predictions (Shi, Blei, and Veitch 2019).

To elaborate, given downstream patient data and corresponding PKGs, we obtain representations for patients, treatment-covariate PKG, and outcome-covariate PKG. We then predict the treatment using a combination of \mathbf{h}_{CLS} , $\mathbf{v}(a)_{\text{CLS}}$, and an attention-based pooling $\mathbf{v}(a)_{\text{POOL}}$ with query $\mathbf{h}_{[\text{CLS}]}$ as the input to a prediction head f_{ϕ^a} . The loss of treatment prediction is computed as:

$$\begin{aligned} \hat{a} &= f_{\phi^a} \circ f_{\theta^*}([\mathbf{h}_{\text{CLS}}, \mathbf{v}(a)_{\text{CLS}}, \mathbf{v}(a)_{\text{POOL}}]), \\ \mathcal{L}_{\text{T}}(\theta^*, \phi^a) &= \text{BCE}(\hat{a}, a), \end{aligned} \quad (12)$$

where θ^* are the optimized parameters of the pre-train model, BCE denotes binary cross entropy loss. Similarly, we predict the outcome by combining \mathbf{h}_{CLS} , $\mathbf{v}(y)_{\text{CLS}}$, and an attention-based pooling $\mathbf{v}(y)_{\text{POOL}}$ with query $\mathbf{h}_{[\text{CLS}]}$ as the input to a prediction head f_{ϕ^y} . We employ separate heads for treated and control potential outcomes, computing the loss of outcome prediction as:

$$\begin{aligned} \hat{y} &= f_{\phi^y} \circ f_{\theta^*}([\mathbf{h}_{\text{CLS}}, \mathbf{v}(y)_{\text{CLS}}, \mathbf{v}(y)_{\text{POOL}}]), \\ \mathcal{L}_{\text{O}}(\theta^*, \phi^y) &= \text{BCE}(\hat{y}, y). \end{aligned} \quad (13)$$

We jointly optimize both treatment prediction and outcome prediction, computing the final loss as follows:

$$\mathcal{L}_{\text{TEE}}(\theta^*, \phi) = \mathcal{L}_{\text{O}}(\theta^*, \phi^y) + \beta \mathcal{L}_{\text{T}}(\theta^*, \phi^a), \quad (14)$$

where β is a hyper-parameter that controls the influence of treatment prediction. Note that only the observational, or factual, outcomes are used to compute outcome prediction loss, as counterfactual outcomes are unavailable. After fine-tuning the model, we infer the treatment effect δ as the difference between two predicted potential outcomes under the treated and control treatment as:

$$\hat{\delta} = \hat{y}_{a=1} - \hat{y}_{a=0} \quad (15)$$

Experimental Setup

Pre-training Data. We extract patient data from MarketScan Commercial Claims and Encounters (CCA)E database² for those diagnosed with coronary artery disease (CAD), producing a dataset of 2,955,399 patient records and 116,661 unique medical codes. We use UMLS³ as external knowledge. We unify the medical codes in patient data and concepts in the UMLS through standard vocabularies⁴ and extract all relevant relationships from UMLS, resulting in a KG with around 300K nodes and 1M edges

Downstream Tasks. Our goal is to evaluate the effects of two treatments on reducing stroke and myocardial infarction risk for CAD patients, given the patient’s covariates and corresponding PKGs. As the ground truth treatment effects are not available in observational data and RCTs are the gold standard for TEE, we specifically create 4 downstream datasets based on CAD-related RCTs. More details of datasets are provided in Appendix Table A3 and Fig. A2).

Baselines. We compare KG-TREAT with state-of-the-art methods, all trained solely on downstream data.

- **TARNet** (Shalit, Johansson, and Sontag 2017) predicts the potential outcomes based on balanced hidden representations among treated and controlled groups.
- **DragonNet** (Shi, Blei, and Veitch 2019) jointly predicts treatment and outcome via a three-head neural network: one for treatment prediction and two for outcomes.
- **DR-CFR** (Hassanpour and Greiner 2020) predicts the counterfactual outcome by learning disentangled representations that the covariates can be disentangled into three components: only contributing to treatment selection, only contributing to outcome prediction, and both.
- **TNet** (Curth and van der Schaar 2021a) is a deep neural network version of T-learner (Künzel et al. 2019) (i.e., decomposes the TEE into two or more sub-regression/classification problems).
- **SNet** (Curth and van der Schaar 2021a) learns disentangled representations and assumes that the covariates can be disentangled into five components by considering two potential outcomes separately.
- **FlexTENet** (Curth and van der Schaar 2021b) assumes inductive bias for the shared structure of two potential outcomes and adaptively learns what to share between the potential outcome functions.
- **TransTEE** (Zhang et al. 2022) is a Transformer-based TEE model, which encodes the covariates and treatments via Transformer and cross-attention.

