

A Targeted Learning Framework for Policy Evaluation with Unobserved Network Interference

Anonymous authors
Paper under double-blind review

Abstract

Estimating causal effects under network interference is a fundamental yet challenging task, especially when the network structure is represented as multiple layers or multiple views. In this paper, we consider a heterogeneous network setting, where the ties from different views of the network might achieve varying levels of interference. Meanwhile, dependence among units is allowed, due to information transmission among network ties and latent traits among units sharing ties (i.e., latent dependency). To the best of our knowledge, this setting has not been studied in literature yet. We propose a novel framework that conducts doubly robust estimation on heterogeneous networks with latent dependency. Our approach relies on a new identification strategy and integrates it with targeted maximum likelihood estimation for robust causal effect estimation from observational data. Crucially, our approach remains valid even when the outcome prediction model or data-generating process is misspecified. It also supports counterfactual inference under hypothetical network interventions using only the observed network structure. Experiments on both synthetic and real-world networks show that our approach consistently outperforms existing baselines and can provide robust estimation towards different intervention policies.

1 Introduction

Estimating causal effects under network interference is fundamental to evidence-based policy-making in social science, public health, and economics. A central challenge is that the outcome of a unit may depend not only on its own treatment, but also on the treatments received by its neighbours (van der Laan, 2014), violating the classical SUTVA assumption. This phenomenon, known as interference, is pervasive in real-world networks: vaccination campaigns create herd immunity, advertising interventions create spillover demand, and job-training programmes affect peers’ employment outcomes. A second, underappreciated challenge is that real-world interference networks are rarely homogeneous or fully observed. Units may be connected via multiple types of relationships—for example, co-viewing and co-purchasing links on an e-commerce platform (He & McAuley, 2016)—and the true interference network is often only partially observed or must be approximated from auxiliary data. These two challenges—heterogeneous multi-view structure and partial observability—are rarely addressed jointly in the causal inference literature.

A common approach to model interference is to define exposure as the proportion of treated neighbors (Forastiere et al., 2022), which allows the estimation of spillover effects, i.e., the effect of varying neighborhood exposure on a unit’s outcome. In contrast to modeling exposure solely as a proportion, Toulis & Kao (2013) introduced a framework that defines exposure through the exact k -level exposure, i.e., the potential outcome of a node when exactly k of its neighbors are treated. Another way to model network interference involves modeling the transmission of information through network ties using graph neural networks (GNN) based methods (Tchetgen et al., 2019; Ma et al., 2021; 2022; Ma & Tresp, 2021) or directly modeling the confounders to address latent network dependency (Guo et al., 2019; Ma et al., 2021; Zhang et al., 2017).

However, real-world networks may be more complex. One complexity comes from the heterogeneous relationships among units. For example, in Amazon dataset (He & McAuley, 2016), ties may represent different types of interaction such as co-viewing and co-purchasing, and the interference may vary between different relational aspects. Lin et al. (2023) proposed HINITE, which aims to estimate the effect of treatment

in a heterogeneous setting by aggregating the covariate information on different aspects of the networks. Related efforts also include modeling heterogeneity in peer influence using structured graph-based causal models (Adhikari & Zheleva, 2025), and modeling heterogeneity in causal pathways across subpopulations with different moderator variables (Watson et al., 2023). These methods work well when the outcome model is correctly specified for the data, however, they are usually sensitive to model misspecification.

To avoid failure caused by misspecification between the data generation process and the modeling process, the estimation of maximum likelihood is a general framework to construct an objective estimate (Van Der Laan & Rubin, 2006; Van der Laan et al., 2011). This method has been widely used in various settings (van der Laan & Gruber, 2012; Kreif et al., 2017; Chen et al., 2023; Balzer et al., 2019). Shi et al. (2019) proposed targeted regularization for binary treatment effect estimation, while Nie et al. (2021) proposed an end-to-end targeted maximum likelihood estimation (TMLE) framework to estimate the continuous treatment effect. Based on previous work, Chen et al. (2024) proposed T-Net, which models the potential outcome as a function of unit treatment and neighborhood exposure in a homogeneous network setting. In this line of research, only one method (i.e., T-Net) is designed for causal inference on networked data, but it still cannot account for heterogeneity and latent dependency in the network structure.

To close the gap, we propose a novel framework, named **Multi-view doubly robust estimator (Mvdr)**, for causal inference under heterogeneous network interference. The challenge of heterogeneity remains valid even when only a single partially observed network is available. Our goal is to estimate the average expected potential outcome under different hypothetical intervention policies over an observed network with heterogeneous ties. Our approach models interference via multi-view representations of the network and integrates these with a TMLE procedure. In this way, the learned estimators are doubly robust, i.e., they remain consistent if either the outcome model or the exposure model is correctly specified. Our Mvdr framework enables robust estimation of average potential outcomes under a variety of hypothetical interventions, such as static, dynamic, or stochastic interventions. The major contributions of this paper can be summarized as follows:

1. We establish an identification and estimation framework for causal inference on partially observed, heterogeneous networks with latent dependency, which implies that the TMLE estimator inherits double robustness even when the summary functions are learned from data rather than pre-specified, as long as the aggregation satisfies Assumptions 1–3. This generalizes prior TMLE-based network estimators, which require a fixed, known exposure mapping.
2. We propose Mvdr, a multi-view doubly robust estimator that integrates graph neural network-based summary learning with TMLE. Unlike existing methods that either assume a correctly specified outcome model or a single fixed network, Mvdr simultaneously handles: (i) heterogeneous network structure via multi-view GCN representations, (ii) latent unit-level dependence via a local dependence assumption, and (iii) model misspecification via the TMLE perturbation step. To our knowledge, this is the first framework to jointly address all three challenges.
3. We design new evaluation settings with dynamic and stochastic interventions beyond the standard static setting, construct a new semi-synthetic DBLP dataset, and conduct extensive experiments demonstrating that MVDR consistently outperforms baselines, particularly under dynamic and stochastic policies where exposure mechanism misspecification is most consequential.

The rest of this paper is organized as follows. In Section 2, we briefly present related work on causal inference under network interference, complex interference structure, and model misspecification. Section 3 introduces preliminary and the problem settings. Section 4 presents our theoretical analysis. In Section 5, we introduce the technical details of the proposed Mvdr framework. Section 6 provides experimental results with discussions, and Section 7 concludes this paper.

2 Related Works

Causal Inference under Network Interference. Classical causal inference frameworks assume that there is no interference between units, but real-world settings often violate this assumption. The foundational work in (Halloran & Struchiner, 1995; Sobel, 2006; Hudgens & Halloran, 2008) formalized potential outcomes

under interference. Manski (2013) and Aronow & Samii (2017) introduced exposure mappings, summarizing neighbor treatments into interpretable variables. Building on this, Karwa & Airoidi (2018) investigated the consequences of misspecified exposure mappings in randomized experiments, while Leung (2022) studied approximate neighborhood interference and provided inference methods under specification error. More recently, Ogburn et al. (2022) developed a semiparametric TMLE framework for causal inference on a single network, allowing dependence to grow with network degree. Sävje (2024) emphasized the distinction between exposures as effect definitions versus structural assumptions, proposing expected exposure effects that remain meaningful under misspecification.

Complex Interference Structure. This motivates the construction of proxy networks based on observed covariates. Egami (2021) considered spillover effects when the interference network is only partially observed, while Li & Wager (2022) studied network interference under latent graphon models and equilibrium settings. Parallel work in graph mining has explored multiview and multiplex networks, where multiple graphs encode different relations among the same nodes. Such methods aggregate heterogeneous views through attention or weighting. These ideas have not yet been fully integrated into causal inference. Our work is motivated by the observation that units may be dependent not only on direct social ties, but also on shared environments, socioeconomic status, or latent similarities. Thus, constructing multiple similarity-based networks from observed covariates provides a flexible way to capture such heterogeneity.

Model Misspecification. A growing literature addresses robustness when models meet with misspecification with data generating process. Leung (2022) proposed alternative estimands for approximate interference, while Sävje (2024) showed that conventional estimators remain unbiased for expected exposure effects under weakly dependent specification errors. Chao et al. (2025) studied misclassified or surrogate networks, introducing methods to correct bias using validation data. Similarly, Hoshino et al. (2024) considered causal inference with noncompliance and unknown network structures. These works highlight that perfect knowledge of the interference graph is rarely available, and meaningful causal estimands can still be defined under approximation. Our approach draws directly on these insights. By treating similarity-based graphs as proxies rather than true causal pathways, we target expected exposure effects that remain interpretable and estimable even if the constructed networks do not fully align with the underlying social process.

3 Preliminary

In this section, we first introduce the notations used in this paper and then describe the preliminary of our work, such as the structural equation model and assumptions. Finally, we present the problem setting and introduce the causal estimands of interest in detail.

3.1 Notations

Let $x_i \in \mathbb{R}^d$ denote the covariates of unit i , $t_i \in \{0, 1\}$ denote the binary treatment assignment, and $y_i \in \mathbb{R}$ denote the observed outcome for unit i under treatment t_i . We then use $X = (x_1, \dots, x_n) \in \mathbb{R}^{n \times d}$, $T = (t_1, \dots, t_n) \in \{0, 1\}^n$, and $Y = (y_1, \dots, y_n) \in \mathbb{R}^n$ denote the covariates, treatment assignments, and observed outcomes for all n units in a network, respectively.

Suppose that the observed network G is of size n and denote the adjacency matrix of G as $A \in \{0, 1\}^{n \times n}$. If there exists an edge between unit i and j , $A_{ij} = 1$; otherwise, $A_{ij} = 0$. In practice, network G can be represented as multiple layers or multiple views (e.g., with adjacency matrices $A^{(1)}, \dots, A^{(K)}$ for K views), and each view captures a specific relationship between units. The observed network G can be seen as an aggregation of these multi-view networks. In this work, we consider the case of undirected edges, where influence between connected units is assumed to be symmetric. A table of all notations used in this paper is provided in Appendix A.

3.2 Structural Equation Model

We assume that the observed data arise from a complete network of size n . Units in this network may exhibit latent variable dependence, meaning that unobserved traits shared by connected units can induce correlations in their covariates, treatments, or outcomes (Shalizi & Thomas, 2011). This type of dependence




Figure 1: Structural equation model for the causal assumptions. The left one is which our model built on, where the G represents learned graph structure.

is typically stronger between units who are closer in the network topology, which is dependent on the network information but not necessarily fully observed from the adjacency matrix.

To formalize the data generating process, we adopt the structural equation model (SEM) framework (Pearl, 2012) and assume that the data are generated by sequentially evaluating the following set of equations:

$$\begin{aligned}
 X_i &= f_X[\epsilon_{X_i}], \quad i = 1, \dots, n \\
 T_i &= f_T[\{X_j : A_{ij} = 1\}, \epsilon_{T_i}], \quad i = 1, \dots, n \\
 Y_i &= f_Y[\{T_j : A_{ij} = 1\}, \{X_j : A_{ij} = 1\}, \epsilon_{Y_i}], \quad i = 1, \dots, n
 \end{aligned} \tag{1}$$

where f_X , f_T , and f_Y are unknown and unspecified functions that may depend on the unit i . ϵ_{X_i} , ϵ_{T_i} , and ϵ_{Y_i} are all exogenous, unobserved errors for the unit i . The errors may be correlated across units. This SEM function implies our assumption, i.e., the outcome of unit i depends on the treatments and covariates of i 's neighborhoods through a function $f_Y(\cdot)$.

