GraphITE: Estimating Individual Effects of Graph-structured Treatments

Shonosuke Harada
sh1108@ml.ist.i.kyoto-u.ac.jp
Kyoto University

Hisashi Kashima
kashima@i.kyoto-u.ac.jp
Kyoto University

ABSTRACT

Outcome estimation of treatments for individual targets is a crucial foundation for decision making based on causal relations. Most of the existing outcome estimation methods deal with binary or multiple-choice treatments; however, in some applications, the number of interventions can be very large, while the treatments themselves have rich information. In this study, we consider one important instance of such cases, that is, the outcome estimation problem of graph-structured treatments such as drugs. Due to the large number of possible interventions, the counterfactual nature of observational data, which appears in conventional treatment effect estimation, becomes a more serious issue in this problem. Our proposed method GraphITE (pronounced ‘graphite’) obtains the representations of the graph-structured treatments using graph neural networks, and also mitigates the observation biases by using the HSIC regularization that increases the independence of the representations of the targets and the treatments. In contrast with the existing methods, which cannot deal with “zero-shot” treatments that are not included in observational data, GraphITE can efficiently handle them thanks to its capability of incorporating graph-structured treatments. The experiments using the two real-world datasets show GraphITE outperforms baselines especially in cases with a large number of treatments.

CCS CONCEPTS

• Computing methodologies → Machine learning approaches.

KEYWORDS

causal inference, treatment effect estimation, graph neural networks

ACM Reference Format:

Shonosuke Harada and Hisashi Kashima. 2021. GraphITE: Estimating Individual Effects of Graph-structured Treatments. In *Proceedings of the 30th ACM International Conference on Information and Knowledge Management (CIKM '21)*, November 1–5, 2021, Virtual Event, QLD, Australia. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3459637.3482349>

1 INTRODUCTION

Estimating causal effects of treatments for individual targets, which is often referred to as individual treatment effect (ITE), is an important foundation for efficient decision making based on observational

Figure 1: Individual effect estimation problem of graph-structured treatments. The possible treatments, i.e., drugs, are associated with graphs representing their molecular structures. In observational data, only one treatment is applied to the target individual; consequently, only the factual outcome is observed, while the counterfactual outcomes for the other treatments are not. Our goal is to predict the outcomes of all treatments for future targets.

data. The scope of applications of causal inference ranges across a wide range of fields, including medicine, education, and economic policy. The main difficulties in estimating ITE are (i) the counterfactual nature of observational data, that is, only the outcome of an actual treatment is observed and (ii) biases in observational data due to biases in past treatment decisions. To address these difficulties, various statistical techniques have been developed, including matching [35], inverse propensity score weighting [34], instrumental variable methods [3], as well as more modern representation learning approaches [22, 42].

Most previous studies dealt with a binary or relatively small number of treatments. However, in some scenarios, the number of treatments can be considerably larger. For example, when modeling drug effects on target cells, the number of candidate drugs (i.e., treatments) can be huge, while the number of observations per drug can be quite small due to the high cost of clinical trials [9, 20]. Each drug is composed of a bunch of atoms, such as carbon, oxygen, and nitrogen, and the number of drugs composed by the combination of atoms are substantially huge. A similar situation can also occur in online advertisements [38]. This scarcity of data exacerbates the aforementioned problems. More seriously, some treatments that never appeared in the observational data, such as new drugs or new ads, may appear for the first time during the test phase. Despite the significant importance of estimating unobserved treatment effects in various applications, existing methods are not capable of dealing with such “zero-shot” treatment effect estimation.

In this study, we consider the individual treatment effect estimation problem with a large number of treatments, for which there is no definitive existing solution due to extremely sparse observations.

This work is licensed under a Creative Commons Attribution International 4.0 License.

CIKM '21, November 1–5, 2021, Virtual Event, Australia.

© 2021 Copyright is held by the owner/author(s).

ACM ISBN 978-1-4503-8446-9/21/11.

<https://doi.org/10.1145/3459637.3482349>

To solve the problem, we focus on auxiliary information that accompanies the treatments. Such auxiliary information is sometimes available in applications. In the drug effect example, each drug is a chemical compound with its own molecular structure, which can be represented as a graph (Figure 1), and it is expected to take advantage of structural patterns contained in the graph structure. The rich structural information of graphs allows us to transfer useful information for predicting outcomes from treatments with many observations to those with less observations, even to “zero-shot” treatments that have not been seen before. Therefore, our challenge in this paper is to incorporate the rich graph structure information of treatments into our treatment effect estimation model, and provide an effective learning method to mitigate biases in sparse observational data.

We propose GraphITE (pronounced “graphite”), which is an outcome prediction model for graph-structured treatments based on biased observational data. It is built upon the recent significant advances in learning representations using graph neural networks (GNNs) [10, 25]. Bias mitigation with a large (possibly infinite) number of treatments is another issue because most existing frameworks [38, 41, 42] are not designed for such cases. To reduce the treatment selection bias depending on the individual target, GraphITE finds representations of the target and treatment that are as independent of each other as possible. This is achieved by Hilbert-Schmidt Independence Criterion (HSIC) regularization, which was recently proposed by Lopez et al. (2020); we extend their framework to exploit the treatment features extracted by GNNs and give theoretical justification on how reducing biases over the representation space extracted from graph space leads to unbiased results. Our formulation makes it possible to reduce the selection bias caused by complex graph structure information, even for the zero-shot treatments that cannot be handled by existing frameworks.

We conduct experiments on two real datasets: one with a relatively small number of treatments and one with over 100 treatments. The results show that the graph structures contribute to improving the predictive performance and that the HSIC regularization is robust to the presence of selection bias.

The contributions of this study can be summarized as follows:

- We propose GraphITE, an outcome prediction model for graph-structured treatments, that can exploit auxiliary information of treatments to deal with a large number of treatments and “zero-shot” treatments.
- In order to train GraphITE from biased observational data, we extend HSIC regularization to cases where treatments have features, and give theoretical justification on how HSIC regularization making use of representations extracted from graph space contributes to mitigating biases.
- Experiments on two real-world datasets empirically demonstrate the benefits of GraphITE for biased observational data and zero-shot treatments.

2 RELATED WORK

Treatment effect estimation. Treatment effect estimation is a practically important task and has been widely studied in various fields ranging from healthcare [8] and economy [27] to education [57].

One of the typical solutions is the matching method [1, 35], which compares the outcomes for pairs with similar covariates but different treatments. The propensity score, which is the probability of a target individual receiving a treatment, is introduced to mitigate the curse of dimensionality and selection bias [34]. Tree-based methods, such as Causal Forest [46] and BART [6, 19], have also been proposed and shown promising performances.