Additionally, we consider several variants of the proposed model for comparison:

- *w/o DIVE*: replacing the proposed deep bi-level attention synergy method DIVE with a simple concatenation of patient and graph representations only at the final prediction.
- *w/o KGs*: pre-trained solely on patient data without KGs.

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

³<https://www.nlm.nih.gov/research/umls/index.html>

⁴[nlm.nih.gov/research/umls/sourcereleasedocs/index.html](https://www.nlm.nih.gov/research/umls/sourcereleasedocs/index.html)

- *w/o pre-train*: directly trained on the downstream patient data and KGs using the same model architecture as KG-TREAT.
- *w/o pre-train & KGs*: directly trained on the downstream patient data with the patient sequence encoder only.

Note that we do not compare with existing pre-training models for clinical risk prediction (Li et al. 2020; Rasmy et al. 2021), as these models are not directly applicable to our context. Clinical risk prediction models focus on forecasting the likelihood of a disease based on variable correlations. TEE mainly quantifies the causal impact of a treatment on an outcome, predicting all potential outcomes as treatment effects, an entirely different objective from risk prediction.

Metrics. We assess factual prediction performance using standard classification metrics: Area under the ROC Curve (AUC) and Area under the Precision-Recall Curve (AUPR). We evaluate counterfactual prediction performance using influence function-based precision of estimating heterogeneous effects (IF-PEHE) (Alaa and van der Schaar 2019), which measures the mean squared error between estimated treatment effects and approximated true treatment effects. Additional details of this metric are in Appendix C.

Implementation Details. The patient sequence encoder uses the BERT-base architecture (Devlin et al. 2019). The PKGs are retrieved with 2-hop bridge nodes, with a maximum node limit of 200. The downstream data is randomly split into training, validation, and test sets with percentages of 90%, 5%, and 5% respectively. All results are reported on the test sets. More implementation details⁵ including parameter tuning and setup are mentioned in Appendix C.

Results

Quantitative Analysis

We quantitatively compare KG-TREAT with state-of-the-art methods in terms of factual outcome prediction and TEE. Table 1 presents the results on 4 downstream datasets. Our key findings include:

- KG-TREAT significantly outperforms the best baseline method, demonstrating an average improvement of 7% in AUC, 12% in AUPR, and 9% in IF-PEHE. This validates the effectiveness of our pre-training approach, which synergizes patient data with KGs.
- The variant of KG-TREAT without the deep bi-level attention synergy method, *w/o DIVE*, shows a drop in performance. This highlights the effectiveness of DIVE in modeling the relationships among covariates, treatments, and outcomes, and in synergizing patient data with KGs.
- Pre-training has a more significant impact on model performance than KGs, as indicated by the greater performance decline in the *w/o pre-train* scenario compared to the *w/o KGs* scenario. And *w/o pre-train & KGs* yields the worst performance among all the model variants.