Based on this SEM model, we can then define summary functions s_X and s_T as well as random variables $W_i = s_{X,i}(\{X_j : A_{ij} = 1\})$ and $V_i = s_{T,i}(\{T_j : A_{ij} = 1\})$. Then, the model in Eq. (1) can be rewritten as:

$$\begin{aligned}
 X_i &= f_X[\epsilon_{X_i}], \quad i = 1, \dots, n \\
 T_i &= f_T[W_i, \epsilon_{T_i}], \quad i = 1, \dots, n \\
 Y_i &= f_Y[V_i, W_i, \epsilon_{Y_i}], \quad i = 1, \dots, n
 \end{aligned} \tag{2}$$

Eq. (2) means that the outcome Y_i depends on (\mathbf{X}, \mathbf{T}) only through summary functions $(s_{X,i}, s_{T,i})$. The summary function $s_{X,i}$ implies that the exposure and outcome of node i only depend on some function of i and its neighbors' covariates. Analogously, $s_{T,i}$ also aggregate the information for T .

3.3 Assumptions

Based on the SEM model, we make the following assumptions:

Assumption 1 (Pre-treatment network construction): Each view-specific network $A^{(m)} = G_m(X)$ is constructed solely from pre-treatment covariates X . In particular, edges are functions of the observed baseline features and do not depend on treatment T or outcome Y .

Assumption 2 (Ignorability): Error vectors $(\epsilon_{T_1}, \dots, \epsilon_{T_N})$, $(\epsilon_{Y_1}, \dots, \epsilon_{Y_N})$, and $(\epsilon_{X_1}, \dots, \epsilon_{X_N})$ are mutually independent.

Assumption 3 (Exchangeability): Errors $\epsilon_{T_1}, \dots, \epsilon_{T_N}$ and $\epsilon_{X_1}, \dots, \epsilon_{X_N}$ are identically distributed.

Assumption 4 (Local Latent Dependence Structure): The error terms $\epsilon_{X_i} \not\perp \epsilon_{X_j}$ (and analogously for $\epsilon_{Y_i}, \epsilon_{C_i}$) are permitted to be dependent if and only if nodes i and j are connected through a common neighbor, i.e., if $\exists k$ such that $A_{ik} = A_{kj} = 1$.

Assumption 5 (Positivity): For all i , $P(V_i = v \mid s_{X,i}(X) = s_{X,i}(x)) > 0$ for all x in the support of X and all v in the support of V .

In addition, the SEM model in Eq. (2) encodes the assumption that \mathbf{X} suffices to control for the confounding effect of \mathbf{T} on \mathbf{Y} . This is a version of the *conditional ignorability* assumption that is typical in *i.i.d* settings. Assumption 1 avoids post-treatment bias and ensures that the network can be regarded as fixed prior to interventions. Assumption 4 allows for the existence of latent dependency among network but limited within two hops. Assumption 5 is critical for the identifiability of treatment effects.

3.4 Problem Setting and Estimands

Our goal is to estimate the causal effects of hypothetical intervention policy T^* applied to the network. We can construct an average expected outcome that evaluates the causal effect of an observed network receiving certain interventions based on this independent relationship (Van der Laan, 2014):

$$E[\bar{Y}_n^*] = \frac{1}{n} \sum_{i=1}^n E[Y_i^*] = \frac{1}{n} \sum_{i=1}^n E[m(s_{T,i}(T^*), s_{X,i}(x))] = \frac{1}{n} \sum_{i=1}^n \sum_w m(v_i^*, w) h_i^*(v_i^*, w), \quad (3)$$

where h refers to the conditional distributions of W and V : $h_i(v | w) = P(V_i = v | W_i = w)$ and $h_i(v, w) = P(V_i = v, W_i = w)$. $m(v, w) = \sum_y y p_Y(y | v, w)$ is the conditional expectation of the outcome Y given $V = v$ and $W = w$.

The causal estimands of interest are identified by functionals of the observed data distribution $P(X, T, Y)$, which depends on N and A in our assumption. This is also known as the identifying function. It is a parameter of the observed data distribution for a network of size N , that is, a parameter of the data generating distribution that gave rise to the data at hand. It is an unknown parameter rather than an observed quantity, because the data we observe comprise a single random draw from $P(X, T, Y)$.

The identifying function of our estimands can be written as: $\mathcal{M} = \mathcal{M}_h \times \mathcal{M}_m$, where \mathcal{M} represents the class of functions that can be used to estimate the two parameters h and m . m is an initial estimation model of the potential outcome, which can be seen as a function of the treatment summary V and covariate summary W . This identifying function shows that our estimand depends on two model classes, i.e., m model class to estimate the potential outcome given a certain variable (v_i, w_i) , and a conditional density estimator h that computes the probability of a certain realization (v_i, w_i) .

4 Theoretical Analysis

4.1 Efficient Influence Curve

The efficient influence curve is a key ingredient in semi-parametric efficient estimation, which defines the linear approximation of any efficient and regular asymptotically linear estimator. Therefore, it provides an asymptotic bound for the variance of all regular asymptotically linear estimations. Under the assumptions made in Section 3.3, the influence function for the causal estimand of interest, which is denoted as ψ_N , can be written as:

$$D_N(o) = \frac{1}{n} \sum_{i=1}^n (E[m(V_i^*, W_i) | X = x] - \psi_n + \frac{\bar{h}^*(v_i, w_i)}{\bar{h}(v_i, w_i)} \{y_i - m(v_i, w_i)\}), \quad (4)$$

where $\bar{h}(v_i, w_i) = \frac{1}{n} \sum_{j=1}^n h_j(v_i, w_i)$ and $\bar{h}^*(v_i, w_i) = \frac{1}{n} \sum_{j=1}^n h_j^*(v_i, w_i)$. The target parameter ψ_n is defined as the unique value satisfying $\mathbb{E}[D_N(o)] = 0$; equivalently, $D_N(o)$ is the estimating equation whose solution identifies ψ_n .

Estimating equations for the parameters indexing a working model for m and h are stacked with the influence function estimating equation for ψ_n (Kennedy, 2016). The key idea of the targeted maximum likelihood estimation (TMLE) is to make the influence function equal to 0. The TMLE procedure consists of four steps: **Initial estimation**, which fits working models $\hat{m}(v, w)$ and $\hat{h}(v | w)$ from the observed data (W, V, Y) ; **clever covariate computation**, which computes the importance weight $H_i = \bar{h}^*(v_i, w_i) / \bar{h}(v_i, w_i)$; **targeted update** which fits the perturbation parameter $\epsilon(v)$ by minimising the influence function loss and finally **policy value estimation**. The estimator is doubly robust: Eq. equation 4 has expectation zero if either \hat{m} or \hat{h} is correctly specified, as long as the other remains bounded.

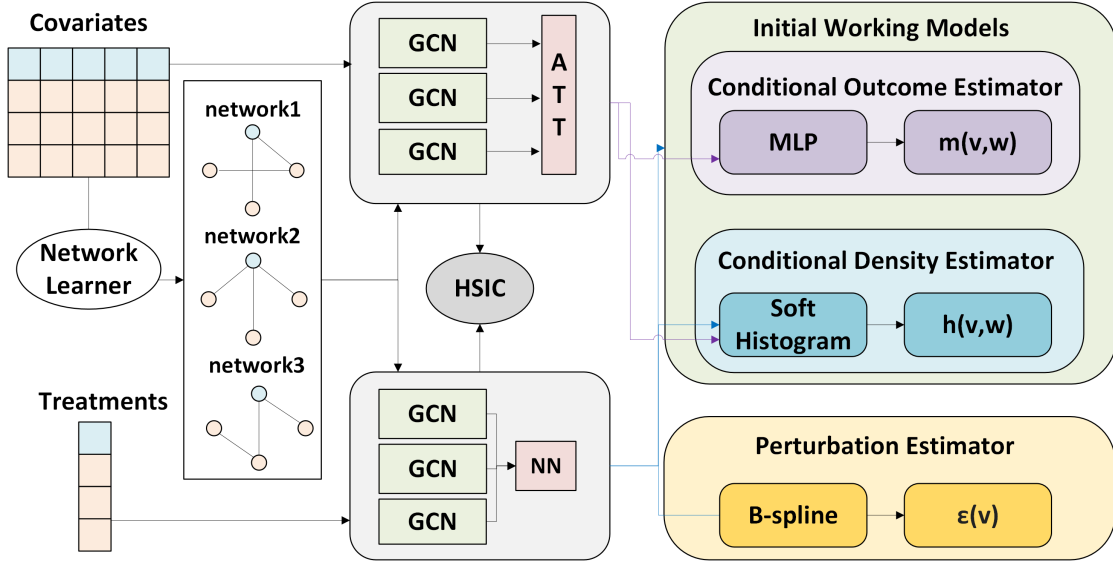


Figure 2: Illustration of our Mvdr framework. Mvdr contains four modules. **Network Learner:** Constructs $K = 3$ view-specific networks from covariates using KNN. The **representation module** learns the summary function of \mathbf{X} and aggregate the information from heterogeneous network. The **initial working model** contains one conditional outcome estimator m and a conditional density estimator h . The **perturbation estimator** is used to conduct TMLE.

The estimators with this influence function are doubly robust. The right hand side of Eq. (4) has an expected value equal to 0 if $m(\cdot)$ is replaced with an arbitrary function of V or if $h(\cdot)$ is replaced with an arbitrary function of W , as long as one of the two remains correctly specified.

Let $O = (W, V, Y)$ denote the observed variables, where W and V are aggregated summaries constructed from multi-view networks. Write $m(v, w) = \mathbb{E}[Y | V = v, W = w]$, $\bar{h}(v | w) = \Pr(V = v | W = w)$ (or density if V is continuous), and p_W the marginal law of W . We consider two common targets.

Target A (Static neighbor intervention). For a fixed v_0 in the support of V , define

$$\psi = \Psi(P) = \mathbb{E}\{m(v_0, W)\} = \int m(v_0, w) p_W(dw).$$

Under the nonparametric model $M = M_m \times M_h \times M_{p_W}$, the efficient influence function (EIF) is:

$$D^*(O) = \frac{\mathbb{1}\{V = v_0\}}{\bar{h}(v_0 | W)} \{Y - m(v_0, W)\} + m(v_0, W) - \psi. \quad (5)$$

Target B (Stochastic/dynamic neighbor intervention). Let $h^*(\cdot | w)$ be a stochastic intervention on V that may depend on W . Define

$$\psi = \Psi(P) = E_W \left[E_{V^* \sim h^*(\cdot | W)} \{m(V^*, W)\} \right] = \iint m(v, w) h^*(dv | w) p_W(dw).$$

The efficient influence function is

$$D^*(O) = \frac{\bar{h}^*(V | W)}{\bar{h}(V | W)} \{Y - m(V, W)\} + \int m(v, W) h^*(dv | W) - \psi. \quad (6)$$

When V is discrete, $\int m(v, W) h^*(dv | W) = \sum_v m(v, W) h^*(v | W)$.