Recently, representation learning based on deep neural networks has been successfully applied to treatment effect estimation and outperformed traditional methods [22, 42, 52]. They encourage the representations of treatment and control groups to get closer to each other to reduce selection bias. In addition, confounding variables are extracted by an additional neural network that predicts treatment assignments [43]. Generative adversarial neural networks (GAN) [11] were also successfully applied to ITE estimation [5, 53]; their key idea is to train a predictive model (i.e., a generator) whose outcomes are difficult to distinguish between factual outcomes and counterfactual outcomes.

Most previous studies have focused on binary treatments, and extensions to multiple types of treatments, especially high numbers, are key research directions. There have been several approaches designed for multiple treatments [6, 41, 46]; however, most of them are limited to a relatively small number of treatments, making it difficult to consider more than a few dozen treatments. Saini et al. [38] whose motivation was somewhat similar to ours, considered combinatorial treatments; however, their focus was on a large number of combinations made from a small number of treatments, whereas we focus on many single treatments with the help of information on the treatments.

Extensions to real-valued treatments are also important for real-world applications, such as estimation of appropriate drug dosages [21, 40]. Wang et al. [47] proposed an interesting approach to learn input representations that cannot distinguish real-valued domains. Lopez et al. [30] considered total-ordered treatment spaces. They proposed HSIC regularization for dealing with biased observational data; the theory of our proposed GraphITE is based on their theoretical framework, but we extend the implications of their framework to representation learning of treatments with rich features.

Graph neural networks. Graph-structured data is one of the most popular data structures and has been widely employed in various domains such as social network analysis, citation analysis, and chemical informatics. The GNN is one of the most successful deep neural network architectures owing to the practical importance of graph-structured data, and it has significantly improved the performance on various graph-structured data analysis tasks, such as node classification [25], graph classification [7, 10], and link prediction [56], beyond conventional methods [14, 18]. In the field of chemo-informatics, GNNs have particularly flourished and played an important role in predicting molecule properties [7, 10, 39], finding interactions between chemical objects [15], and generating desirable and unique molecules [54, 55]. GraphITE also relies on their powerful ability to extract features from graph-structured treatments.

Theoretical analysis of the expressive power of GNNs has been of great interest to researchers, for example, in their invertibility [29,

32, 49, 55]. GraphITE theoretically requires this property although it does not hold for most practical GNNs; however, a non-invertible GNN shows satisfactory performance in practice, as shown in the experimental section.

Recently, several studies have considered causal inference in *graph-structured input domains* [2, 13, 16, 31, 45], where the input space has a graph structure representing proximal relations among target individuals. However, to the best of our knowledge, no study has explored treatment effect estimation with *graph-structured treatments*, which is at the intersection of the above two topics of practical importance.

3 PROBLEM DEFINITION

We consider the problem of estimating the outcomes of treatments with graph structures from biased observational data. Let $\mathcal{D} = \{(x_i, t_i, y_i^{t_i})\}_{i=1}^N \in \mathcal{X} \times \mathcal{T} \times \mathcal{Y}$ be a biased observational dataset, where $x_i \in \mathcal{X}$ is the covariate vector of the i -th target individual, $t_i \in \mathcal{T}$ is the treatment performed on the target individual, and $y_i^{t_i}$ is the outcome.¹ We assume the covariate space $\mathcal{X} = \mathbb{R}^D$, treatment space $\mathcal{T} = \{1, 2, \dots, |\mathcal{T}|\}$, and outcome space $\mathcal{Y} = \mathbb{R}$. In addition to $\mathcal{D}$, we assume each treatment $j \in \mathcal{T}$ is associated with a graph $G_j = (\mathcal{V}_j, \mathcal{E}_j)$, where $\mathcal{V}_j$ denotes a set of nodes and $\mathcal{E}_j \subseteq \mathcal{V}_j \times \mathcal{V}_j$ denotes a set of edges. We denote the set of the treatment graphs by $\mathcal{G} = \{G_j\}_{j=1}^{|\mathcal{T}|}$. Our goal is to, given $\mathcal{D}$ and $\mathcal{G}$, estimate an outcome prediction function $f : \mathcal{X} \times \mathcal{T} \rightarrow \mathcal{Y}$.

Figure 1 illustrates our problem setting in the context of medical treatments. In the observational data, there are multiple drugs that could be applied to the target individual, where each drug corresponds to a treatment and is associated with a graph representing its molecular structure. Only the outcome $y_i^{t_i}$ for the actual prescribed drug t_i (i.e., the factual outcome) is observed, and those for the other drugs (i.e., counterfactual outcomes) are not observed. Because a doctor prescribes a drug based on the condition of the target patient x_i , there is a bias in the choice of t_i in the observational data.

A potential difficulty in our problem is that the number of treatments can be large, say $|\mathcal{T}| > 100$; it is clear that this can cause a data scarcity issue. In our scenario, graphs are available as auxiliary information for the treatments, which potentially help in dealing with such a large number of treatments.

Following the existing work, we make the typical assumptions in the Rubin-Neyman framework [36]: (i) Stable unit treatment value; the outcome of each instance is not affected by the treatments assigned to other instances. (ii) Unconfoundedness; the treatment assignment to an instance is independent of the outcome given the covariates (i.e., the confounder variables). (iii) Overlap; each instance has a positive probability of treatment assignment, i.e., $\forall x, t, p(x, t) > 0$.

4 GRAPHITE

Previous studies on individual treatment effect estimation have not considered rich information associated with treatments, which in

Figure 2: The model architecture of GraphITE. A target individual x and graph-structured treatment G_t are the inputs. The ϕ and ψ map them to the low-dimensional representations, where ϕ is a standard feed-forward neural network and ψ is a graph neural network. The two representation vectors $\phi(x)$ and $\psi(G_t)$ are concatenated to be an input to another feed-forward network g , which predicts the outcome.

our case is given as graphs. We expect that the use of such auxiliary information will be effective, especially when the number of treatments is relatively high and the training dataset is biased. We propose GraphITE, which utilizes graph-structured treatments while reducing selection bias effectively. We first introduce the network architecture of GraphITE, and then apply HSIC regularization to estimate outcomes appropriately from a biased dataset.

4.1 Model

The model of GraphITE consists of three components: two mapping functions $\phi : \mathcal{X} \rightarrow \Phi$ and $\psi : \mathcal{T} \rightarrow \Psi$ for extracting representations of the input and treatment graph, respectively, and a prediction function $g : \Phi \times \Psi \rightarrow \mathcal{Y}$ for predicting the outcome, where Φ and Ψ are the latent representation spaces of the inputs and treatments induced by ϕ and ψ , respectively. Figure 2 illustrates the overview of the neural network architecture of GraphITE.

As mentioned earlier, the mapping function ψ extracts representations that capture the features of graph-structured treatments. If we simply take ψ as a one-hot encoding of discrete treatments, it coincides with the standard setting with multiple treatments; however, this approach cannot take advantage of the rich structural information that the graph treatments have, and therefore suffers from a large number of treatments.