⁵Code: <https://github.com/ruoqi-liu/KG-TREAT>

Method	Rivaroxaban v.s. Aspirin			Valsartan v.s. Ramipril		
	AUC \uparrow	AUPR \uparrow	IF-PEHE \downarrow	AUC \uparrow	AUPR \uparrow	IF-PEHE \downarrow
TARNet	0.746	0.404	0.277	0.743	0.333	0.289
DragonNet	0.761	0.424	0.261	0.742	0.337	0.271
DR-CFR	0.764	0.426	0.259	0.747	0.341	0.276
TNet	0.726	0.401	0.321	0.730	0.322	0.299
SNet	0.761	0.430	0.254	0.747	0.354	0.268
FlexTENet	0.729	0.403	0.301	0.739	0.328	0.285
TransTEE	0.751	0.411	0.272	0.753	0.379	0.264
KG-TREAT	0.828	0.556	0.171	0.858	0.526	0.149
<i>w/o DIVE</i>	0.811	0.522	0.202	0.829	0.495	0.160
<i>w/o KGs</i>	0.805	0.518	0.219	0.813	0.481	0.165
<i>w/o pre-train</i>	0.786	0.488	0.231	0.778	0.373	0.189
<i>w/o pre-train & KGs</i>	0.769	0.470	0.239	0.749	0.371	0.198
Method	Ticagrelor v.s. Aspirin			Apixaban v.s. Warfarin		
	AUC \uparrow	AUPR \uparrow	IF-PEHE \downarrow	AUC \uparrow	AUPR \uparrow	IF-PEHE \downarrow
TARNet	0.755	0.460	0.282	0.760	0.515	0.325
DragonNet	0.762	0.464	0.289	0.766	0.534	0.309
DR-CFR	0.764	0.460	0.278	0.769	0.526	0.287
TNet	0.741	0.433	0.311	0.753	0.509	0.330
SNet	0.765	0.465	0.265	0.769	0.527	0.273
FlexTENet	0.750	0.452	0.313	0.756	0.512	0.324
TransTEE	0.770	0.471	0.255	0.803	0.534	0.267
KG-TREAT	0.851	0.609	0.160	0.843	0.639	0.191
<i>w/o DIVE</i>	0.839	0.571	0.176	0.830	0.611	0.210
<i>w/o KGs</i>	0.830	0.552	0.181	0.829	0.605	0.217
<i>w/o pre-train</i>	0.815	0.524	0.200	0.807	0.569	0.241
<i>w/o pre-train & KGs</i>	0.808	0.497	0.213	0.801	0.550	0.247

Table 1: Comparison with state-of-the-art methods on 4 downstream datasets. DIVE is our proposed deep bi-level attention method for synergizing patient data with KGs. The results are averaged over 20 random runs.

Target v.s. Compared	Estimated Effect	P value	Model Conclusion	RCT Conclusion
Rivaroxaban v.s. Aspirin	[-0.010, 0.009]	0.952	No significant difference	No significant difference (Anand et al. 2018)
Valsartan v.s. Ramipril	[-0.015, 0.007]	0.564	No significant difference	No significant difference (Pfeffer et al. 2021)
Ticagrelor v.s. Aspirin	[-0.006, 0.021]	0.436	No significant difference	No significant difference (Sandner et al. 2020)
Apixaban v.s. Warfarin	[-0.006, -0.001]	0.001	A. is more effective than W.	A. is more effective than W. (Granger et al. 2011)

Table 2: Comparison of the estimated treatment effects with corresponding ground truth RCT. The estimated effects are shown in 95% confidence intervals (CI) under 20 bootstrap runs. The RCT conclusions are obtained from published articles.

Qualitative Analysis

Besides the quantitative analysis, we demonstrate the model’s ability in facilitating randomized controlled trials (RCTs) by providing an accurate estimation of treatment effects and generating consistent conclusions. Additionally, we show that KGs help improve performance by identifying a more comprehensive set of confounders for adjustment.

Validation with RCT Conclusion. We compare the estimated treatment effects to the corresponding RCT results. First, we compute the average treatment effects as the mean of the differences between the treated outcomes and controlled outcomes (Hernán 2004). Then, we test the significance of the estimated effects against zero using a T-test with

significance level $\alpha = 0.05$ for model conclusion generation. Table 2 shows that model-generated conclusions align with their corresponding RCT conclusions, suggesting that KG-TREAT can be served as an effective computational tool to emulate RCTs using large-scale patient data and KGs. A thorough comparison of our method with all the baseline methods is shown in Appendix Table A7.

Attention Visualization. We use case studies to illustrate how our model identifies potential confounders from patient data and KGs for adjusting bias and accurate estimation. We visualize the model attention weights of each PKG in Fig. 2 and observe that KG-TREAT successfully identifies key confounders as medical codes with high attention weights from

