LEARNING EXPOSURE MAPPING FUNCTIONS FOR INFERRING HETEROGENEOUS PEER EFFECTS

Anonymous authorsPaper under double-blind review

000

001

002003004

006

008 009

010 011

012

013

014

016

017

018

019

021

024

025

026

027

028

029

031

033 034

037

040

041

042 043

044

046

047

048

050 051

052

ABSTRACT

Peer effect refers to the difference in counterfactual outcomes for a unit resulting from different levels of peer exposure, the extent to which the unit is exposed to the treatments, actions, or behaviors of its peers. Peer exposure is typically captured through an explicitly defined exposure mapping function that aggregates peer treatments and outputs peer exposure. Exposure mapping functions range from simple functions like the number or fraction of treated friends to more sophisticated functions that allow for different peers to exert different degrees of influence. However, the true function is rarely known in practice and when the function is misspecified, this leads to biased causal effect estimation. To address this problem, the focus of our work is to move away from the need to explicitly define an exposure mapping function and instead introduce a framework that allows learning this function automatically. We develop EGONETGNN, a graph neural network (GNN), for heterogeneous peer effect estimation that automatically learns the appropriate exposure mapping function and allows for complex peer exposure mechanisms that involve not only peer treatments but also attributes of the local neighborhood, including node, edge, and structural attributes. We theoretically and empirically show that GNN models that use peer exposure based on the number or fraction of treated peers or learn peer exposure naively face difficulty accounting for such influence mechanisms. Our evaluation on synthetic and semi-synthetic network data shows that our method is more robust to different unknown underlying influence mechanisms when compared to state-of-the-art baselines.

1 Introduction

In networked environments, the outcome of a unit can be influenced by the treatments or outcomes of other units, a phenomenon known as interference. For example, in a contact network, the smoking habits of peers may affect an individual's respiratory health, and in a social network the political affiliations of peers may influence one's stance on a policy issue like immigration. Peer effects capture this influence by comparing an individual's outcomes under different peer network conditions (e.g., having no smoker peers versus some smoker peers, or observed peer political affiliations versus counterfactual, flipped affiliations). Peer effect estimation is important for policy-making and targeted intervention design in many domains, including healthcare (Barkley et al., 2020), online advertisement (Nabi et al., 2022), and education (Patacchini et al., 2017).

Peer network conditions are typically captured through an *explicitly defined* exposure mapping function (Aronow and Samii, 2017) that summarizes the peer treatments and peer network and outputs peer exposure, which is the equivalent to a composite peer treatment value. The peer effect is defined as the difference in outcomes under two distinct levels of peer exposure. Different peer exposure mapping functions capture different possible underlying influence mechanisms. Typically, domain experts define exposure mapping functions appropriate to the causal question and domain of interest. The advantage of exposure mapping functions is that they reduce the high dimensionality of peer network attributes and that they are invariant to irrelevant contexts (e.g., permutation of peers).

Figure 1 presents examples of prominent exposure mapping functions and the resulting peer exposure values for a toy peer network. The first graph shows Gaby's peer network along with the observed (i.e., factual) treatments for Gaby's peers. The second graph shows hypothetical (i.e., counterfactual) treatments for the peers. The peers in the treatment group (e.g., smokers) and control group (e.g.,

Example exposure mapping functions	Peer exposure value			
Example exposure mapping functions	Factual	Counterfactual		
Binary (at least a peer treated)	1	1		
Fraction of treated peers	3/6	2/6		
Linear threshold (40%)	1	0		
Weighted fraction (tie strengths)	4/8	3/8		
Weighted fraction (attribute similarity: female)	1/3	0/3		
Local structure: Causal network motifs E.g., Closed triads: Count Count Count	< 3, 1, 2 >	< 0, 4, 2 >		

Figure 1: Illustration of different possible peer exposure representations for a node (Gaby) in a toy peer network. Red nodes represent peers in the treatment group, and blue nodes represent peers in the control group. Gray star node represents the node that has a fixed treatment.

non-smokers) are depicted as red and blue nodes, respectively. The edge weights capture the tie strengths in the network. Binary peer exposure mapping is the simplest and it summarizes peer treatments to 0 or 1, e.g., whether any peers have been treated (Bargagli-Stoffi et al., 2025) or whether the weighted treatment of peers has reached a linear threshold (Tran and Zheleva, 2022). Some exposure mapping functions assume that all peers influence equally (e.g., fraction of treated peers (Hudgens and Halloran, 2008; Jiang and Sun, 2022)), while others consider that different peers can exert different degrees of influence (e.g., weighted fraction (Forastiere et al., 2021) or sum (Zhao et al., 2024) of treated peers). Peer exposure has also been modeled with counts of different causal network motifs, i.e., recurrent subgraphs in a unit's peer network with treatment assignments as attributes (Yuan et al., 2021). We discuss the related work in more detail in the Appendix A.2.