As the mapping function ψ of the treatment graphs, we employ a GNN. GNNs have been successfully applied in various domains [7, 10, 25] and are capable of extracting features of graphs owing to the flexible expressive power of neural networks optimized in an end-to-end manner.

The representation vector of graph-structured treatment $G = (\mathcal{V}, \mathcal{E})$ is obtained as follows. First, for each node $v_k \in \mathcal{V}$, the representation of v_k is initialized to a low-dimensional vector $\mathbf{v}_k^{(0)} \in \mathbb{R}^{D_\psi}$ determined by randomized initialization depending on the node label, such as the atom type. At the c -th layer of the GNN, the node representations are updated using

$$\mathbf{v}_k^{(c)} = \sigma^{\mathcal{V}} \left(\mathbf{W} \mathbf{v}_k^{(c-1)} + \sum_{v_m \in \mathcal{N}_k} \mathbf{M} \mathbf{v}_m^{(c-1)} \right), \quad (1)$$

where $\sigma^{\mathcal{V}}$ is an activation function, such as the ReLU function, $\mathcal{N}_k$ is the set of nodes adjacent to v_k , and $\mathbf{W}$ and $\mathbf{M}$ are transformation

¹Owing to the counterfactual nature of the problem, we are unable to observe the outcomes of the other treatments as they are not performed.

Figure 3: Training of GraphITE using the HSIC regularization. In addition to the prediction loss function ℓ between the prediction $g(\phi(x), \psi(G_t))$ and the true outcome y , the HSIC regularization term encourages the two representations $\phi(x)$ and $\psi(G_t)$ to be independent of each other in order to mitigate selection biases. Theorem 1 gives a theoretical guarantee on how the HSIC regularization contributes to the bias mitigation.

matrices. After the updates through C layers, the representations of all the nodes are aggregated into a graph-level representation $\psi(G)$ as

$$\psi(G) = \sum_{v_k \in \mathcal{V}} \sigma_G \left(\sum_{c=0}^C \mathbf{v}_k^{(c)} \right), \quad (2)$$

where σ_G is an activation function, such as the softmax function.

Note that, as we will see later, we require the treatment mapping function to be invertible, i.e., one-to-one, which most GNNs are not; however, some recent studies have proposed GNNs with the one-to-one property [29, 32, 55]. In our experiments, we use a non-invertible GNN which exhibits satisfactory performance.

4.2 Bias mitigation using HSIC regularization

With an unbiased dataset collected through randomized controlled trials (RCT), it suffices to minimize the objective function

$$\sum_{i=1}^N \ell(y_i^{t_i}, g(\phi(x_i), \psi(G_{t_i}))) \quad (3)$$

to estimate the components of GraphITE, ϕ , ψ , and g , where ℓ is a loss function, such as mean squared error (MSE). However, the objective function is biased when the training dataset is biased, and we must adjust it to mitigate the negative effect.

We first propose our approach from an intuitive viewpoint. The main source of the bias is that, in contrast with RCT, the treatments in the observational data are selected depending on the target individuals (i.e., the covariates). Our idea for mitigating the bias is to reduce the dependency, i.e., to find representations of the target and treatment that are as independent of each other as possible. To implement this idea, we employ HSIC [12] to measure the independence; the HSIC is defined as

$$\text{HSIC}(\phi, \psi) = (N-1)^{-2} \text{tr}(\mathbf{K}^\Phi \mathbf{H} \mathbf{K}^\Psi \mathbf{H}), \quad (4)$$

where $\mathbf{K}^\Phi$ and $\mathbf{K}^\Psi$ are the kernel matrices of the representations of the targets and treatments, respectively, and $\mathbf{H}$ is the centering matrix $\mathbf{H} = \mathbf{I} - \frac{1}{N} \mathbf{1}\mathbf{1}^T$. If the kernel function is characteristic, HSIC becomes 0 in expectation if and only if the two representations

Algorithm 1: GraphITE training procedure

Input: Observational data: $\mathcal{D} = \{(x_i, t_i, y_i^{t_i})\}_{i=1}^N \sim p^{\text{train}}$, a set of graph-structured treatments $\mathcal{G} = \{G_j\}_{j=1}^{|\mathcal{T}|}$, and a hyperparameter $\lambda \geq 0$.

Output: An outcome prediction model $f = (g, \phi, \psi)$.

while not converged do

 Sample a mini batch $\mathcal{B} = \{(x_{i_o}, t_{i_o}, y_{i_o}^{t_{i_o}})\}_{o=1}^B \subset \mathcal{D}$

 # Mini-batch approximation of the supervised loss (Eq. (3))

 Compute $L_{\mathcal{B}}(g, \phi, \psi) = \sum_{o=1}^B \ell(y_{i_o}, g(\phi(x_{i_o}), \psi(G_{t_{i_o}})))$

 # Mini-batch approximation of the HSIC loss (Eq. (4))

 Compute $\text{HSIC}_{\mathcal{B}}(\phi, \psi) = \text{HSIC}(\phi_{x \in \mathcal{B}}, \psi_{t \in \mathcal{B}})$

 # Update the parameters of f

 Minimize $L_{\mathcal{B}}(g, \phi, \psi) + \lambda \cdot \text{HSIC}_{\mathcal{B}}(\phi, \psi)$ using SGD

end

are independent; we use the Gaussian kernel as the kernel function. HSIC is somewhat computationally expensive, which requires $O(N^2)$ time and space complexity, and does not scale to the sample size. Therefore, for the sake of computational convenience, we compute the HSIC loss in a mini-batch manner.

With the HSIC as a regularization term, our objective function is modified to

$$\sum_{i=1}^N \ell(y_i^{t_i}, g(\phi(x_i), \psi(G_{t_i}))) + \lambda \text{HSIC}(\phi, \psi), \quad (5)$$

where λ is the regularization hyper-parameter. The ϕ is implemented as a standard feed-forward neural network, while ψ is a GNN; the ϕ and ψ are concatenated as an input to g that is another feed-forward network. Figure 3 illustrates the training of GraphITE using the HSIC regularization.

The objective function (5) is optimized in a mini-batch manner using Adam [24]; more specifically, each epoch divides the entire training dataset $\mathcal{D}$ into mini batches without overlapping, and approximates the loss function and HSIC term with them. Recent theoretical analysis reveals that minimizing the HSIC loss in a mini-batch manner is equivalent to bagging block HSIC method [33, 50], which ensures that it converges to the same value. The training procedure for GraphITE is outlined in Algorithm 1.