The multi-view construction only affects the definitions of V and W . We now establish the large-sample properties of our estimator. Under Assumptions 1–5, the target parameter is identifiable and the TMLE

estimator is consistent and asymptotically normal. In particular, Assumptions 1–3 ensure identifiability, Assumption 5 guarantees positivity so that causal contrasts are well-defined with finite variance, and Assumption 4 restricts dependence to local neighborhoods, allowing the use of network central limit theorem arguments. As is standard in semiparametric theory, these results hold under mild regularity conditions such as bounded moments and smoothness of the relevant estimators. The detailed conditions we used to obtain the CAN of our estimator are provided in Appendix C.

Lemma 4.1 (Identification invariance under aggregation). *Let (W, V) be any Borel-measurable aggregations of $(X, T, A^{(1)}, \dots, A^{(M)})$ with W depending only on pre-treatment inputs. Under Assumptions 1–3, both targets $\psi_{\text{stat}}, \psi_{\text{stoch}}$ are identified from $P(W, V, Y)$ via*

$$\psi_{\text{stat}} = \int m(v_0, w) p_W(dw), \quad \psi_{\text{stoch}} = \iint m(v, w) h^*(dv | w) p_W(dw).$$

Theorem 4.1 (Consistency and Asymptotic Normality of TMLE). *Let $\hat{\psi}_n$ denote the TMLE for the target parameter ψ . Under Assumptions 1–3 and Regularity Conditions (1) and (2),*

$$\sqrt{C_n} (\hat{\psi}_n - \psi) \xrightarrow{d} N(0, \sigma^2),$$

where C_n is an effective sample size satisfying $n/K_{\max, n}^2 \leq C_n \leq n$. The asymptotic variance σ^2 equals the variance of the efficient influence function under the true data-generating process.

5 Methodology

In this section, we introduce our *Mvdr* framework, which estimates the causal estimands of interest discussed above. As illustrated in Figure 2, our framework consists of three components: confounder representation learning with multi-view feature balancing, density estimator, and outcome prediction. Our motivation is to find a way to integrate information from different views of the network. Inspired by the SEM framework, we propose a novel way to present the transmission of information to units through summary functions and construct them as multi-view networks. Technically, our method offers a semi-parametric way to estimate two classes of parameters h and m included in the function class \mathcal{M} that decides our model. More specifically, the outcome predictor m is a parametric model dependent on the covariate summary w and the exposure summary v , and the density estimator h is estimated through every realization of (w, v) .

5.1 Multi-view Representation of Network

Based on the assumptions of causal relationship, the outcome Y is affected by covariates X through the summary function W and by exposure T through another summary function V . In the confounder representation learning module, we use K graph convolutional network (Kipf & Welling, 2017) blocks to represent the summary function $s_{x,i}$ that aggregates information of covariates of unit i and its neighbors on multi-view networks (i.e. $X_i, A^{(k)}, k = 1, \dots, K$):

$$h_i^{(k)} = \sigma\left(\sum_{j \in N_i^{(k)}} \frac{1}{\sqrt{d_i^{(k)} d_j^{(k)}}} W_k^T x_j; \theta_k\right), \quad (7)$$

where $\sigma(\cdot)$ is a sigmoid activation function, $d_i^{(k)}$ is the degree of unit i in the k -th network, and W_k is the weight matrix of GCN parameterized by θ_k . To integrate multiple views of structural information and identify the importance of different views of the networks, we use an attention-based fusion module. Given node embeddings $h_i^{(k)} \in \mathbb{R}^d$, $k = 1, \dots, K$ from the GCNs, the scores can be computed as:

$$\alpha_{ik} = \frac{\exp(\mathbf{w}^T h_i^{(k)} + b)}{\sum_{j=1}^K \exp(\mathbf{w}^T h_i^{(j)} + b)}, \quad (8)$$

and the final multi-view representation of unit i can be written as: $w_i = \sum_{k=1}^K \alpha_{ik} \cdot h_i^{(k)}$.

For the exposure summary, we use another stack of GCN layers to learn the aggregation of neighborhood treatment, and then an MLP layer is adopted to learn the weights of exposure from different views. The final exposure mapping V is then combined with two variables: unit treatment T and multi-view neighborhood exposure S .

To ensure that the GCNs for different views of the network can learn different aspects of information, we use a **multi-view network representation discrepancy** module to maximize the information learned from the multiple GCNs. Specifically, we hope that the summary functions W_i , which is used to learn the information from covariates, and S_i , which aims to aggregate information from multiple views, can learn different information for the final estimation. In particular, we design a loss function based on the Hilbert-Schmidt independence criterion (HSIC) as follows:

$$L_{HSIC}(W, S) = \frac{1}{n^2} \text{tr}(KMLM), \quad M = \mathbb{I}_n - \frac{1}{n} \mathbf{1}_n \mathbf{1}_n^T, \quad (9)$$

where n is the number of training units, \mathbb{I}_n represents the identity matrix, and $\mathbf{1}_n$ is the vector of all ones. K and L represent the Gaussian kernel applied to W and S , respectively, i.e.,

$$K_{ij} = \exp\left(-\frac{\|w_i - w_j\|_2^2}{2}\right), \quad L_{ij} = \exp\left(-\frac{(s_i - s_j)^2}{2}\right). \quad (10)$$

5.2 Essential Predictors

As we discussed earlier, three essential parts are used to achieve our doubly robust estimator, including the conditional outcome estimator, the conditional density estimator, and the perturbation estimator.

For the **conditional outcome estimator** m , given the covariate representations W and treatment representations V obtained from the multi-view representation module, we use two multi-layer perceptron (MLP) blocks that consist of three feed-forward layers to infer the outcome y_i for the treated and control groups respectively. This potential outcome predictor is used for initial outcome prediction, representing m in the identifying function of the SEM model. Specifically, m_0 denotes the predictor for the treated units, and m_1 denotes the predictor for the controlled units. These potential outcome predictors are optimized by the mean square error (MSE) between the predicted outcomes and observed outcomes:

$$L_1 = \frac{1}{n} \sum_{i=1}^n (m_{t_i}(v_i, w_i) - y_i)^2. \quad (11)$$

The **conditional density estimator** h is the conditional distribution of the joint exposure mapping $V = (T, S)$ given covariates W , denoted by:

$$h((t, s) | w) = \Pr(T = t | W = w) \cdot f_{S|T,W}(s | t, w). \quad (12)$$

This factorization separates the discrete ego treatment T from the continuous neighbor summary S .

Ego Treatment Model. We estimate $\Pr(T = 1 | W = w)$ using an MLP classifier. Given input w , the network outputs $\hat{e}(w) = \sigma(g_\theta(w))$, where g_θ is an MLP and σ is the sigmoid function. The model is trained by minimizing the binary cross-entropy loss:

$$L_e(\theta) = -\frac{1}{n} \sum_{i=1}^n \left[T_i \log \hat{e}(W_i) + (1 - T_i) \log(1 - \hat{e}(W_i)) \right]. \quad (13)$$

Neighbor Summary Model. For the second factor $f_{S|T,W}$, we approximate the conditional density of S given (T, W) . Following the conditional histogram approach (Díaz Muñoz & Van Der Laan, 2011), we normalize v into $[0, 1]$ and partition its support into C intervals. For each interval, we fit a multinomial logistic regression model that outputs bin-wise probabilities:

$$\pi_i = \text{softmax}(W_{\text{dens}}^\top(x_i \| w_i) + b) \in \mathbb{R}^C, \quad (14)$$

where W_{dens} are learnable parameters and $\|$ denotes concatenation. The conditional density at v_i is then obtained by interpolating between the corresponding lower and upper bins:

$$\hat{h}(v_i | x_i, w_i) = \pi_{i, L_i} + (v_i \cdot C - L_i) \cdot (\pi_{i, U_i} - \pi_{i, L_i}), \quad (15)$$

where L_i and U_i denote the indices of the bins bracketing v_i . This estimator can be viewed as a flexible, data-adaptive approximation to the Radon–Nikodým derivative of the target intervention distribution with respect to the observed distribution. It is trained by minimizing the negative log-likelihood:

$$L_n = - \sum_{i=1}^n \log \hat{h}(v_i | x_i, w_i). \quad (16)$$

Putting all the two density models together, the final density model is trained by minimizing the negative log-likelihood:

$$L_2 = L_e + L_n = - \sum_{i=1}^n \log \hat{h}((x_i, s_i) | w_i). \quad (17)$$

To achieve doubly robust estimation, we use an MLP block to learn ϵ that makes the influence function equal to 0.¹ The loss function of our *perturbation estimator* is:

$$L_3 = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{m}_i - \frac{\epsilon}{\hat{h}_i})^2. \quad (18)$$

Combining the three modules for our estimation together, we have the final loss function:

$$Loss = \sum_{i=1}^N (y_i - \tilde{y}_i)^2 - \sum_{i=1}^n \log \hat{h}((x_i, s_i) | w_i) + \sum_{i=1}^n (y_i - \hat{m}_i - \frac{\epsilon}{\hat{h}_i})^2 + \lambda L_{HSIC}. \quad (19)$$

5.3 Inference under Interventions Policies

After training the initial outcome predictor m and the density estimator h , we purpose the procedure of predicting average potential outcome under certain intervention policy. Suppose our intervention towards the network is $T^* \in \{0, 1\}^n$. We first calculate the initial potential outcome $\hat{y}_i = m_{t_i^*}(v_i^*, w_i)$. Secondly, to achieve the robust estimator of the final result, we need to compute the auxiliary weights as:

$$H_i = \frac{\hat{h}^*(v_i^*, w_i)}{\hat{h}(v_i, w_i)} \quad (20)$$

For the last step, we construct the final outcome prediction with targeted learning. The final estimation for the unit i can be written as:

$$\tilde{Y}_i^* = m_{t_i^*}(v_i^*, w_i) + \epsilon(v_i^*, w_i) \cdot H_i, \quad (21)$$

and the average potential outcome under the intervention T^* is given by:

$$\bar{Y}^* = \frac{1}{n} \sum_{i=1}^n \tilde{Y}_i^*. \quad (22)$$

This estimator is doubly robust, as it will be consistent for the inference function if either h or m is correctly specified for the model.

6 Experiment

In this section, we validate the proposed method *Mvdr* on two commonly used semi-synthetic datasets. We also construct a new semi-synthetic dataset based on the DBLP dataset. We verify the effectiveness of our method and further evaluate the correctness of our analysis using these semi-synthetic datasets. In particular, we aim to answer the following research questions (RQ):

- RQ1: How does the proposed method compare with existing methods in terms of effect estimation performance?
- RQ2: How does the perturbation estimator module affect the performance of our methods?
- RQ3: Does our method stably perform well under different types of misspecification?

¹Standard TMLE for finite-dimensional parameters fits a scalar ϵ . In our setting, the exposure summary $V = (T, S)$ contains a continuous component S , so the clever covariate varies with v . Following Nie et al. (2021) and the continuous-treatment TMLE literature, we parameterize $\epsilon(v)$ as a B-spline network to flexibly cover the continuous support of V . When V is discrete (static intervention, Target A), ϵ reduces to a finite-dimensional vector, consistent with the standard formulation.