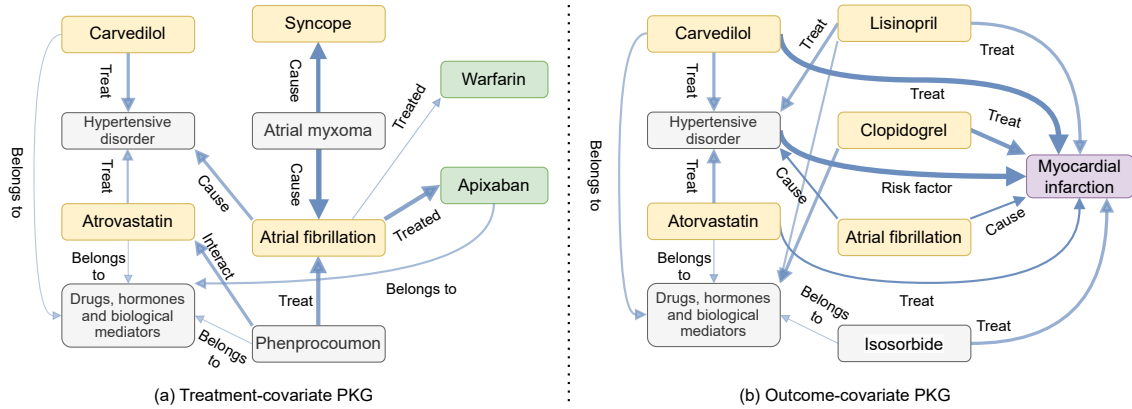


Figure 2: Visualization of the graph attention weights for (a) treatment-covariate PKG and (b) outcome-covariate PKG. The patient is from “Apixaban v.s. Warfarin” dataset. Higher attention weights are denoted as thicker and darker edges in the graph. The extra 2-hop bridge nodes are in gray color to distinguish them from the initial set of nodes in yellow color.

Method		Rivaroxaban v.s. Aspirin		Valsartan v.s. Ramipril		Ticagrelor v.s. Aspirin		Apixaban v.s. Warfarin	
		AUC \uparrow	IF-PEHE \downarrow	AUC \uparrow	IF-PEHE \downarrow	AUC \uparrow	IF-PEHE \downarrow	AUC \uparrow	IF-PEHE \downarrow
Pre-train Task	MCP only	0.815	0.204	0.840	0.160	0.840	0.169	0.831	0.200
	LP only	0.791	0.225	0.813	0.174	0.816	0.186	0.812	0.232
Score Function	RotatE	0.825	0.176	0.851	0.152	0.848	0.164	0.838	0.198
	TransE	0.823	0.177	0.851	0.155	0.846	0.166	0.832	0.203
KG-TREAT		0.828	0.171	0.858	0.149	0.851	0.160	0.843	0.191

Table 3: Ablation study results of using different pre-train tasks and score functions in link prediction. KG-TREAT adopts both MCP and LP as pre-training tasks and DistMult as score function.

PKGs. For example, the common medical codes, such as “Carvedilol”, “Hypertensive disorder”, “Atorvastatin”, are identified as the potential confounders and also mentioned in related literature (Stolk et al. 2017; Yusuf et al. 2004). Notably, with the help of external knowledge, our model can recover potential confounders that are not observed in the patient data. For example, “Hypertensive disorder” is a potentially missing confounding factor added to PKG through 2-hop bridge node searching. This result indicates that solely relying on the patient data may fail to recognize a more comprehensive set of confounders for adjustment.

Ablation Studies

In the above experiments, we find that pre-training plays a critical role in model performance. Therefore, we focus on important model choices in pre-training tasks. An additional ablation study of the influence of downstream data size on model performance is provided in Appendix Fig. A3.

Pre-training Tasks. We analyze the impact of pre-training tasks by excluding the LP and MCP tasks. As shown in Table 3, unifying both MCP and LP tasks (KG-TREAT) yields the best performance on 4 downstream datasets. The MCP task plays a crucial role in pre-training, with its exclusion leading to larger performance drops (3% of AUC and 4% of IF-PEHE) than the LP task (1% of AUC and IF-PEHE). This demonstrates the importance of integrating both tasks

to jointly learn from both data modalities.

Link Prediction Head Choice. We evaluate various score functions for the LP task. As shown in Table 3, DistMult, which is adopted in our model, offers better performance than other score functions. While, the differences among different functions are not significant, indicating that the pre-training model on large-scale data is not particularly sensitive to the scoring function selection.