Finally, we give some historical remarks explaining why we specifically chose the HSIC as the regularization term. Previous studies have considered a broad class of regularization terms, integral probability metrics (IPM) [22, 42]; however, they basically assume a binary treatment or a relatively small number of treatments, and they cannot be directly applied to a large number of treatments. For example, typical IPM such as the maximum mean discrepancy (MMD) and Wasserstein distance, require many regularization terms for all pairs of treatments; otherwise, an expedient “pivot” control treatment must be set, which is not effective, as demonstrated in our experiments. The use of the HSIC is proposed by Lopez et al. [30] as it naturally allows multiple treatments. However, they did not consider learning representations of treatments.

4.3 Theoretical justification of HSIC regularization

Now we consider the theoretical justification for using the HSIC regularization in GraphITE. Our discussion is based on generalizations of the theories [23, 30] to treatments with features. We also discuss the benefits of our formulation and how it makes the prediction model more flexibly deal with complex situations than the existing approaches.

Denote by p^{train} the probability distribution on $\mathcal{X} \times \mathcal{T}$ the training dataset $\mathcal{D}$ follows, and by p^{test} the one for the test dataset. We assume that the test distribution has the form of $p^{\text{test}}(x, t) = p^{\mathcal{X}}(x)p^{\mathcal{T}}(t)$ because we want our model to perform well on the distribution where treatments do not depend on the covariates.

For some (unknown target) function $f^* : \mathcal{X} \times \mathcal{T} \rightarrow \mathcal{Y}$ and probability distribution p over $\mathcal{X} \times \mathcal{T}$, let the expected risk of our prediction model $f : \mathcal{X} \times \mathcal{T} \rightarrow \mathcal{Y}$ with the mapping functions ϕ, ψ and a predictive function g be

$$\epsilon_p(f_{g, \phi, \psi}) = \mathbb{E}_{(x, t) \sim p}[\ell(f(g(\phi(x), \psi(G_t))), f^*(x, t))]. \quad (6)$$

Then we have the following theoretical upper bound of the expected risk on the test distribution in terms of that for the training distribution and the HSIC regularization term.

THEOREM 1. *Let $\epsilon_{p^{\text{train}}}(f_g)$ and $\epsilon_{p^{\text{test}}}(f_g)$ be the expected risk for the training distribution and the test distribution, respectively. Let the IPM between two distributions p and q with respect to a function family $\mathcal{H}$ be*

$$\text{IPM}_{\mathcal{H}}(p, q) = \sup_{h \in \mathcal{H}} |\mathbb{E}_p[h] - \mathbb{E}_q[h]|. \quad (7)$$

Let $J_{\phi}^{-1}(\xi), J_{\psi}^{-1}(\tau)$ be the Jacobian matrices of ϕ^{-1} and ψ^{-1} at ξ and τ , respectively. Assume that there exist positive constants A and B that satisfy $|J_{\phi^{-1}(\xi)}||J_{\psi^{-1}(\tau)}| \leq A$ and $\frac{\ell_{f_{g, \phi, \psi}}}{B} \in \mathcal{H} \subseteq \{g : \Phi \times \Psi \rightarrow \mathcal{Y}\}$. Then the expected risk for the test distribution is upper-bounded by

$$\epsilon_{p^{\text{test}}}(f_{g, \phi, \psi}) \leq \epsilon_{p^{\text{train}}}(f_{g, \phi, \psi}) + \lambda \cdot \text{HSIC}(\phi, \psi). \quad (8)$$

PROOF.

$$\begin{aligned} & \epsilon_{p^{\text{test}}}(f_{g, \phi, \psi}) - \epsilon_{p^{\text{train}}}(f_{g, \phi, \psi}) \\ &= \int_{\mathcal{X} \times \mathcal{T}} \ell_{f_{g, \phi, \psi}}(x, t)(p(x)p(t) - p(x, t)) dx dt - \int_{\mathcal{X} \times \mathcal{T}} \ell_{f_{g, \phi, \psi}}(x, t)p(x, t) dx dt \\ &= \int_{\mathcal{X} \times \mathcal{T}} \ell_{f_{g, \phi, \psi}}(x, t)(p(x)p(t) - p(x, t)) dx dt \\ &= \int_{\Phi \times \Psi} \ell_{f_{g, \phi, \psi}}(\phi^{-1}(\xi), \psi^{-1}(\tau))(p(\xi)p(\tau) - p(\xi, \tau)) \\ & \quad \cdot |J_{\phi^{-1}(\xi)}||J_{\psi^{-1}(\tau)}| d\xi d\tau \\ &\leq A \int_{\Phi \times \Psi} \ell_{f_{g, \phi, \psi}}(\phi^{-1}(\xi), \psi^{-1}(\tau))(p(\xi)p(\tau) - p(\xi, \tau)) d\xi d\tau \\ &\leq A \cdot B \sup_{g \in \mathcal{H}} \left| \int_{\Phi \times \Psi} g(\phi^{-1}(\xi), \psi^{-1}(\tau))(p(\xi)p(\tau) - p(\xi, \tau)) d\xi d\tau \right| \\ &\leq A \cdot B \cdot \text{IPM}_{\mathcal{H}}(p(\xi)p(\tau), p(\xi, \tau)) \\ &\leq A \cdot B \cdot C \cdot \text{HSIC}(p(\xi), p(\tau)), \end{aligned}$$

Table 1: Summary statistics of the datasets.

Dataset	#Units	#Treatments	#Interactions
CCLE	491	24	11,054
GDSC	925	117	105,694

where $\ell_{f_{g, \phi, \psi}} = \mathbb{E}_y[\ell(y^t, g(\phi(x), \psi(G_t))) \mid x, t]$. C is a constant that denotes the radius of the function space. By setting $\lambda = A \cdot B \cdot C$, we obtain the inequality (8). $\square$

The theorem states that minimizing the HSIC between the representations of the targets and graph-structured treatments leads to minimizing the expected risk for the test distribution, making the predictive model to handle even unseen graph-structured treatments unbiasedly. For the inequality (8) to hold, we require several conditions: ϕ and ψ must be twice-differentiable one-to-one mapping functions, and the HSIC must be defined using continuous, bounded, positive semi-definite kernels $k^{\Phi} : \Phi \times \Phi \rightarrow \mathbb{R}$ and $k^{\Psi} : \Psi \times \Psi \rightarrow \mathbb{R}$. The λ must also be theoretically determined based on the radius of the function space in which f lies, but empirically, we simply treat it as a hyper-parameter. Our choices of the kernels and the hyper-parameters are detailed in Section 5.

Note that MMD and Wasserstein distance are special cases of IPM when the function family includes the set of 1-Lipschitz functions and the set of unit norm functions in a universal reproducing norm Hilbert space [42], respectively. Hence, they are obviously bounded by the inequality (8).

5 EXPERIMENTS

We experimentally investigate the performance of the proposed GraphITE and its merits of using the GNN and the HSIC regularization compared with various baseline methods on two real-world datasets.