In the following, we first introduce the experimental setup and then answer the questions above by conducting the corresponding experiments.

Datasets. In our experiment, we utilize two widely used semi-synthetic datasets, i.e., *BlogCatalog* and *Flickr*. Another contribution of this paper is that we construct a new semi-synthetic dataset based on the *DBLP* dataset. The DBLP dataset contains rich information of conference and journal publications, such as detailed information of authors. We construct a coauthorship network by matching the author information and the publication information. Detailed statistics of three datasets are provided in Appendix B.

Multi-view network construction. A simple idea of constructing multi-view networks is to compute the similarity scores of covariates between units and construct the relational networks based on their similarity. We follow the multi-view network construction method from (Lin et al., 2023), in which multi-view networks were constructed for the BlogCatalog and Flickr datasets based on the similarity of covariates. The constructed latent networks can then be potential for the upcoming multi-view information aggregation. Notice that these latent multi-view networks are learned from the covariates and are not directly observed from the network data.

Since it is impossible to observe the counterfactual outcomes, we need to mimic the unknown output. Following prior studies on ITE and ATE, we simulate outcomes as ground-truth values for counterfactual outcomes that are not available using the following equation:

$$Y_i = \beta_T \cdot T_i + \beta_X \cdot (\mathbf{w}^\top \mathbf{x}_i) + \beta_{NX} \cdot \sum_{k=1}^3 \frac{1}{|\mathcal{N}_i^{(k)}|} \sum_{j \in \mathcal{N}_i^{(k)}} \mathbf{w}^\top \mathbf{x}_j + \beta_{NT} \cdot \left(\sum_{k=1}^3 \alpha_k \cdot \frac{\sum_{j \in \mathcal{N}_i^{(k)}} T_j}{|\mathcal{N}_i^{(k)}|} \right) + \varepsilon_i. \quad (23)$$

This construction method assumes that the exposure mapping is a simple function of the ratio of treated neighbors across views. Here, $\mathbf{x}_i \in \mathbb{R}^d$ denotes the covariate vector for unit i , and $T_i \in \mathbb{R}$ is the assigned treatment. The neighborhood sets $\mathcal{N}_i^{(k)}$ are derived from three adjacency matrices $A^{(1)}, A^{(2)}, A^{(3)}$, representing multi-view network structures. Each network layer $k \in \{1, 2, 3\}$ contributes to the potential outcome through aggregated covariate signals and treatment exposures, weighted by coefficients α_k . We set $\alpha_1 = \alpha_2 = 1$ and $\alpha_3 = 0.5$ in our experiments to reflect the varying network influence. Lastly, $\varepsilon_i \sim \mathcal{N}(0, 1)$ introduces stochastic variation.

6.1 Three Types of Interventions

To evaluate the stability of our method and baselines under different types of interventions, we design three interventions to mimic different hypothetical treatment policies in real life. **Static intervention** assigns all the units in the network as treated or not treated. **Dynamic intervention** assigns exposures as a user-specified, deterministic function of covariates. This scenario often appears when there are some constraints in resources. For example, the economic incentive was resource constrained and could only be allocated to up to 10% of the whole group. In our experiments, we allocate the top 10% most connected members of the community. **Stochastic intervention** assigns exposures as a user-specified, random function. In our experiments, we assign each unit to treatment with a constant probability of 0.35.

6.2 Baselines and Metrics

Baselines. We compare our method with several representative baselines in causal inference, including methods that deal with a single observed network and multi-view networks. CFR (Johansson et al., 2023) and ND (Guo et al., 2019) are modified with additionally inputting the exposure. TARNet (Shalit et al., 2017) has a similar model architecture as CFR but removes the balance term. [For fair comparison, we augment CFR, ND, and TARNet with the exposure summary \$z_i = \frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i} T_j\$ as an additional input feature, denoted CFR+z, ND+z, and TARNet+z respectively.](#) HINITE (Lin et al., 2023) considers several observed heterogeneous networks. NetEst (Jiang & Sun, 2022) is designed for treatment effect estimation under interference. TNet (Chen et al., 2024) uses the proportion of treated neighbors as the exposure mapping, which is a doubly robust method.

Evaluation Metric. The main purpose of our model is to predict the average expected outcome of hypothetical intervention policies. We evaluate the accuracy of the model predicting the average expected

Table 1: Comparison of methods on causal effect of network estimand with different interventions. ‘Static0’ sets all units to control; ‘Static1’ sets all to treated; ‘Stochastic’ applies treatment to 35% randomly; ‘Dynamic’ assigns treatment based on covariate-dependent probabilities. Best performance in each row is **bolded**. Results are reported as the mean and standard deviation over five runs for each setting.

Dataset	Intervention	CFR	CFR+z	ND	ND+z	TARNet	TARNet+z	NetEst	HINITE	TNet	Mvdr(w/o. \mathcal{L}_3)	Mvdr	
BlogCata	Static0	0.442 _{0.2618}	0.403 _{0.2625}	0.4466 _{0.2508}	0.4093 _{0.2576}	0.4457 _{0.2442}	0.4082 _{0.2726}	0.3861 _{0.2576}	1.0292 _{0.7670}	16.346 _{17.3932}	0.4968 _{0.2366}	0.4891 _{0.1899}	
	Static1	0.5381 _{0.1829}	0.4302 _{0.2203}	0.5766 _{0.1959}	0.3953 _{0.2357}	0.5873 _{0.2077}	0.5045 _{0.1933}	0.8412 _{0.2356}	1.2021 _{0.9884}	7999.1908 _{12974.5585}	0.3035 _{0.1633}	0.257 _{0.1515}	
	Stochastic	0.5833 _{0.2865}	0.5747 _{0.2845}	0.5821 _{0.3006}	0.5797 _{0.2920}	0.5881 _{0.3175}	0.5813 _{0.3119}	0.2751 _{0.1782}	0.6227 _{0.2099}	81.2073 _{123.1866}	0.2818 _{0.1699}	0.2668 _{0.1847}	
	Dynamic	0.4858 _{0.2952}	0.4869 _{0.3001}	0.4910 _{0.3025}	0.4974 _{0.3033}	0.4919 _{0.3108}	0.5065 _{0.3007}	0.5001 _{0.3032}	0.5585 _{0.2657}	23.1303 _{47.0434}	0.4480 _{0.2433}	0.4401 _{0.3636}	
	Out-of-Sample	Static0	0.4251 _{0.2345}	0.3723 _{0.2388}	0.4316 _{0.1992}	0.3283 _{0.2780}	0.4758 _{0.2521}	0.5224 _{0.2921}	0.4005 _{0.348}	0.9796 _{0.7588}	15.6192 _{14.4018}	0.4509 _{0.2701}	0.4174 _{0.3391}
	Static1	0.5011 _{0.1543}	0.3958 _{0.2316}	0.4179 _{0.4720}	0.5031 _{0.3988}	0.4843 _{0.3408}	0.3502 _{0.3040}	0.7291 _{0.3454}	1.1732 _{0.9409}	5735.9332 _{2219.6814}	0.2940 _{0.2183}	0.2279 _{0.1549}	
Flickr	Static0	0.4759 _{0.2916}	0.4590 _{0.2745}	0.4478 _{0.3245}	0.4914 _{0.3420}	0.5313 _{0.3889}	0.5067 _{0.4134}	0.4446 _{0.3048}	0.5065 _{0.3123}	112.701 _{133.1668}	0.1974 _{0.1821}	0.1847 _{0.1784}	
	Static1	0.6421 _{0.3101}	0.6385 _{0.3204}	0.6324 _{0.3334}	0.5830 _{0.3148}	0.6721 _{0.3047}	0.6414 _{0.3250}	0.6623 _{0.3347}	0.7148 _{0.3062}	18.7178 _{9.9793}	0.4785 _{0.3788}	0.4500 _{0.3886}	
	Stochastic	0.1710 _{0.0766}	0.0887 _{0.0988}	0.1747 _{0.0768}	0.1002 _{0.0231}	0.1816 _{0.0700}	0.1700 _{0.0676}	0.1014 _{0.0717}	0.3486 _{0.3168}	5.9296 _{6.5908}	0.1311 _{0.0371}	0.1086 _{0.0996}	
	Dynamic	0.1751 _{0.1185}	0.0994 _{0.0927}	0.2522 _{0.1844}	0.0923 _{0.0613}	0.2683 _{0.1961}	0.1347 _{0.0902}	0.4473 _{0.0796}	1.1022 _{0.8834}	18.1979 _{6.6327}	0.3878 _{0.0494}	0.3788 _{0.0658}	
	Out-of-Sample	Static0	0.1410 _{0.0777}	0.1368 _{0.0967}	0.1647 _{0.0974}	0.1908 _{0.1021}	0.1779 _{0.1059}	0.2181 _{0.0993}	0.3243 _{0.0986}	0.2045 _{0.1504}	10.8568 _{11.5352}	0.1995 _{0.0742}	0.1316 _{0.0375}
	Static1	0.1793 _{0.1497}	0.2085 _{0.1607}	0.2109 _{0.1635}	0.2305 _{0.0678}	0.2279 _{0.1714}	0.1905 _{0.0563}	0.5980 _{0.1442}	0.1436 _{0.1705}	2934.7736 _{3022.9666}	0.0401 _{0.041}	0.0338 _{0.0325}	
DBLP	Static0	0.1241 _{0.0905}	0.1407 _{0.1073}	0.1419 _{0.0812}	0.1391 _{0.1172}	0.1642 _{0.0788}	0.1289 _{0.1133}	0.1376 _{0.0892}	0.3111 _{0.2683}	14.9494 _{19.1093}	0.2731 _{0.0744}	0.2163 _{0.0486}	
	Static1	0.1927 _{0.1098}	0.1565 _{0.1069}	0.2960 _{0.1432}	0.2018 _{0.0665}	0.3354 _{0.1231}	0.2048 _{0.0773}	0.2352 _{0.0867}	1.0144 _{0.8732}	26.4602 _{21.5615}	0.6340 _{0.1683}	0.4648 _{0.1063}	
	Stochastic	0.1424 _{0.1079}	0.1598 _{0.0834}	0.2542 _{0.1024}	0.1864 _{0.0732}	0.1471 _{0.1169}	0.2180 _{0.1140}	0.3173 _{0.0793}	0.2248 _{0.1756}	4.875 _{5.1017}	0.1245 _{0.0469}	0.1172 _{0.0718}	
	Dynamic	0.1642 _{0.1354}	0.1703 _{0.1496}	0.2032 _{0.1661}	0.2905 _{0.1676}	0.2036 _{0.1606}	0.1293 _{0.0574}	0.7368 _{0.1807}	0.3059 _{0.2397}	2141.9216 _{912.0244}	0.0470 _{0.0152}	0.0402 _{0.0287}	
	Out-of-Sample	Static0	0.6801 _{0.6428}	0.7406 _{0.6214}	0.8354 _{0.5585}	0.6759 _{0.5753}	0.9776 _{0.8482}	0.7069 _{0.6272}	0.9383 _{0.593}	0.8162 _{0.5422}	7.2816 _{17.7863}	0.7142 _{0.4798}	0.5933 _{0.5328}
	Static1	0.6128 _{0.4203}	0.6682 _{0.4753}	0.6194 _{0.3556}	0.7100 _{0.6311}	0.5597 _{0.3182}	0.6378 _{0.4699}	0.697 _{0.2554}	1.4879 _{1.1256}	25.8228 _{50.5348}	1.1797 _{0.5966}	1.1800 _{0.5476}	

Table 2: Stochastic Policy Evaluation. The proportion of treated units ranges from 0.20 to 0.80, reflecting real-world policy settings where the treatment size may fluctuate according to budget constraints. Best performance in each row is **bolded**.