Related Work

Deep learning for TEE. Deep learning has been extensively used for TEE and achieved improved performance than classical linear methods due to its flexibility of modeling non-linearity (Shalit, Johansson, and Sontag 2017; Shi, Blei, and Veitch 2019; Hassanpour and Greiner 2020; Curth and van der Schaar 2021a,b). For instance, TARNet (Shalit, Johansson, and Sontag 2017) employs shared representations to simultaneously predict two potential outcomes, encouraging the similarity between treated and control distributions. SNet (Curth and van der Schaar 2021a) learns disentangled representations for flexible information sharing among treatment prediction and outcome prediction. Recent Transformer-based models have been introduced as backbones for TEE to help handle various data modalities (e.g., graph, text, etc.) (Zhang et al. 2022; Guo et al. 2021). However, these existing methods are mainly trained on small-

scale, task-specific labeled data. This may limit model performance due to insufficient learning of the complex relationships among covariates, treatments, and outcomes.

Knowledge Integration in Healthcare. Various works in healthcare incorporate biomedical KGs to enrich patient data (Choi et al. 2017; Ma et al. 2018a,b). Recent works have shown that personalized KGs (PKG), constructed by retrieving relevant medical features of individual patients from KG, can encode personalized information and mitigate noise of the entire KG (Ye et al. 2021; Xu et al. 2023; Yang et al. 2023). However, these approaches often fail to distinguish among different types of medical codes and encode relationships among covariates, treatments, and outcomes. This can lead to confounding bias and spurious correlations in TEE. Additionally, existing methods often integrate patient data and KGs only in the final prediction, resulting in shallow combination and inefficient information utilization.

Pre-training in Healthcare. The foundation models have been successfully applied to various domains including healthcare and patient data (Huang, Altosaar, and Ranganath 2019; Li et al. 2020; Rasmy et al. 2021). These methods typically convert a patient’s medical records into a sequence of tokens for pre-training, followed by fine-tuning for healthcare-related downstream tasks such as clinical risk prediction. However, learning deep, domain-specific representations of complex medical features from sparse patient data can be challenging, even with large-scale data. Furthermore, existing methods often fail to handle complex relationships among covariates, treatments, and outcomes, potentially leading to biased estimation. To the best of our knowledge, ours is the first pre-training model by synergizing both observational patient data and KGs.

Conclusion

In this paper, we propose KG-TREAT, a pre-training and fine-tuning framework for TEE by synergizing patient data with KGs. We construct *dual-focus* personalized KGs that incorporate key relationships among covariates, treatments, and outcomes for addressing potential bias in TEE. We propose a novel synergy method (DIVE) to achieve deep information exchange between patient data and PKGs, and encourage complex relationship encoding for TEE. We jointly pre-train the model via two self-supervised tasks and fine-tune it on downstream TEE datasets. Thorough experiments on real-world patient data show the effectiveness of KG-TREAT compared to state-of-the-art methods. We further demonstrate the estimated treatment effects are well consistent with corresponding published RCTs.

Acknowledgments

This work was funded in part by the National Institute of General Medical Sciences (NIGMS) of NIH under award number R01GM141279.

References

Alaa, A. M.; and van der Schaar, M. 2019. Validating Causal Inference Models via Influence Functions. In Chaudhuri, K.;

and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, 191–201.

Anand, S. S.; Bosch, J.; Eikelboom, J. W.; Connolly, S. J.; Diaz, R.; Widimsky, P.; Aboyans, V.; Alings, M.; Kakkar, A. K.; Keltai, K.; et al. 2018. Rivaroxaban with or without aspirin in patients with stable peripheral or carotid artery disease: an international, randomised, double-blind, placebo-controlled trial. *The Lancet*, 391(10117): 219–229.

Bodenreider, O. 2004. The unified medical language system (UMLS): integrating biomedical terminology. *Nucleic acids research*, 32(suppl_1): D267–D270.

Bommasani, R.; Hudson, D. A.; Adeli, E.; Altman, R.; Arora, S.; von Arx, S.; Bernstein, M. S.; Bohg, J.; Bosselet, A.; Brunskill, E.; et al. 2021. On the opportunities and risks of foundation models. *ArXiv preprint*, abs/2108.07258.

Bordes, A.; Usunier, N.; García-Durán, A.; Weston, J.; and Yakhnenko, O. 2013. Translating Embeddings for Modeling Multi-relational Data. In Burges, C. J. C.; Bottou, L.; Ghahramani, Z.; and Weinberger, K. Q., eds., *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, 2787–2795.

Brown, T. B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; Krueger, G.; Henighan, T.; Child, R.; Ramesh, A.; Ziegler, D. M.; Wu, J.; Winter, C.; Hesse, C.; Chen, M.; Sigler, E.; Litwin, M.; Gray, S.; Chess, B.; Clark, J.; Berner, C.; McCandlish, S.; Radford, A.; Sutskever, I.; and Amodei, D. 2020. Language Models are Few-Shot Learners. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.