A key challenge in peer effect estimation is that the true exposure mapping function is rarely known in practice and when the function is misspecified, this leads to biased causal effect estimation. The focus of this paper is to move away from the need to explicitly define an exposure mapping function and instead learn this function automatically from data. This has the advantage of reducing subjectivity and allowing for automated representation of peer exposure under unknown and complex peer influence mechanisms. More specifically, we study the problem of exposure mapping function learning in the context of heterogeneous peer effect estimation. Heterogeneous peer effects (HPE) denote variation in peer effects across individuals that may originate from personal attributes or from characteristics of their peer networks. For example, while having a friend who smokes may have a negative effect on health for some people, it may make no difference for others.

We propose EGONETGNN, a novel graph neural network (GNN) architecture, that automatically learns a relevant exposure mapping function under appropriate identifiability assumptions. EGONET-GNN allows for complex peer influence mechanisms that, in addition to peer treatments, can involve the local neighborhood structure, node, and edge attributes. Our work builds upon the success of utilizing neural networks (NNs) (Shalit et al., 2017; Im et al., 2021; Shi et al., 2019) and, recently, graph neural networks (GNNs) (Jiang and Sun, 2022; Cai et al., 2023; Chen et al., 2024; Khatami et al., 2024) for end-to-end learning of counterfactual outcome models or causal effect estimators. Few studies have utilized GNNs to learn the exposure mapping function (Mao et al., 2025; Wu et al., 2025) or to derive peer exposure embedding by aggregating feature embeddings and peer treatments (Adhikari and Zheleva, 2025; Zhao et al., 2024). However, these works use off-the-shelf GNNs like GCN (Kipf and Welling, 2016) or GIN (Xu et al., 2018) and prior work (Chen et al., 2020) has shown such architectures lack expressiveness for counting subgraphs with cycles and for capturing mechanisms involving local neighborhood structure. On the other hand, counts of such subgraphs, like causal network motifs, are rich features for capturing local structural contexts (Yuan et al., 2021), but they are expensive to compute, inflexible, and may not capture every local structural context (e.g., edge weights).

One of the biggest strengths of EGONETGNN is the ability to capture the exposure mapping functions studied in previous works, including finding relevant causal network motifs and scaling to higher-order motifs. To add robustness to the downstream peer effect estimation task, EGONETGNN is designed to learn the exposure mapping function to produce representation that is *expressive* to differentiate between different peer exposure conditions and *invariant* to irrelevant contexts. Moreover, EGONETGNN is designed to promote bounded representation with substantial coverage of possible peer exposure values. Figure 2 shows an overview of EGONETGNN. While most peer effect estimation frameworks contain a feature mapping and a counterfactual outcome model component,

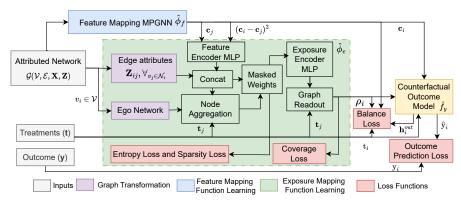


Figure 2: An overview of the proposed EGONETGNN model to learn exposure mapping function for peer effect estimation. EGONETGNN extracts ego networks, for each node v_i , with peer treatments along with feature embedding and its edge attributes as node attributes. Then, node-level aggregations are performed to capture local neighborhood contexts. These contexts are passed through a masked weight layer and encoded by an multi-layer perceptron (MLP) to learn relevant influence mechanisms and summarized with graph-level aggregation. The learned peer exposure embeddings (ρ_i) , along with the feature embeddings (c_i) , and treatment (t_i) are passed to a counterfactual outcome model that is used to infer peer effects. The graph transformation ensures expressiveness, while balance, coverage, entropy, and sparsity losses promote the robustness of the peer exposure representation.

the novel additional component in ours is the custom-designed exposure mapping function learning (marked in green in the figure). We design this component to excel in counting attributed subgraphs, such as causal network motifs, enhancing its expressiveness to capture unknown underlying peer exposure mechanisms. We theoretically and empirically show that, unlike EGONETGNN, existing GNN-based approaches that solely rely on homogeneous peer exposure or learn heterogeneous peer exposure naively lack expressiveness in capturing heterogeneous peer influence mechanisms based on local neighborhood structure.