Dataset	Proportion	CFR	CFR+z	ND	ND+z	TARNet	TARNet+z	NetEst	HINITE	TNet	Mvdr
BlogCata	0.20	0.3197 _{0.305}	0.3562 _{0.3169}	0.5202 _{0.4264}	0.5617 _{0.4721}	0.4834 _{0.364}	0.4491 _{0.3787}	0.4672 _{0.4064}	0.6987 _{0.2483}	57.2356 _{57.4217}	0.3016 _{0.2646}
	0.50	0.2598 _{0.2665}	0.2552 _{0.261}	0.3095 _{0.2291}	0.3846 _{0.2485}	0.3743 _{0.3218}	0.2980 _{0.3026}	0.3130 _{0.4055}	0.5063 _{0.2674}	8.5232 _{11.3317}	0.2948 _{0.2602}
	0.65	0.4678 _{0.3421}	0.494 _{0.3619}	0.4957 _{0.2699}	0.6111 _{0.3331}	0.4117 _{0.2646}	0.4747 _{0.2324}	0.5033 _{0.3818}	1.0180 _{0.5412}	7.3843 _{6.2696}	0.3990 _{0.3953}
	0.80	0.2743 _{0.175}	0.2728 _{0.1868}	0.2859 _{0.2169}	0.3844 _{0.2375}	0.2581 _{0.1462}	0.3422 _{0.1296}	0.4369 _{0.3256}	0.9691 _{0.6849}	3033.8552 _{2888.6561}	0.2300 _{0.1513}
Flickr	0.20	0.0956 _{0.0763}	0.1181 _{0.0968}	0.1745 _{0.0364}	0.1864 _{0.0732}	0.0992 _{0.0541}	0.1294 _{0.055}	0.1724 _{0.0824}	0.4369 _{0.3551}	320.0 _{634.0}	0.0872 _{0.0694}
	0.50	0.1150 _{0.0419}	0.1161 _{0.0507}	0.1455 _{0.0486}	0.1277 _{0.0537}	0.1414 _{0.0649}	0.1293 _{0.0574}	0.2080 _{0.1165}	0.1928 _{0.2158}	320.5 _{545.6}	0.0914 _{0.0688}
	0.65	0.1028 _{0.0692}	0.0966 _{0.0706}	0.1118 _{0.0111}	0.1002 _{0.0527}	0.1071 _{0.0273}	0.1036 _{0.026}	0.2436 _{0.0739}	0.4003 _{0.3475}	333.2976 _{295.757}	0.0908 _{0.0614}
	0.80	0.1196 _{0.0886}	0.0933 _{0.0495}	0.1781 _{0.0703}	0.137 _{0.0744}	0.2265 _{0.110}	0.2008 _{0.0558}	0.5136 _{0.1013}	0.5119 _{0.3149}	1073.9346 _{1198.002}	0.0850 _{0.0700}
DBLP	0.20	0.6884 _{0.6434}	0.4900 _{0.6684}	0.5301 _{0.6866}	0.4749 _{0.7014}	0.5244 _{0.6407}	0.5582 _{0.7820}	0.6222 _{0.7431}	0.8006 _{0.7482}	121.7685 _{240.3945}	0.4284 _{0.5858}
	0.50	0.4794 _{0.3694}	0.4749 _{0.3688}	0.7325 _{0.5254}	0.5381 _{0.4129}	0.5989 _{0.4951}	0.5228 _{0.3430}	0.5057 _{0.3798}	0.8234 _{0.6519}	3.0800 _{4.7400}	0.4363 _{0.3654}
	0.65	0.5841 _{0.5569}	0.6119 _{0.5669}	0.5515 _{0.2723}	0.5838 _{0.6110}	0.5397 _{0.2044}	0.5593 _{0.4905}	0.5499 _{0.4583}	1.4967 _{1.0754}	17.7013 _{4.764}	0.5140 _{0.4782}
	0.80	0.6107 _{0.6391}	0.6028 _{0.6532}	0.7433 _{0.2997}	0.5800 _{0.6448}	0.7180 _{0.2523}	0.6690 _{0.5468}	0.6785 _{0.3982}	1.7752 _{1.2415}	9.3042 _{7.7427}	0.4026 _{0.5439}

outcome of a group receiving treatment under different intervention policies, written as:

$$\varepsilon_{intervention} = |\widetilde{Y}_n^* - Y_n^*|, \text{ where } Y_n^* = \frac{1}{n} \sum_{i=1}^n Y_i^*.$$

6.3 Results and Ablation Studies

Comparisons with Baselines. Results in Table 1 show that our Mvdr method outperforms baselines in most cases, especially for the stochastic and dynamic policy settings, where the treatment assignment is related to units’ covariates and neighborhood exposure. Under static interventions, strong i.i.d. baselines can be competitive or best, suggesting that static policies are less sensitive to misspecification of the exposure mechanism. We also report the results of some other metrics in Appendix D.

The Perturbation Estimator. As shown in Table 1, we have conducted experiments to compare the performance between Mvdr and Mvdr(w/o. \mathcal{L}_3) across the three datasets. Overall, Mvdr outperforms Mvdr(w/o. \mathcal{L}_3), which is as expected. Results show that the TMLE framework introduces double robustness to Mvdr and mitigates bias through the three modules. When the initial outcome model includes bias, the conditional density works as a re-biasing step to produce robust results.

The Stochastic Policy. We evaluate the performance when the treated proportion of stochastic assignment policy changes from 0.20 to 0.80. As shown in Table 2, Mvdr outperforms baselines across different assignment policy settings.




Figure 3: Evaluation on bias caused by misspecification on treatment mechanism.




Figure 4: Evaluation on bias caused by misspecification on outcome model.

Unbiasness under Misspecification. To show that our model can perform stably when either the working model m and h is misspecified, we simulate data accordingly to mimic the different misspecification problem. (1) To model the misspecification of treatment mechanism, we include data that randomly flip treatment with rate from 0.25 to 1. (2) For the misspecification of outcome mechanism, we include higher nonlinearity by introducing the tanh function toward covariates to outcome function. This will lead to the exposure summary being no longer sufficient for outcome prediction but still following the treatment generation mechanism. Figure 3 and Figure 4 show that the prediction with misspecification is close to the original results.

7 Discussions and Conclusions

This paper proposes a framework to estimate the average expected outcome on an observed heterogeneous network with latent network dependency through a targeted maximum likelihood estimation method. We provided theorems showing the consistency and double robustness. Experiments show that our method is able to perform well on estimating average expected outcome under hypothetical intervention policies on networks. Our method aims to solve the problem of latent network dependency in causal inference under heterogeneous network setting, by modeling potential outcome as conditional expectation of two network related variable, W and V . Though it achieves relatively good performance, we acknowledge that this method has its limitations. For example, the exposure mapping, where we limited its dimension to 1, may not be sufficient to cover the interference mechanism in complex network settings. We plan to explore other semi-parametric methods to model the exposure mapping and achieve better identification of interference in our future work.

References

- Shishir Adhikari and Elena Zheleva. Inferring individual direct causal effects under heterogeneous peer influence. *Machine Learning*, 114(4), March 2025. ISSN 1573-0565. doi: 10.1007/s10994-024-06729-2. URL <http://dx.doi.org/10.1007/s10994-024-06729-2>.
- Peter M Aronow and Cyrus Samii. Estimating average causal effects under general interference, with application to a social network experiment. 2017.
- Laura B Balzer, Wenjing Zheng, Mark J van der Laan, and Maya L Petersen. A new approach to hierarchical data analysis: Targeted maximum likelihood estimation for the causal effect of a cluster-level exposure. *Statistical methods in medical research*, 28(6):1761–1780, 2019.
- Ariel Chao, Donna Spiegelman, Ashley Buchanan, and Laura Forastiere. Estimation and inference for causal spillover effects in egocentric-network randomized trials in the presence of network membership misclassification. *Biostatistics*, 26(1):kxaf009, 2025.
- David Chen, Maya L Petersen, Helene Charlotte Rytgaard, Randi Grøn, Theis Lange, Søren Rasmussen, Richard E Pratley, Steven P Marso, Kajsa Kvist, John Buse, et al. Beyond the cox hazard ratio: a targeted learning approach to survival analysis in a cardiovascular outcome trial application. *Statistics in Biopharmaceutical Research*, 15(3):524–539, 2023.
- Weilin Chen, Ruichu Cai, Zeqin Yang, Jie Qiao, Yuguang Yan, Zijian Li, and Zhifeng Hao. Doubly robust causal effect estimation under networked interference via targeted learning, 2024. URL <https://arxiv.org/abs/2405.03342>.
- Iván Díaz Muñoz and Mark J. Van Der Laan. Super Learner Based Conditional Density Estimation with Application to Marginal Structural Models. *The International Journal of Biostatistics*, 7(1):1–20, January 2011. ISSN 1557-4679. doi: 10.2202/1557-4679.1356. URL <https://www.degruyter.com/document/doi/10.2202/1557-4679.1356/html>.
- Naoki Egami. Spillover effects in the presence of unobserved networks. *Political Analysis*, 29(3):287–316, 2021.
- Laura Forastiere, Fabrizia Mealli, Albert Wu, and Edoardo M Airoidi. Estimating causal effects under network interference with bayesian generalized propensity scores. *Journal of Machine Learning Research*, 23(289):1–61, 2022.
- Ruocheng Guo, Jundong Li, and Huan Liu. Learning Individual Causal Effects from Networked Observational Data, December 2019. URL <http://arxiv.org/abs/1906.03485>. arXiv:1906.03485 [cs].
- M Elizabeth Halloran and Claudio J Struchiner. Causal inference in infectious diseases. *Epidemiology*, pp. 142–151, 1995.
- Ruining He and Julian McAuley. Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. In *Proceedings of the 25th International Conference on World Wide Web, WWW '16*, pp. 507–517, Republic and Canton of Geneva, CHE, 2016. International World Wide Web Conferences Steering Committee. ISBN 9781450341431. doi: 10.1145/2872427.2883037. URL <https://doi.org/10.1145/2872427.2883037>.
- Tadao Hoshino, , and Takahide Yanagi. Causal Inference with Noncompliance and Unknown Interference. *Journal of the American Statistical Association*, 119(548):2869–2880, October 2024. ISSN 0162-1459. doi: 10.1080/01621459.2023.2284413. URL <https://doi.org/10.1080/01621459.2023.2284413>. Publisher: ASA Website _eprint: <https://doi.org/10.1080/01621459.2023.2284413>.
- Michael G Hudgens and M Elizabeth Halloran. Toward causal inference with interference. *Journal of the american statistical association*, 103(482):832–842, 2008.