Chen, T.; and Guestrin, C. 2016. XGBoost: A Scalable Tree Boosting System. In Krishnapuram, B.; Shah, M.; Smola, A. J.; Aggarwal, C. C.; Shen, D.; and Rastogi, R., eds., *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016*, 785–794.

Choi, E.; Bahadori, M. T.; Song, L.; Stewart, W. F.; and Sun, J. 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, Halifax, NS, Canada, August 13 - 17, 2017*, 787–795.

Curth, A.; and van der Schaar, M. 2021a. Nonparametric Estimation of Heterogeneous Treatment Effects: From Theory to Learning Algorithms. In Banerjee, A.; and Fukumizu, K., eds., *The 24th International Conference on Artificial Intelligence and Statistics, AISTATS 2021, April 13-15, 2021, Virtual Event*, volume 130 of *Proceedings of Machine Learning Research*, 1810–1818.

- Curth, A.; and van der Schaar, M. 2021b. On Inductive Biases for Heterogeneous Treatment Effect Estimation. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y. N.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, 15883–15894.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186.
- Feng, Y.; Chen, X.; Lin, B. Y.; Wang, P.; Yan, J.; and Ren, X. 2020. Scalable Multi-Hop Relational Reasoning for Knowledge-Aware Question Answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1295–1309.
- Glass, T. A.; Goodman, S. N.; Hernán, M. A.; and Samet, J. M. 2013. Causal inference in public health. *Annual review of public health*, 34: 61–75.
- Granger, C. B.; Alexander, J. H.; McMurray, J. J.; Lopes, R. D.; Hylek, E. M.; Hanna, M.; Al-Khalidi, H. R.; Ansell, J.; Atar, D.; Avezum, A.; et al. 2011. Apixaban versus warfarin in patients with atrial fibrillation. *New England Journal of Medicine*, 365(11): 981–992.
- Guo, Z.; Zheng, S.; Liu, Z.; Yan, K.; and Zhu, Z. 2021. CETransformer: Casual Effect Estimation via Transformer Based Representation Learning. In *Chinese Conference on Pattern Recognition and Computer Vision (PRCV)*, 524–535. Springer.
- Hassanpour, N.; and Greiner, R. 2020. Learning Disentangled Representations for Counterfactual Regression. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*.
- Hernán, M. A. 2004. A definition of causal effect for epidemiological research. *Journal of Epidemiology & Community Health*, 58(4): 265–271.
- Huang, K.; Altosaar, J.; and Ranganath, R. 2019. Clinicalbert: Modeling clinical notes and predicting hospital readmission. *ArXiv preprint*, abs/1904.05342.
- Künzel, S. R.; Sekhon, J. S.; Bickel, P. J.; and Yu, B. 2019. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the national academy of sciences*, 116(10): 4156–4165.
- Li, Y.; Rao, S.; Solares, J. R. A.; Hassaine, A.; Ramakrishnan, R.; Canoy, D.; Zhu, Y.; Rahimi, K.; and Salimi-Khorshidi, G. 2020. BEHRT: transformer for electronic health records. *Scientific reports*, 10(1): 7155.
- Ma, F.; Gao, J.; Suo, Q.; You, Q.; Zhou, J.; and Zhang, A. 2018a. Risk Prediction on Electronic Health Records with Prior Medical Knowledge. In Guo, Y.; and Farooq, F., eds., *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2018, London, UK, August 19-23, 2018*, 1910–1919.
- Ma, F.; You, Q.; Xiao, H.; Chitta, R.; Zhou, J.; and Gao, J. 2018b. KAME: Knowledge-based Attention Model for Diagnosis Prediction in Healthcare. In Cuzzocrea, A.; Allan, J.; Paton, N. W.; Srivastava, D.; Agrawal, R.; Broder, A. Z.; Zaki, M. J.; Candan, K. S.; Labrinidis, A.; Schuster, A.; and Wang, H., eds., *Proceedings of the 27th ACM International Conference on Information and Knowledge Management, CIKM 2018, Torino, Italy, October 22-26, 2018*, 743–752.
- Murahari, V.; Batra, D.; Parikh, D.; and Das, A. 2020. Large-scale pretraining for visual dialog: A simple state-of-the-art baseline. In *European Conference on Computer Vision*, 336–352. Springer.
- Pfeffer, M. A.; Claggett, B.; Lewis, E. F.; Granger, C. B.; Køber, L.; Maggioni, A. P.; Mann, D. L.; McMurray, J. J.; Rouleau, J.-L.; Solomon, S. D.; et al. 2021. Angiotensin receptor–neprilysin inhibition in acute myocardial infarction. *New England Journal of Medicine*, 385(20): 1845–1855.
- Rasmy, L.; Xiang, Y.; Xie, Z.; Tao, C.; and Zhi, D. 2021. Med-BERT: pretrained contextualized embeddings on large-scale structured electronic health records for disease prediction. *NPJ digital medicine*, 4(1): 1–13.
- Rubin, D. B. 2005. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, 100(469): 322–331.
- Sandner, S. E.; Schunkert, H.; Kastrati, A.; Wiedemann, D.; Misfeld, M.; Boening, A.; Tebbe, U.; Nowak, B.; Stritzke, J.; Laufer, G.; et al. 2020. Ticagrelor monotherapy versus aspirin in patients undergoing multiple arterial or single arterial coronary artery bypass grafting: insights from the TiCAB trial. *European Journal of Cardio-Thoracic Surgery*, 57(4): 732–739.
- Shalit, U.; Johansson, F. D.; and Sontag, D. A. 2017. Estimating individual treatment effect: generalization bounds and algorithms. In Precup, D.; and Teh, Y. W., eds., *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, 3076–3085.
- Shi, C.; Blei, D. M.; and Veitch, V. 2019. Adapting Neural Networks for the Estimation of Treatment Effects. In Wallach, H. M.; Larochelle, H.; Beygelzimer, A.; d’Alché-Buc, F.; Fox, E. B.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, 2503–2513.
- Stolk, L. M.; de Vries, F.; Ebbelaar, C.; de Boer, A.; Schalekamp, T.; Souverein, P.; ten Cate-Hoek, A.; and Burden, A. M. 2017. Risk of myocardial infarction in patients with atrial fibrillation using vitamin K antagonists, aspirin or direct acting oral anticoagulants. *British journal of clinical pharmacology*, 83(8): 1835–1843.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In Guyon, I.; von Luxburg, U.; Bengio, S.; Wallach, H. M.; Fergus, R.; Vishwanathan, S.