5.1 Datasets

We use two real-world datasets on drug responses: the Cancer Cell-Line Encyclopedia (CCLE) dataset [4] and Genomics of Drug Sensitivity in Cancer (GDSC) dataset [51]. Table 1 lists their basic statistics. #Units, #Treatments, and #Interactions represent the number of units, the number of treatments, and the number of labeled data in a dataset. CCLE is a relatively small dataset with a moderate number of treatments, while GDSC has more than 100 treatments. Both of the datasets include IC_{50} values for drug-cell pairs, which are known to be closely related to drug sensitivity. Namely, we define the drug sensitivity as $y = -\log \text{IC}_{50}$ following previous studies [28, 44], which is the regression target in our experiments. We use the similarity matrices of each cell line as the input features. Both datasets are publicly available² [44].

Because the two original datasets are fairly close to complete observations (specifically, their observation rates are about 94% and 98%, respectively), we simply assume that they are complete, and introduce synthetic observation biases to extract biased training datasets, and then test the predictors obtained from them on the remainder. We introduce synthetic treatment bias that assigns treatment t using $t \sim \text{Categorical}(\text{softmax}(py))$ following previous

²<https://github.com/CSB5/CaDRReS>

studies [5, 40, 41]. The $\rho = \frac{\eta}{100\sigma}$ is a bias coefficient, where η is the magnitude of selection bias and σ is the standard deviation of target values. A larger η indicates a higher selection probability; intuitively, this indicates that scientists are more likely to conduct experiments for drug–cell pairs with higher sensitivity values. Note that although this bias generation procedure does not necessarily satisfy the typical unconfoundedness assumption, it is more practical and reasonable in the sense that the scientists likely to select promising experimental targets based on their knowledge and experience. In other words, we assume scientists are not incompetent.

5.2 Baseline methods

We compare GraphITE with the following six baselines. (i) Ordinary least squares linear regression (OLS) concatenates two vectors, the covariate vector and treatment vector coded as a one-hot vector, which is used as the input. (ii) Bayesian additive regression trees (BART) [6, 19] predicts the outcomes by an ensemble of multiple regression trees; we used a Python implementation of BART³. (iii) Treatment embedding method exploits low-dimensional representations of treatments to deal with a large number of treatments. Each treatment is associated with a low-dimensional vector, which is input to a neural network, as well as a covariate vector. Note that this method does not use the graph structures of the treatments at all. (iv) TARNet [42] is a deep neural network model with shared layers for representation learning and different layers for outcome prediction for treatment and control instances. (v) Counterfactual regression (CFR) [42] is one of the state-of-the-art deep neural network models based on balanced representations between treatment and control instances; we used the MMD as its IPM. Following previous studies [38, 53], we extend the CFR [42] to the multiple-treatment setting; we regard the most frequent treatment as the control treatment. (vi) GANITE [53] is another state-of-the-art deep neural network model based on GAN. It trains a TARNet-like generator that generates counterfactual outcomes, and a discriminator tells whether outcomes come from the generator or the real distribution. In the original GANITE, the discriminator just tries to solve a binary classification; on the other hand in our setting, GANITE has to solve multi-class classification to tell which outcome is the genuine one.

In addition to the use of graph structured treatments, one of the key features of GraphITE is the bias mitigation using HSIC regularization; therefore, we use the versions without it our baseline methods for the ablation study. We also tested several variants of GraphITE: (vii) a variant with no bias mitigation that does not have the HSIC regularization term and only uses a GNN, which we refer to as “GNN” hereafter and (viii) another variant using MMD regularization instead of HSIC regularization. We used the same approach as CFR to deal with multiple treatments. We denote it by “GNN+MMD”.

5.3 Experimental setting

As the evaluation metrics, we employ the root mean square error (RMSE) of all target–treatment pairs in the test set defined

as:

$$\text{RMSE} = \sqrt{\frac{1}{N^{\text{test}}} \frac{1}{|\mathcal{T}|} \sum_{i=1}^{N^{\text{test}}} \sum_{t=1}^{|\mathcal{T}|} (y_i^t - f(x_i, t))^2}, \quad (9)$$

where N^{test} is the number of target individuals included in the test dataset. We also employ the concordance index (CI) [17] to evaluate predictive performance in terms of ranking accuracy, which has been widely used in previous studies [26, 37]. The CI is defined as

$$\text{CI} = \frac{1}{N^{\text{test}}} \sum_{i=1}^{N^{\text{test}}} \sum_{t, u | y_i^t > y_i^u} \frac{\theta(f(x_i, t) - f(x_i, u))}{|\{t, u | y_i^t > y_i^u\}|}, \quad (10)$$

where γ is the Heaviside step function defined as

$$\theta(x) = \begin{cases} 1, & x > 0 \\ 0.5, & x = 0 \\ 0 & x < 0 \end{cases}. \quad (11)$$

Note that the CI is identical to ROC-AUC when all outcomes are binary.

We split the whole individuals into 80%, 10%, and 10% for training, validation, and testing sets, respectively. We report the average results of 50 different trials. Note that while we sample the factual treatments in the training and validation sets following the biased sampling scenario explained in the Dataset section, all of the treatments are included in the test set because we want our prediction model to perform uniformly well on all treatments.

In GraphITE, to promote effective feature extraction from small data, we pre-train the GNN ψ on regression tasks using three popular molecular datasets: ESOL, FreeSolv, and Lipophilicity, provided by MoleculeNet [48]⁴. To the best of our observation, GraphITE performs slightly better with pre-training than the one without pre-training. We believe that this phenomenon may be caused by the small size of the training data, and is the one of limitations of this study to be addressed in future work. For the HSIC regularization, we use the normalized version of the HSIC (nHSIC), defined as

$$\text{nHSIC}(\phi, \psi) = \frac{\text{tr}(\mathbf{K}^\phi \mathbf{H} \mathbf{K}^\psi \mathbf{H})}{\|\mathbf{K}^\phi \mathbf{H}\| \|\mathbf{K}^\psi \mathbf{H}\|}, \quad (12)$$

where $\|\cdot\|$ denotes the Frobenius norm. The regularization parameter λ is optimized in $\{10^{-3}, 10^{-2}, \dots, 10^3\}$ based on the RMSE for the validation set. We set the number of representation dimensions for target individuals and graph-structured treatments to 64 because we did not observe the significant differences in $\{16, 32, 64, 128\}$ in the both datasets. Similarly, we set the numbers of layers of ϕ , ψ , and g to 3.

5.4 Results

Table 2 summarizes the predictive performances of the different methods for $\eta = 40$ (i.e., the strongest bias). In the remainder, we report the experimental results when we set $\eta = 40$ unless otherwise stated.