- Song Jiang and Yizhou Sun. Estimating Causal Effects on Networked Observational Data via Representation Learning. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pp. 852–861, Atlanta GA USA, October 2022. ACM. ISBN 978-1-4503-9236-5. doi: 10.1145/3511808.3557311. URL <https://dl.acm.org/doi/10.1145/3511808.3557311>.
- Fredrik D. Johansson, Uri Shalit, Nathan Kallus, and David Sontag. Generalization bounds and representation learning for estimation of potential outcomes and causal effects, 2023. URL <https://arxiv.org/abs/2001.07426>.
- Vishesh Karwa and Edoardo M Airoidi. A systematic investigation of classical causal inference strategies under mis-specification due to network interference. *arXiv preprint arXiv:1810.08259*, 2018.
- Edward H. Kennedy. Semiparametric theory and empirical processes in causal inference, 2016. URL <https://arxiv.org/abs/1510.04740>.
- Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks, 2017. URL <https://arxiv.org/abs/1609.02907>.
- Noémi Kreif, Linh Tran, Richard Grieve, Bianca De Stavola, Robert C Tasker, and Maya Petersen. Estimating the comparative effectiveness of feeding interventions in the pediatric intensive care unit: a demonstration of longitudinal targeted maximum likelihood estimation. *American journal of epidemiology*, 186(12):1370–1379, 2017.
- Michael P Leung. Causal inference under approximate neighborhood interference. *Econometrica*, 90(1): 267–293, 2022.
- Shuangning Li and Stefan Wager. Random graph asymptotics for treatment effect estimation under network interference. *The Annals of Statistics*, 50(4):2334–2358, 2022.
- Xiaofeng Lin, Guoxi Zhang, Xiaotian Lu, Han Bao, Koh Takeuchi, and Hisashi Kashima. Estimating Treatment Effects Under Heterogeneous Interference, September 2023. URL <http://arxiv.org/abs/2309.13884>. arXiv:2309.13884 [cs].
- Jing Ma, Ruocheng Guo, Chen Chen, Aidong Zhang, and Jundong Li. Deconfounding with Networked Observational Data in a Dynamic Environment. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, pp. 166–174, Virtual Event Israel, March 2021. ACM. ISBN 978-1-4503-8297-7. doi: 10.1145/3437963.3441818. URL <https://dl.acm.org/doi/10.1145/3437963.3441818>.
- Jing Ma, Mengting Wan, Longqi Yang, Jundong Li, Brent Hecht, and Jaime Teevan. Learning Causal Effects on Hypergraphs, July 2022. URL <http://arxiv.org/abs/2207.04049>. arXiv:2207.04049 [cs].
- Yunpu Ma and Volker Tresp. Causal Inference under Networked Interference and Intervention Policy Enhancement, May 2021. URL <http://arxiv.org/abs/2002.08506>. arXiv:2002.08506 [stat].
- Charles F Manski. Identification of treatment response with social interactions. *The Econometrics Journal*, 16(1):S1–S23, 2013.
- Lizhen Nie, Mao Ye, Qiang Liu, and Dan Nicolae. VCNET AND FUNCTIONAL TARGETED REGULARIZATION FOR LEARNING CAUSAL EFFECTS OF CONTINUOUS TREATMENTS. 2021.
- Elizabeth L Ogburn, Ilya Shpitser, and Youjin Lee. Causal inference, social networks and chain graphs. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 183(4):1659–1676, 2020.
- Elizabeth L. Ogburn, Oleg Sofrygin, Ivan Diaz, and Mark J. van der Laan. Causal inference for social network data, 2022. URL <https://arxiv.org/abs/1705.08527>.
- Judea Pearl. The causal foundations of structural equation modeling. *Handbook of structural equation modeling*, pp. 68–91, 2012.

- Fredrik Sävje. Causal inference with misspecified exposure mappings: separating definitions and assumptions. *Biometrika*, 111(1):1–15, 2024.
- Uri Shalit, Fredrik D. Johansson, and David Sontag. Estimating individual treatment effect: generalization bounds and algorithms, 2017. URL <https://arxiv.org/abs/1606.03976>.
- Cosma Rohilla Shalizi and Andrew C. Thomas. Homophily and contagion are generically confounded in observational social network studies. *Sociological Methods & Research*, 40(2):211–239, May 2011. ISSN 1552-8294. doi: 10.1177/0049124111404820. URL <http://dx.doi.org/10.1177/0049124111404820>.
- Claudia Shi, David Blei, and Victor Veitch. Adapting neural networks for the estimation of treatment effects. *Advances in neural information processing systems*, 32, 2019.
- Michael E Sobel. What do randomized studies of housing mobility demonstrate? causal inference in the face of interference. *Journal of the American Statistical Association*, 101(476):1398–1407, 2006.
- Eric J. Tchetgen Tchetgen, Isabel Fulcher, and Ilya Shpitser. Auto-g-computation of causal effects on a network, 2019. URL <https://arxiv.org/abs/1709.01577>.
- Panos Toulis and Edward Kao. Estimation of causal peer influence effects. In Sanjoy Dasgupta and David McAllester (eds.), *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pp. 1489–1497, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR. URL <https://proceedings.mlr.press/v28/toulis13.html>.
- Mark J Van der Laan. Causal inference for a population of causally connected units. *Journal of Causal Inference*, 2(1):13–74, 2014.
- Mark J. van der Laan. Causal Inference for a Population of Causally Connected Units. *Journal of causal inference*, 2(1):13–74, March 2014. ISSN 2193-3677. doi: 10.1515/jci-2013-0002. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4500386/>.
- Mark J van der Laan and Susan Gruber. Targeted minimum loss based estimation of causal effects of multiple time point interventions. *The international journal of biostatistics*, 8(1), 2012.
- Mark J Van Der Laan and Daniel Rubin. Targeted maximum likelihood learning. *The international journal of biostatistics*, 2(1), 2006.
- Mark J Van der Laan, Sherri Rose, et al. *Targeted learning: causal inference for observational and experimental data*, volume 4. Springer, 2011.
- Richard A Watson, Hengrui Cai, Xinming An, Samuel McLean, and Rui Song. On heterogeneous treatment effects in heterogeneous causal graphs, 2023. URL <https://arxiv.org/abs/2301.12383>.
- Kun Zhang, Biwei Huang, Jiji Zhang, Clark Glymour, and Bernhard Schölkopf. Causal discovery from non-stationary/heterogeneous data: Skeleton estimation and orientation determination. In *IJCAI: Proceedings of the Conference*, volume 2017, pp. 1347, 2017.

A Notations

Symbol	Type / domain	Description
<i>Units, data, and network</i>		
$i \in \{1, \dots, n\}$	index	Unit (node) index; n is network size.
x_i	\mathbb{R}^d	Covariates of unit i (column vector).
t_i	$\{0, 1\}$	Binary treatment of unit i .
y_i	\mathbb{R}	Observed outcome for unit i under t_i .
X	$\mathbb{R}^{n \times d}$	Covariate matrix $X = (x_1, \dots, x_n)^\top$.
T	$\{0, 1\}^n$	Treatment vector $T = (t_1, \dots, t_n)^\top$.
Y	\mathbb{R}^n	Outcome vector $Y = (y_1, \dots, y_n)^\top$.
G	graph on n nodes	Observed (aggregated) network.
A	$\{0, 1\}^{n \times n}$	Adjacency matrix of G (symmetrized). $A_{ij} = 1$ iff edge (i, j) exists.
$A^{(m)}$	$\{0, 1\}^{n \times n}$	Adjacency of the m -th view/layer, $m = 1, \dots, K$.
$N_i^{(m)}$	subset of $\{1, \dots, n\}$	Neighbor set of node i in view m .
$d_i^{(m)}$	\mathbb{N}	Degree of node i in view m (in $A^{(m)}$).
<i>Structural equation model (SEM)</i>		
f_X, f_T, f_Y	functions	Unknown structural functions in the SEM.
$\varepsilon_{X_i}, \varepsilon_{T_i}, \varepsilon_{Y_i}$	r.v.	Exogenous errors for unit i (may be dependent across units).
$s_{X,i}$	summary map	Summary of $\{X_j : A_{ij} = 1\}$ for unit i .
$s_{T,i}$	summary map	Summary of $\{T_j : A_{ij} = 1\}$ for unit i .
W_i	r.v.	Covariate summary $W_i = s_{X,i}(\{X_j : A_{ij} = 1\})$.
S_i	r.v.	Neighbor exposure summary of treatments (continuous or discrete).
V_i	r.v.	Joint exposure summary, typically $V_i = (T_i, S_i)$.
<i>Statistical models and conditionals</i>		
$m(v, w)$	\mathbb{R}	Outcome regression: $m(v, w) = \mathbb{E}[Y \mid V = v, W = w]$.
$h(v \mid w)$	density / pmf	Exposure mechanism: $h(v \mid w) = \mathbb{P}(V = v \mid W = w)$.
$h(v, w)$	joint	Joint law of (V, W) ; $h(v, w) = \mathbb{P}(V = v, W = w)$.
$h^*(\cdot \mid w)$	intervention kernel	Target (stochastic) intervention on V given $W = w$.
p_W	law on \mathcal{W}	Marginal distribution of W .
O	tuple	Observed variables $O = (W, V, Y)$.
$\bar{h}(\cdot)$	average	Node-wise average: e.g. $\bar{h}(v, w) = \frac{1}{n} \sum_{j=1}^n h_j(v, w)$.
$\bar{h}^*(\cdot)$	average under h^*	$\bar{h}^*(v, w) = \frac{1}{n} \sum_{j=1}^n h_j^*(v, w)$.
<i>Targets / estimands</i>		
T^*	$\{0, 1\}^n$	Hypothetical intervention assignment over the network.
V^*	r.v.	Exposure under intervention, e.g. $V^* \sim h^*(\cdot \mid W)$.
v_0	value in $\text{supp}(V)$	Fixed (static) neighbor exposure level.
ψ	\mathbb{R}	Target parameter (e.g., $\mathbb{E}[m(v_0, W)]$ or $\mathbb{E}_W \int m(v, W) h^*(dv \mid W)$).
H_i	weight	“Clever covariate”: $\frac{\bar{h}^*(V_i \mid W_i)}{\bar{h}(V_i \mid W_i)}$ (stochastic) or $\frac{\mathbf{1}\{V_i=v_0\}}{\bar{h}(v_0 \mid W_i)}$ (static).
$D^*(O)$	r.v.	Efficient influence function (EIF) for ψ .
<i>Asymptotics / dependence</i>		
$K_{\max, n}$	\mathbb{N}	Maximum degree of (aggregated) network.
C_n	\mathbb{N}	Effective sample size with $n/K_{\max, n}^2 \leq C_n \leq n$.