V. N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, 5998–6008.

Xu, Y.; Chu, X.; Yang, K.; Wang, Z.; Zou, P.; Ding, H.; Zhao, J.; Wang, Y.; and Xie, B. 2023. SeqCare: Sequential Training with External Medical Knowledge Graph for Diagnosis Prediction in Healthcare Data. In *Proceedings of the ACM Web Conference 2023*, 2819–2830.

Yang, B.; Yih, W.; He, X.; Gao, J.; and Deng, L. 2015. Embedding Entities and Relations for Learning and Inference in Knowledge Bases. In Bengio, Y.; and LeCun, Y., eds., *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.

Yang, K.; Xu, Y.; Zou, P.; Ding, H.; Zhao, J.; Wang, Y.; and Xie, B. 2023. KerPrint: Local-Global Knowledge Graph Enhanced Diagnosis Prediction for Retrospective and Prospective Interpretations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 5357–5365.

Yasunaga, M.; Leskovec, J.; and Liang, P. 2022. LinkBERT: Pretraining Language Models with Document Links. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 8003–8016.

Ye, M.; Cui, S.; Wang, Y.; Luo, J.; Xiao, C.; and Ma, F. 2021. Medpath: Augmenting health risk prediction via medical knowledge paths. In *Proceedings of the Web Conference 2021*, 1397–1409.

Yusuf, S.; Hawken, S.; Ôunpuu, S.; Dans, T.; Avezum, A.; Lanas, F.; McQueen, M.; Budaj, A.; Pais, P.; Varigos, J.; et al. 2004. Effect of potentially modifiable risk factors associated with myocardial infarction in 52 countries (the INTERHEART study): case-control study. *The lancet*, 364(9438): 937–952.

Zhang, Y.-F.; Zhang, H.; Lipton, Z. C.; Li, L. E.; and Xing, E. P. 2022. Can Transformers be Strong Treatment Effect Estimators? *ArXiv preprint*, abs/2202.01336.