The deep learning-based methods (Treatment Embedding, TARNet, CFR) outperform the naïve methods, such as Mean and OLS. BART also gives the comparable performance. However, we observe GANITE performs poorly in comparison with the other

³<https://github.com/JakeColtman/bartpy>

⁴<http://moleculenet.ai/>

Table 2: Performance comparison of different methods on the CCLE and GDSC dataset in terms of RMSE and CI. [†] and [‡] indicate statistically significantly better performance of the proposed GraphITE than the baseline by the paired t -test with $p < 0.05$ and $p < 0.01$, respectively. The bold results indicate the statistically significant best results. The shaded rows indicate the GNN-based methods. Lower RMSEs are better, and higher CIs are better.

	CCLE		GDSC	
Method	RMSE	CI	RMSE	CI
Mean	$\ddagger 3.777 \pm 0.101$	—	$\ddagger 4.030 \pm 0.102$	—
OLS	$\ddagger 4.861 \pm 0.755$	$\ddagger 0.642 \pm 0.021$	$\ddagger 6.463 \pm 0.493$	$\ddagger 0.602 \pm 0.018$
BART	$\ddagger 2.993 \pm 0.203$	$\ddagger 0.711 \pm 0.016$	$\ddagger 3.965 \pm 0.102$	$\ddagger 0.632 \pm 0.015$
Treatment Embedding	$\ddagger 2.662 \pm 0.161$	$\ddagger 0.724 \pm 0.013$	$\ddagger 3.642 \pm 0.131$	$\ddagger 0.670 \pm 0.015$
TARNet	$\ddagger 2.831 \pm 0.123$	$\ddagger 0.711 \pm 0.013$	$\ddagger 3.813 \pm 0.135$	$\ddagger 0.663 \pm 0.009$
CFR	$\ddagger 2.822 \pm 0.121$	$\ddagger 0.712 \pm 0.013$	$\ddagger 3.792 \pm 0.134$	$\ddagger 0.664 \pm 0.009$
GANITE	$\ddagger 3.652 \pm 0.211$	$\ddagger 0.651 \pm 0.023$	$\ddagger 7.739 \pm 1.394$	$\ddagger 0.613 \pm 0.018$
GNN	$\ddagger 2.652 \pm 0.123$	$\ddagger 0.720 \pm 0.010$	$\ddagger 3.553 \pm 0.126$	$\ddagger 0.681 \pm 0.010$
GNN+MMD	$\ddagger 2.596 \pm 0.162$	$\ddagger 0.726 \pm 0.014$	$\ddagger 3.531 \pm 0.136$	$\ddagger 0.683 \pm 0.013$
GraphITE (Proposed)	2.561 ± 0.112	0.732 ± 0.009	3.421 ± 0.135	0.695 ± 0.015

Figure 4: Sensitivity of the results to the regularization strength λ . Although the optimal choices significantly improve the performance, at worst the other choices do not harm the performance.

Figure 5: Predictive performance depending on the bias coefficient η . A large η indicates a larger selection bias. GraphITE shows its strong and stable tolerance to the biases and consistently performs the best in the whole range.

deep learning-based methods. We believe this is because of the difficulty in learning the GAN models, that is, GANITE has to solve many-class classification with a limited amount of training data. The existing bias mitigation methods that are simply extended to multiple treatments (CFR and GNN+MMD) do not show significant improvements over the corresponding original ones (TARNet and GNN). By contrast, GraphITE achieves the best performance, which is statistically significant against all the

baselines, and it demonstrates its effectiveness, especially on the larger dataset (GDSC). The merit of exploiting the graph structures associated with treatments can also be seen in Table 2. The GNN-based methods (shaded rows) perform better than the methods that neglect graph-structured information.

Figure 4 shows the sensitivity of the results to the strength of HSC regularization; although the optimal choices significantly improves the performance, none of the other choices harm the

Figure 6: Predictive performance depending on the the number of treatments. Whereas the baseline methods get degraded as the number of treatments increase, graphite shows relatively robust to its increase and achieves the best performances, especially in the larger dataset (GDSC).

Figure 7: Predictive performance depending on treatment popularity. From the RMSE results, the methods that do not rely on treatment graph information (CFR, TARNet, BART) suffer from a lack of data especially for unpopular treatments. From the CI results, the methods that have no bias mitigation mechanism put too much attentions on popular treatments (i.e., difficult in terms of ranking) treatments, and perform suboptimally. GraphITE shows the most stable and best performance on every group.

Table 3: Performance comparison of different methods on the CCLE and GDSC dataset in terms of RMSE and CI. [†] and [‡] indicate statistically significantly better performance of the proposed GraphITE than the baseline by the paired t -test with $p < 0.05$ and $p < 0.01$, respectively. The bold results indicate the statistically significant best results. The shaded rows indicate the GNN-based methods. Lower RMSEs are better, and higher CIs are better.

	CCLE		GDSC	
Method	RMSE	CI	RMSE	CI
Mean	3.458 $\pm$ 1.301	—	[†] 4.705 $\pm$ 0.702	—
GNN	3.920 $\pm$ 0.932	0.551 $\pm$ 0.130	[‡] 4.646 $\pm$ 0.631	0.570 $\pm$ 0.061
GNN+MMD	3.903 $\pm$ 0.923	0.549 $\pm$ 0.132	[‡] 4.640 $\pm$ 0.674	0.574 $\pm$ 0.061
GraphITE	3.637 $\pm$ 0.905	0.545 $\pm$ 0.114	4.482$\pm$0.595	0.569 $\pm$ 0.054

performance. The results also highlight the effectiveness of the HSIC as the choice for the regularization term; MMD shows no remarkable improvement over the plain GNN because it cannot handle many treatments efficiently, whereas GraphITE using HSIC regularization shows distinct improvements.

Now, we investigate the robustness of GraphITE against selection bias. Figure 5 shows the performances for different bias strengths, where a larger η represents a larger selection bias. GraphITE shows its strong and stable tolerance to the biases and consistently performs best under all bias strength settings.

The impacts of the number of treatments are shown in Figure 6. For small numbers of treatments, both GraphITE and the other baseline methods perform similarly well; however, the baseline methods, especially BART, degrade the performances as the number

increase. GNN+MMD does not show improvements from the plain GNN, particularly on the larger dataset (GDSC). GraphITE shows the remarkable robustness to the selection bias, even with large numbers of treatments.

Next, we investigate the predictive performance based on treatment popularity, as shown in Figure 7. We focus on the treatment groups that are grouped by their popularity, namely, the top 20%, 20–40%, 40–60%, 60–80% most popular treatment groups. As can be seen from the RMSE results, the methods that do not rely on treatment graph information (CFR, TARNet, and BART) suffer from a lack of training data especially for unpopular treatments; On the other hand, the methods exploiting the auxiliary information mitigate this problem.