B Data Statics

Dataset	Graph	Nodes	Edges	AvgDeg	StdDeg	MinDeg	MaxDeg	Density	Components
BC	A1	5196	171743	66	54.83	5	769	0.012725	1
	A2	5196	32036	12	1.77	10	20	0.002374	1
	A3	5196	34957	13	2.12	10	23	0.002590	1
Flickr	A1	7301	238086	65	132.66	2	1834	0.008934	1
	A2	7301	45727	12	2.07	10	30	0.001716	1
	A3	7301	48203	13	2.18	10	32	0.001809	1
DBLP	A1	2828	3988527	2820	91.90	2	2824	0.997786	2
	A2	2828	17868	12	2.24	10	23	0.004470	7
	A3	2828	17615	12	1.95	10	22	0.004407	1

C Theory

Lemma C.1 (Identification invariance under aggregation). *Let (W, V) be any Borel-measurable aggregations of $(X, T, A^{(1)}, \dots, A^{(M)})$ with W depending only on pre-treatment inputs. Under Assumptions 1–3, both targets $\psi_{\text{stat}}, \psi_{\text{stoch}}$ are identified from $P(W, V, Y)$ via*

$$\psi_{\text{stat}} = \int m(v_0, w) p_W(dw), \quad \psi_{\text{stoch}} = \iint m(v, w) h^*(dv | w) p_W(dw).$$

Proof. Multiple view-specific networks $\{A^{(m)} : m = 1, \dots, M\}$ are constructed from baseline covariates. For each unit i , we form view-specific summaries

$$W_i^{(m)} = s_C^{(m)}(\{C_j : A_{ij}^{(m)} = 1\}), \quad V_i^{(m)} = s_X^{(m)}(\{X_j : A_{ij}^{(m)} = 1\}),$$

and aggregate them into

$$W_i = \text{Agg}_W(W_i^{(1)}, \dots, W_i^{(M)}), \quad V_i = \text{Agg}_V(V_i^{(1)}, \dots, V_i^{(M)}).$$

Importantly, both constructions produce (V_i, W_i) that are functions of baseline covariates C and treatment X of neighbors. As such, the statistical model remains nonparametric in (V, W) , and the target parameter is defined entirely through $m(v, w) = E[Y | V = v, W = w]$ and $h(v | w) = \Pr(V = v | W = w)$. Therefore the pathwise derivative and canonical gradient retain the same form given in Equations equation 5–equation 6. Only the dimension and support of (V, W) change under different aggregation rules; the EIF expressions remain unchanged. \square

Condition 1 (Bounded dependence / degree growth). *The maximum degree of the (aggregated) network grows slowly with sample size:*

$$K_{\max, n}^2 / n \rightarrow 0.$$

This ensures that each unit’s dependence neighborhood does not become too large relative to the overall sample, permitting a central limit theorem under Stein’s method (Ogburn et al., 2022).

Condition 2. *Either the outcome regression $m(v, w)$ or the exposure mechanism $h(v | w)$ is consistently estimated at a sufficiently fast rate, or their estimation errors satisfy*

$$\|\hat{m} - m\| \cdot \|\hat{h} - h\| = o_p(n^{-1/2}).$$

This guarantees that the second-order remainder in the influence function expansion vanishes.

Proposition C.1 (EIF form is invariant; only m, \bar{h} change). *Under Assumptions 1–3 and positivity, the efficient influence functions are*

$$D_{\text{stat}}^*(O) = \frac{\mathbb{1}\{V = v\}}{\bar{h}(v | W)} \{Y - m(v, W)\} + (m(v, W) - \psi_{\text{stat}}),$$

$$D_{\text{stoch}}^*(O) = \frac{\bar{h}^*(V | W)}{\bar{h}(V | W)} \{Y - m(V, W)\} + \left(\int m(v, W) h^*(dv | W) - \psi_{\text{stoch}} \right),$$

with expectations 0 and finite variances. The expressions are identical to the single-view case, with m, \bar{h} computed for the aggregated (W, V) .

Lemma C.2 (Dependence and effective sample size under aggregation). *Let $G^{(m)}$ have maximum degree $K_{\max,n}^{(m)}$. Define the aggregated support-graph \tilde{G} by connecting $i \sim j$ if they are connected in at least one view. Then*

$$K_{\max,n}(\tilde{G}) \leq \sum_{m=1}^M K_{\max,n}^{(m)}.$$

Under Assumption 4 (two-hop local dependence per view), the dependency graph for $\{D^(O_i)\}_{i=1}^n$ on \tilde{G} has maximum degree*

$$\Delta_n \lesssim K_{\max,n}(\tilde{G})^2 \leq \left(\sum_{m=1}^M K_{\max,n}^{(m)} \right)^2.$$

By Lemma C.1 and Proposition C.1, the multi-view reparametrization preserves both the target functional and the EIF, so the EIF coincides with the forms in Ogburn et al. (2020). Hence, under the stated regularity (positivity, bounded moments) and $K_{\max,n}^2/n \rightarrow 0$, the CAN result can be obtained.

Proposition C.2 (CLT with multi-view aggregation). *Let $Z_i := D^*(O_i)$ with $\mathbb{E}|Z_i|^3 < \infty$. Assume $K_{\max,n}(\tilde{G})^2/n \rightarrow 0$. Then, with $C_n := n/(1 + \Delta_n)$ and $\Delta_n \lesssim K_{\max,n}(\tilde{G})^2$,*

$$\sqrt{C_n}(\mathbb{P}_n - P_0)D^*(O) \Rightarrow \mathcal{N}(0, \sigma^2), \quad \sigma^2 = \text{Var}\{D^*(O)\}.$$

Theorem C.1 (CAN with multi-view aggregation). *Under Assumptions 1–5, Lemma C.1, Proposition C.1, the CLT of Proposition C.2, and the nuisance-rate condition $\|\hat{m} - m\| \|\hat{h} - \bar{h}\| = o_p(C_n^{-1/2})$, the TMLE $\hat{\psi}_n$ for either target satisfies*

$$\sqrt{C_n}(\hat{\psi}_n - \psi) \Rightarrow \mathcal{N}(0, \sigma^2), \quad C_n = \frac{n}{1 + \Delta_n}, \quad \Delta_n \lesssim K_{\max,n}(\tilde{G})^2.$$

D Additional Experimental Results

Implementation details. All GCN layers use hidden dimension 64. The attention fusion module uses a single linear layer. The soft histogram density estimator uses $C = 10$ bins. The B-spline perturbation network uses degree-3 splines with 10 knots. All models are trained with Adam optimizer (learning rate 10^{-3}) for 500 epochs. Results are reported as mean \pm standard deviation over 5 independent runs with different random seeds.

D.1 Ablation Study

Ablation: contribution of multi-view aggregation. To directly assess the benefit of aggregating multiple network views, we compare MVDR against three single-view variants (MVDR-n1, MVDR-n2, MVDR-n3) and one variant with the HSIC regularisation disabled (MVDR-noHSIC), all trained with identical hyperparameters. Results on Flickr are reported in Table 3. On the Flickr dataset, MVDR (full) outperforms all single-view variants in 4 evaluation metrics, demonstrating that aggregating complementary information from multiple network views consistently improves causal effect estimation. This suggests that while each individual view captures some useful structural information, the full multi-view representation provides complementary signals that no single view alone can recover. Removing the HSIC regularisation (MVDR-noHSIC) also leads to degradation in 4 metrics, confirming that encouraging diversity across view-specific representations is beneficial.

D.2 Additional Metrics

The main purpose of this paper is to evaluate the average potential outcome $E[\bar{Y}_n^*]$, which means the outcome for an observed network with latent dependency if certain hypothetical intervention is assigned to it. In Table 4, we also report the results for average main effect, average spillover effect, average total effect and individual main effect, individual spillover effect, individual total effect by pre-defining the exposure mapping v as the average ratio of treated neighbors across views.

Define ψ_i as $\psi_i(t, z) = m_i(z, w_i) + \epsilon(z, w_i) \cdot H_i$, and ψ as $\psi(t, z) = \frac{1}{n} \sum_{i=1}^n \psi_i(t, z)$. Then we have the following definitions:

- The average main effect (AME) measures the difference in the mean outcomes between units assigned to $T = t, Z = 0$ and assigned to $T = t', z = 0$: $\tau^{(t,0),(t',0)} = \psi(t, 0) - \psi(t', 0)$.

Table 3: Ablation study on Flickr. MVDR-view k uses only the k -th view; MVDR-noHSIC disables HSIC regularisation. All variants use identical hyperparameters. Best result in each column is **bolded**.

Method	ATE-indiv	ATE-total	PEHE-indiv	PEHE-total
MVDR-N1	0.2038 \pm 0.1099	0.2153 \pm 0.1297	3.0374 \pm 0.0579	3.0175 \pm 0.0856
MVDR-N2	0.2212 \pm 0.1508	0.2382 \pm 0.1533	3.0459 \pm 0.0591	3.0284 \pm 0.0716
MVDR-N3	0.2091 \pm 0.1132	0.2156 \pm 0.1046	3.0366 \pm 0.0619	3.0420 \pm 0.0875
MVDR-noHSIC	0.1632 \pm 0.1053	0.1783 \pm 0.2052	3.0373 \pm 0.0532	3.0273 \pm 0.0698
MVDR (full)	0.1623 \pm 0.1050	0.1782 \pm 0.1800	3.0354 \pm 0.0547	3.0261 \pm 0.0711

- The average spillover effect (ASE) measures the difference in mean outcomes between units assigned to $T = 0, Z = z$ and assigned to $T = 0, Z = z'$: $\tau^{(0,z),(0,z')} = \psi(0, z) - \psi(0, z')$.
- The average total effect (ATE) measures the difference in mean outcomes between units assigned to $T = t, Z = z$ and assigned $T = t', Z = z'$: $\tau^{(t,z),(t',z')} = \psi(t, z) - \psi(t', z')$.
- The individual main effect (IME) measures the difference in mean outcomes of a particular unit x_i assigned to $T = t, Z = 0$ and assigned $T = t', Z = 0$: $\tau_i^{(t,0),(t',0)} = \psi_i(t, 0) - \psi_i(t', 0)$.
- The individual spillover effect (ISE) measures the difference in mean outcomes of a particular unit x_i assigned to $T = 0, Z = z$ and assigned $T = 0, Z = z'$: $\tau_i^{(0,z),(0,z')} = \psi_i(0, z) - \psi_i(0, z')$.
- The individual total effect (ITE) measures the difference in mean outcomes of a particular unit x_i assigned to $T = t, Z = z$ and assigned $T = t', Z = z'$: $\tau_i^{(t,z),(t',z')} = \psi_i(t, z) - \psi_i(t', z')$.

For the average effects, we use mean absolute error (MAE) as metric: $\epsilon_{average} = |\hat{\tau} - \tau|$. For the individual effects, we use Precision in Estimation of Heterogeneous Effect (PEHE) as metric: $\epsilon_{individual} = \frac{1}{n} \sum_{i=1}^n (\tau_i - \hat{\tau}_i)^2$. Specifically, HINITE does not indentify spillover effects and main effects, so the results are not reported.

Table 4: Comparison of estimation errors across different models and datasets. Best performance in each row is **bolded**. ‘-’ represents the results are not reported by that method.