Now we turn to the CI results. From its definition (10), CI measures ranking performance. It is more difficult to estimate accurate ranking for popular treatments rather than less popular treatments, because in our bias setting, more popular treatments have larger outcomes, and they are more heterogeneous sets than less popular ones. On the other hand, it is rather easy to obtain small RMSEs for popular treatments because their outcomes have small variances. The methods that have no bias mitigation mechanism put too much attentions on such difficult treatments (in terms of ranking), and eventually perform suboptimally, while GraphITE shows the most stable and best performance on every group.

Finally, we investigate the predictive performance on the unobserved “zero-shot” treatments that are not included in training data. We keep 30% of the entire treatments aside in advance as the validation and target unobserved treatments. Table 3 shows the prediction accuracy for the unobserved treatments when we set $\eta = 40$. Note that the existing methods such as CFR are incapable of dealing with this setting. In the smaller dataset, CCLE, the selection bias prevents the models from working appropriately, and all of the variants perform worse than the simply baseline taking the mean of training data in terms of RMSE. On the other hand in the larger dataset, GDSC, GraphITE achieves the best RMSE, while the other methods still suffer from the bias. However, we do not observe significant improvements in terms of CI in the both datasets.

6 CONCLUSION

In this study, we proposed GraphITE, which can handle graph-structured treatments in order to achieve better treatment effect estimation even when the number of treatments is large. GraphITE is based on the recent developments of deep neural networks and representation learning, namely, GNNs and HSIC regularization, which contribute to improving estimation accuracy of complex-structured treatments from biased observational data. In addition, GraphITE is applicable to previously unobserved “zero-shot” treatments, which the existing ITE estimation methods are intrinsically not capable of dealing with. In the experiments on two real-world drug response datasets, GraphITE achieved the best performances in terms of RMSE and CI when compared to the various baselines. In particular, we observed a significant improvement when the effect of selection bias and the number of treatments were large. A potential future direction is to consider other types of complex structured data, such as texts, images, and videos. We also plan to apply GraphITE to much larger datasets, in which we expect further improvements.

ACKNOWLEDGMENTS

This work was supported by JSPS KAKENHI Grant Number 21J14882 and 20H04244.

REFERENCES

- [1] Alberto Abadie and Guido W Imbens. 2006. Large sample properties of matching estimators for average treatment effects. *Econometrica* (2006).
- [2] Hamidreza Alviri, Elham Shaabani, Soumajyoti Sarkar, Ghazaleh Beigi, and Paulo Shakarian. 2019. Less is more: Semi-supervised causal inference for detecting pathogenic users in social media. In *Proceedings of the Web Conference*.
- [3] Michael Baiocchi, Jing Cheng, and Dylan S Small. 2014. Instrumental variable methods for causal inference. *Statistics in Medicine* (2014).
- [4] Jordi Barretina, Giordano Caponigro, Nicolas Stransky, Kavitha Venkatesan, Adam A Margolin, Sungjoon Kim, Christopher J Wilson, Joseph Lehar, Gregory V Kryukov, Dmitriy Sonkin, et al. 2012. The Cancer Cell Line Encyclopedia enables predictive modelling of anticancer drug sensitivity. *Nature* (2012).
- [5] Ioana Bica, Ahmed M Alaa, Craig Lambert, and Mihaela van der Schaar. 2021. From real-world patient data to individualized treatment effects using machine learning: current and future methods to address underlying challenges. *Clinical Pharmacology & Therapeutics* (2021).
- [6] Hugh A Chipman, Edward I George, Robert E McCulloch, et al. 2010. BART: Bayesian additive regression trees. *The Annals of Applied Statistics* (2010).
- [7] David K Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P Adams. 2015. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in Neural Information Processing Systems*.
- [8] H-G Eichler, Brigitte Bloechl-Daum, Peter Bauer, Frank Bretz, Jeffrey Brown, Lisa Victoria Hampson, Peter Honig, Michael Krams, Hubert Leufkens, Robyn Lim, et al. 2016. “Threshold-crossing”: a useful way to establish the counterfactual in clinical trials? *Clinical Pharmacology & Therapeutics* (2016).
- [9] Brigitte Ganter, Stuart Tugendreich, Cecelia I Pearson, Eser Ayanoglu, Susanne Baumhueter, Keith A Bostian, Lindsay Brady, Leslie J Browne, John T Calvin, Gwo-Jen Day, et al. 2005. Development of a large-scale chemogenomics database to improve drug candidate selection and to understand mechanisms of chemical toxicity and action. *Journal of Biotechnology* (2005).
- [10] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. 2017. Neural Message Passing for Quantum Chemistry. In *Proceedings of the 34th International Conference on Machine Learning*.
- [11] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial networks. In *Advances in Neural Information Processing Systems*.
- [12] Arthur Gretton, Kenji Fukumizu, Choon H Teo, Le Song, Bernhard Schölkopf, and Alex J Smola. 2001. A kernel statistical test of independence. In *Advances in Neural Information Processing Systems*.
- [13] Ruocheng Guo, Jundong Li, and Huan Liu. 2020. Learning Individual Causal Effects from Networked Observational Data. In *Proceedings of the 13th International Conference on Web Search and Data Mining*.
- [14] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*.
- [15] Shonosuke Harada, Hirotaka Akita, Masashi Tsubaki, Yukino Baba, Ichigaku Takigawa, Yoshihiro Yamanishi, and Hisashi Kashima. 2020. Dual graph convolutional neural network for predicting chemical networks. *BMC Bioinformatics* (2020).
- [16] Shonosuke Harada and Hisashi Kashima. 2020. Counterfactual Propagation for Semi-Supervised Individual Treatment Effect Estimation. In *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*.
- [17] Frank E Harrell Jr, Kerry L Lee, and Daniel B Mark. 1996. Multivariable prognostic models: issues in developing models, evaluating assumptions and adequacy, and measuring and reducing errors. *Statistics in Medicine* (1996).
- [18] Mikael Henaff, Joan Bruna, and Yann LeCun. 2015. Deep convolutional networks on graph-structured data. *arXiv preprint arXiv:1506.05163* (2015).
- [19] Jennifer L Hill. 2011. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics* (2011).
- [20] Swen Hoelder, Paul A Clarke, and Paul Workman. 2012. Discovery of small molecule cancer drugs: successes, challenges and opportunities. *Molecular oncology* (2012).
- [21] Guido W Imbens. 2000. The role of the propensity score in estimating dose-response functions. *Biometrika* (2000).
- [22] Fredrik Johansson, Uri Shalit, and David Sontag. 2016. Learning representations for counterfactual inference. In *Proceedings of the 33rd International Conference on Machine Learning*.
- [23] Fredrik D Johansson, Uri Shalit, Nathan Kallus, and David Sontag. 2020. Generalization bounds and representation learning for estimation of potential outcomes and causal effects. *arXiv preprint arXiv:2001.07426* (2020).
- [24] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [25] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. In *Proceedings of the 5th International Conference on Learning Representations*.
- [26] Roman Kurilov, Benjamin Haibe-Kains, and Benedikt Brors. 2020. Assessment of modelling strategies for drug response prediction in cell lines and xenografts. *Scientific Reports* (2020).
- [27] Robert J LaLonde. 1986. Evaluating the econometric evaluations of training programs with experimental data. *The American Economic Review* (1986).
- [28] Alex P Lind and Peter C Anderson. 2019. Predicting drug activity against cancer cells by random forest models based on minimal genomic information and chemical properties. *PLoS ONE* (2019).
- [29] Jenny Liu, Aviral Kumar, Jimmy Ba, Jamie Kiros, and Kevin Swersky. 2019. Graph normalizing flows. In *Advances in Neural Information Processing Systems*.