Split	Metric	dataset	effect	CFR	CFR+z	ND	ND+z	TARNet	TARNet+z	NetEst	HINITE	TNet	MVDR
With-in-Sample	$\epsilon_{\text{average}}$	BlogCata	AME	0.1425 _{0.1516}	0.2549 _{0.1329}	0.1590 _{0.0859}	0.3080 _{0.1620}	0.1550 _{0.0683}	0.3335 _{0.1443}	0.3533 _{0.1410}	-	40.3076 _{92.4736}	0.1720 _{0.0995}
			ASE	0.1735 _{0.1516}	0.2509 _{0.2208}	0.1735 _{0.1516}	0.2430 _{0.2278}	0.1734 _{0.0696}	0.2241 _{0.2090}	0.2191 _{0.2231}	-	145.1494 _{214.8604}	0.1905 _{0.1016}
			ATE	0.2449 _{0.1704}	0.1650 _{0.0786}	0.1326 _{0.1208}	0.1105 _{0.1060}	0.1326 _{0.1076}	0.1308 _{0.1385}	0.5076 _{0.1546}	0.6230 _{0.4354}	5707.6897 _{9251.3347}	0.1639 _{0.1457}
	Flickr	AME	0.2691 _{0.1097}	0.1692 _{0.1467}	0.2726 _{0.2432}	0.4251 _{0.1537}	0.2918 _{0.2166}	0.3492 _{0.0897}	0.7693 _{0.0524}	-	12.0787 _{11.3391}	0.0593 _{0.0470}	
		ASE	0.0887 _{0.0211}	0.1767 _{0.0620}	0.1078 _{0.0822}	0.1964 _{0.0883}	0.1078 _{0.0822}	0.2736 _{0.0700}	0.1948 _{0.0679}	-	158.5197 _{175.1475}	0.2296 _{0.1129}	
		ATE	0.2617 _{0.1099}	0.1887 _{0.1861}	0.3048 _{0.1872}	0.3013 _{0.1364}	0.2833 _{0.2067}	0.3701 _{0.1785}	1.0245 _{0.1374}	0.1751 _{0.1004}	2923.5093 _{3001.1783}	0.4908 _{0.2424}	
	DBLP	AME	0.2153 _{0.1557}	0.2606 _{0.1441}	0.5965 _{0.2617}	0.2916 _{0.1325}	1.0216 _{0.6321}	0.3672 _{0.2286}	0.8758 _{0.2556}	-	3.5683 _{11.7916}	0.0336 _{0.0348}	
		ASE	0.2606 _{0.1769}	0.329 _{0.2398}	0.3108 _{0.1768}	0.3311 _{0.1755}	0.1381 _{0.0438}	0.1376 _{0.0429}	0.2619 _{0.1827}	-	7.5224 _{11.0523}	0.2611 _{0.0314}	
		ATE	0.076 _{0.0833}	0.1739 _{0.0943}	0.8423 _{0.2742}	0.4265 _{0.0574}	1.1592 _{0.5784}	0.2585 _{0.1698}	1.0077 _{0.0836}	0.6815 _{0.3553}	11.9821 _{6.2407}	0.4408 _{0.0527}	
	BlogCata	IME	3.0597 _{0.1232}	3.0358 _{0.1232}	2.9224 _{0.094}	2.9359 _{0.0842}	2.9427 _{0.0894}	2.9570 _{0.0808}	2.9412 _{0.0865}	-	140.5076 _{127.1403}	2.8966 _{0.0583}	
		ISE	2.8973 _{0.0845}	2.9072 _{0.0845}	2.8973 _{0.0852}	2.9069 _{0.0842}	2.8973 _{0.0826}	2.9044 _{0.0836}	2.9088 _{0.0800}	-	270.6221 _{341.4499}	2.7668 _{0.0404}	
		ITE	3.0540 _{0.097}	3.0182 _{0.0970}	2.9004 _{0.0510}	2.8967 _{0.0496}	2.9216 _{0.0529}	2.9185 _{0.0591}	2.9353 _{0.0535}	3.3113 _{0.2458}	12038.2843 _{15704.8492}	2.8692 _{0.0482}	
	$\epsilon_{\text{individual}}$	Flickr	IME	9.3637 _{5.1163}	2.9457 _{0.0596}	6.0954 _{1.8223}	2.9607 _{0.0690}	2.9683 _{0.0648}	2.9124 _{0.0499}	2.9537 _{0.0498}	-	35.5421 _{32.8715}	0.6581 _{0.2529}
			ISE	2.8980 _{0.0439}	2.8965 _{0.0976}	3.0014 _{0.0237}	3.0342 _{0.0964}	2.9228 _{0.0547}	2.8865 _{0.0989}	2.8896 _{0.1010}	-	440.6491 _{463.7988}	0.6886 _{0.2660}
			ITE	9.3395 _{5.1269}	3.0072 _{0.641}	6.1200 _{1.8291}	2.9525 _{0.0830}	2.9506 _{0.0622}	2.9470 _{0.0449}	3.0514 _{0.0463}	3.0427 _{0.0662}	9314.1054 _{9633.2401}	0.8652 _{0.2978}
	DBLP	IME	3.3294 _{0.2079}	3.3316 _{0.2094}	3.0448 _{0.0829}	2.6218 _{0.0341}	2.8543 _{0.2530}	2.6366 _{0.0541}	3.4669 _{0.2156}	-	6.3263 _{11.4395}	0.0506 _{0.0353}	
		ISE	3.4736 _{0.1455}	3.483 _{0.1517}	2.9793 _{0.0928}	2.6923 _{0.0325}	2.6921 _{0.0321}	2.6921 _{0.0322}	3.508 _{0.1563}	-	12.3642 _{21.6718}	0.2852 _{0.0557}	
		ITE	3.3944 _{0.0849}	3.3985 _{0.0839}	3.0996 _{0.0805}	2.5758 _{0.0772}	3.2055 _{0.1500}	2.5984 _{0.1022}	3.5871 _{0.0966}	2.9779 _{0.2013}	13.5564 _{15.9518}	0.4881 _{0.0460}	
Out-of-Sample	ϵ_{ATE}	BlogCata	AME	0.2103 _{0.154}	0.203 _{0.1861}	0.5117 _{0.2551}	0.6367 _{0.3633}	0.4508 _{0.3217}	0.5380 _{0.4065}	0.2312 _{0.2088}	-	40.3076 _{92.4736}	0.1289 _{0.0416}
			ASE	0.1277 _{0.0965}	0.1882 _{0.1349}	0.1277 _{0.0965}	0.1825 _{0.1322}	0.1277 _{0.0965}	0.1736 _{0.1373}	0.1209 _{0.098}	-	145.1494 _{214.8604}	0.1678 _{0.0751}
			ATE	0.2449 _{0.1704}	0.2726 _{0.2062}	0.5484 _{0.3592}	0.7826 _{0.4412}	0.4874 _{0.2824}	0.6466 _{0.3515}	0.3724 _{0.2632}	0.3574 _{0.2483}	5707.6897 _{9251.3347}	0.1639 _{0.1055}
	Flickr	AME	0.2691 _{0.1097}	0.2908 _{0.2120}	0.1662 _{0.1123}	0.4614 _{0.2634}	0.1626 _{0.1219}	0.4107 _{0.2165}	0.7582 _{0.1641}	-	9.0943 _{10.9997}	0.2605 _{0.0590}	
		ASE	0.0887 _{0.0211}	0.2074 _{0.0858}	0.0887 _{0.0211}	0.2475 _{0.1138}	0.0887 _{0.0211}	0.2106 _{0.1111}	0.1406 _{0.1115}	-	92.6766 _{87.5788}	0.2426 _{0.0766}	
		ATE	0.2617 _{0.1099}	0.1940 _{0.1716}	0.2669 _{0.1859}	0.2878 _{0.1566}	0.3395 _{0.1229}	0.2618 _{0.1089}	0.7807 _{0.1527}	0.1674 _{0.1618}	2152.0205 _{1924.3221}	0.5516 _{0.3424}	
	DBLP	AME	0.2153 _{0.1557}	0.2606 _{0.1441}	0.7047 _{0.2382}	0.284 _{0.1825}	0.8104 _{0.4552}	0.4669 _{0.1738}	0.8758 _{0.2556}	-	3.5683 _{11.7916}	0.0506 _{0.0353}	
		ASE	0.2006 _{0.1769}	0.329 _{0.2398}	0.2606 _{0.1769}	0.2718 _{0.1793}	0.2606 _{0.1769}	0.2627 _{0.176}	0.2619 _{0.1827}	-	7.5224 _{11.0523}	0.2852 _{0.0557}	
		ATE	0.076 _{0.0833}	0.1739 _{0.0943}	0.8558 _{0.0891}	0.3097 _{0.1377}	0.9616 _{0.325}	0.5847 _{0.1165}	1.0077 _{0.0836}	0.7259 _{0.4971}	11.9821 _{6.2407}	0.4881 _{0.0460}	
	BlogCata	IME	3.6311 _{1.2841}	3.9781 _{1.7998}	5.6662 _{2.2252}	7.2773 _{3.4133}	4.7532 _{2.3242}	5.5978 _{4.2786}	3.1134 _{0.1372}	-	140.5076 _{127.1403}	2.9013 _{0.0460}	
		ISE	2.9072 _{0.0845}	2.8033 _{0.0519}	2.7983 _{0.0453}	2.8027 _{0.0516}	2.7983 _{0.0453}	2.8022 _{0.0522}	3.0264 _{0.124}	-	270.6221 _{341.4499}	2.7966 _{0.0467}	
		ITE	3.6261 _{1.255}	3.9765 _{1.7725}	5.6698 _{2.2135}	7.2954 _{3.3815}	4.7462 _{2.3115}	5.5797 _{4.2738}	3.0796 _{0.1691}	3.2063 _{0.3024}	12038.2843 _{15704.8492}	2.8580 _{0.0460}	
	ϵ_{PEHE}	Flickr	IME	9.3637 _{5.1163}	3.573 _{0.6824}	4.0064 _{1.4430}	3.1551 _{0.1926}	3.2343 _{0.1757}	3.0985 _{0.0172}	3.0916 _{0.0363}	-	68.2053 _{39.6463}	0.7660 _{0.2204}
			ISE	3.0457 _{0.0475}	3.0295 _{0.0976}	3.0457 _{0.0475}	3.0566 _{0.0409}	3.0457 _{0.0475}	3.0314 _{0.0964}	3.0338 _{0.0290}	-	352.7498 _{303.2289}	0.7214 _{0.2673}
			ITE	9.3395 _{5.1269}	3.5895 _{0.641}	4.0251 _{1.4578}	3.1328 _{0.1616}	3.242 _{0.1979}	3.0533 _{0.0709}	3.0834 _{0.0812}	3.1153 _{0.0732}	8165.584 _{7729.401}	0.9543 _{0.3456}
	DBLP	IME	3.3294 _{0.2079}	3.3316 _{0.2094}	3.4027 _{0.1876}	3.3357 _{0.2167}	3.450.1486	3.3558 _{0.2131}	3.4669 _{0.2156}	-	6.3263 _{11.4395}	0.5119 _{0.0155}	
		ISE	3.4736 _{0.1455}	3.483 _{0.1517}	3.4736 _{0.1455}	3.4745 _{0.146}	3.4736 _{0.1455}	3.4737 _{0.1455}	3.508 _{0.1563}	-	12.3642 _{21.6718}	0.5175 _{0.0370}	
		ITE	3.3944 _{0.0849}	3.3985 _{0.0839}	3.5013 _{0.0936}	3.4116 _{0.0815}	3.5403 _{0.1502}	3.4462 _{0.0883}	3.5871 _{0.0966}	0.8227 _{0.219}	13.5564 _{15.9518}	0.7080 _{0.0249}	