- [30] Romain Lopez, Chenchen Li, Xiang Yan, Junwu Xiong, Michael I Jordan, Yuan Qi, and Le Song. 2020. Cost-Effective Incentive Allocation via Structured Counterfactual Inference. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*.
- [31] Yunpu Ma, Yuyi Wang, and Volker Tresp. 2020. Causal Inference under Networked Interference. *arXiv preprint arXiv:2002.08506* (2020).
- [32] Ryan Murphy, Balasubramaniam Srinivasan, Vinayak Rao, and Bruno Ribeiro. 2019. Relational Pooling for Graph Representations. In *Proceedings of the 36th International Conference on Machine Learning*.
- [33] Thao Nguyen, Maithra Raghu, and Simon Kornblith. 2021. Do wide and deep networks learn the same things? Uncovering how neural network representations vary with width and depth. *arXiv:cs.LG/2010.15327*
- [34] Paul R Rosenbaum and Donald B Rubin. 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika* (1983).
- [35] Donald B Rubin. 1973. Matching to remove bias in observational studies. *Biometrics* (1973).
- [36] Donald B Rubin. 2005. Causal inference using potential outcomes: Design, modeling, decisions. *J. Amer. Statist. Assoc.* (2005).
- [37] Zhaleh Safikhani, Petr Smirnov, Kelsie L Thu, Jennifer Silvester, Nehme El-Hachem, Rene Quevedo, Mathieu Lupien, Tak W Mak, David Cescon, and Benjamin Haibe-Kains. 2017. Gene isoforms as expression-based biomarkers predictive of drug response in vitro. *Nature Communications* (2017).
- [38] Shiv Kumar Saini, Sunny Dhamnani, Akil Arif Ibrahim, and Prithviraj Chavan. 2019. Multiple Treatment Effect Estimation using Deep Generative Model with Task Embedding. In *Proceedings of the Web Conference*.
- [39] Kristof T Schütt, Huziel E Sauceda, P-J Kindermans, Alexandre Tkatchenko, and K-R Müller. 2018. SchNet—a deep learning architecture for molecules and materials. *The Journal of Chemical Physics* (2018).
- [40] Patrick Schwab, Lorenz Linhardt, Stefan Bauer, Joachim M Buhmann, and Walter Karlen. 2020. Learning counterfactual representations for estimating individual dose-response curves. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*.
- [41] Patrick Schwab, Lorenz Linhardt, and Walter Karlen. 2018. Perfect match: A simple method for learning representations for counterfactual inference with neural networks. *arXiv preprint arXiv:1810.00656* (2018).
- [42] Uri Shalit, Fredrik D Johansson, and David Sontag. 2017. Estimating individual treatment effect: generalization bounds and algorithms. In *Proceedings of the 34th International Conference on Machine Learning*.
- [43] Claudia Shi, David Blei, and Victor Veitch. 2019. Adapting Neural Networks for the Estimation of Treatment Effects. In *Advances in Neural Information Processing Systems*.
- [44] Chayaporn Suphavitai, Denis Bertrand, and Niranjan Nagarajan. 2018. Predicting cancer drug response using a recommender system. *Bioinformatics* (2018).
- [45] Victor Veitch, Yixin Wang, and David Blei. 2019. Using embeddings to correct for unobserved confounding in networks. In *Advances in Neural Information Processing Systems*.
- [46] Stefan Wager and Susan Athey. 2018. Estimation and inference of heterogeneous treatment effects using random forests. *J. Amer. Statist. Assoc.* (2018).
- [47] Hao Wang, Hao He, and Dina Katabi. 2020. Continuously Indexed Domain Adaptation. In *Proceedings of the 37th International Conference on Machine Learning*.
- [48] Zhenqin Wu, Bharath Ramsundar, Evan N Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S Pappu, Karl Leswing, and Vijay Pande. 2018. MoleculeNet: a benchmark for molecular machine learning. *Chemical Science* (2018).
- [49] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2018. How Powerful are Graph Neural Networks?. In *Proceedings of the 6th International Conference on Learning Representations*.
- [50] Makoto Yamada, Yuta Umez, Kenji Fukumizu, and Ichiro Takeuchi. 2018. Post selection inference with kernels. In *International Conference on Artificial Intelligence and Statistics*.
- [51] Wanjun Yang, Jorge Soares, Patricia Greninger, Elena J Edelman, Howard Lightfoot, Simon Forbes, Nidhi Bindal, Dave Beare, James A Smith, I Richard Thompson, et al. 2012. Genomics of Drug Sensitivity in Cancer (GDSC): a resource for therapeutic biomarker discovery in cancer cells. *Nucleic Acids Research* (2012).
- [52] Liuyi Yao, Sheng Li, Yaliang Li, Mengdi Huai, Jing Gao, and Aidong Zhang. 2018. Representation learning for treatment effect estimation from observational data. In *Advances in Neural Information Processing Systems*.
- [53] Jinsung Yoon, James Jordon, and Mihaela van der Schaar. 2018. GANITE: Estimation of individualized treatment effects using generative adversarial nets. In *Proceedings of the 6th International Conference on Learning Representations*.
- [54] Jiaxuan You, Bowen Liu, Zhitao Ying, Vijay Pande, and Jure Leskovec. 2018. Graph convolutional policy network for goal-directed molecular graph generation. In *Advances in Neural Information Processing Systems*.
- [55] Chengxi Zang and Fei Wang. 2020. Moflow: An invertible flow model for generating molecular graphs. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- [56] Muhan Zhang and Yixin Chen. 2018. Link Prediction Based on Graph Neural Networks. In *Advances in Neural Information Processing Systems*.
- [57] Siyuan Zhao and Neil Heffernan. 2017. Estimating Individual Treatment Effect from Educational Studies with Residual Counterfactual Networks. In *Proceedings of the 10th International Conference on Educational Data Mining*.