2 Causal Inference Problem Setup

Notations. We represent the network as an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with a set of $n = |\mathcal{V}|$ nodes, a set of edges \mathcal{E} , node attributes \mathbf{X} , and edge attributes \mathbf{Z} . Let $\mathbf{t} = \langle t_1, ..., t_i, ..., t_n \rangle$ be a random variable comprising the treatment variables t_i for each node $v_i \in \mathcal{V}$ in the network and y_i be a random variable for v_i 's outcome. Let $\pi = \langle \pi_1, ..., \pi_i, ..., \pi_n \rangle$ be an assignment to \mathbf{t} with $\pi_i \in \{0, 1\}$ assigned to t_i . Let $\mathbf{t}_{-i} = \mathbf{t} \setminus t_i$ and $\pi_{-i} = \pi \setminus \pi_i$ denote random variable and its value for treatment assignment to other units except v_i , which we refer to as peer treatments for convenience.

Peer exposure reflects how much a unit is exposed to peer treatments and is defined as follows.

Definition 1 (Peer exposure and exposure mapping function). Peer exposure for unit v_i is defined as $\rho_i \in [0,1]^d = \phi_e(\pi_{-i}, \mathcal{G}, \mathbf{X}, \mathbf{Z})$, where ϕ_e is the *exposure mapping function* that maps high-dimensional contexts $\{\pi_{-i}, \mathcal{G}, \mathbf{X}, \mathbf{Z}\}$ to a d-dimensional peer exposure representation bounded between 0 and 1 such that $\mathbf{y}_i(\mathbf{t}_i = \pi_i, \mathbf{t}_{-i} = \pi_{-i})|\{\mathcal{G}, \mathbf{X}, \mathbf{Z}\} = \mathbf{y}_i(\mathbf{t}_i = \pi_i, \mathbf{p}_i = \rho_i)|\{\mathcal{G}, \mathbf{X}, \mathbf{Z}\}$.

Definition 1 maps peer treatments $\mathbf{t}_{-i} = \pi_{-i}$ and peer network contexts $\{\mathcal{G}, \mathbf{X}, \mathbf{Z}\}$ to peer exposure $\mathbf{p}_i = \rho_i$ in terms of equivalence of counterfactual outcomes $\mathbf{y}_i(\mathbf{t}_i = \pi_i, \mathbf{t}_{-i} = \pi_{-i})$ and $\mathbf{y}_i(\mathbf{t}_i = \pi_i, \mathbf{p}_i = \rho_i)$. Here, $\mathbf{y}_i(\mathbf{t}_i = \pi_i, \mathbf{p}_i = \rho_i)$, captures that, in interference settings, the counterfactual outcome of a unit v_i is influenced by both unit's treatment $\mathbf{t}_i = \pi_i$ and peer exposure $\mathbf{p}_i = \rho_i$. Note that the exposure mapping function could map different contexts to the same peer exposure.

Peer effect refers to the difference in counterfactual outcomes for different values of peer exposure. Heterogeneous peer effects (HPE) refers to different units having different peer effects dependent on their contexts. For any given unit v_i , its heterogeneous peer effect is described through its context, i.e., for peer exposures $\mathbf{p}_i = \boldsymbol{\rho}_i$ versus $\mathbf{p}_i = \boldsymbol{\rho}_i'$ and unit's treatment $\mathbf{t}_i = \pi_i$ conditioned on the unit's contexts \mathbf{c}_i , it is defined as:

$$\delta_i(\boldsymbol{\rho}_i, \boldsymbol{\rho}_i') = \mathbb{E}[\mathbf{y}_i(\mathbf{t}_i = \boldsymbol{\pi}_i, \mathbf{p}_i = \boldsymbol{\rho}_i)|\mathbf{c}_i] - \mathbb{E}[\mathbf{y}_i(\mathbf{t}_i = \boldsymbol{\pi}_i, \mathbf{p}_i = \boldsymbol{\rho}_i')|\mathbf{c}_i], \tag{1}$$

where expectation is over units with similar contexts \mathbf{c}_i , referred to as *effect modifiers* (e.g., unit's degree or node attributes), defined by a feature mapping function of contexts $\{\mathcal{G}, \mathbf{X}, \mathbf{Z}\}$ from v_i 's

perspective, i.e., $\mathbf{c}_i = \phi_f(v_i, \mathcal{G}, \mathbf{X}, \mathbf{Z})$. Substituting peer exposures ρ_i and ρ_i' with corresponding exposure mapping functions for two peer treatment assignments π_{-i} versus π_{-i}' in Eq. 1, we get:

$$\delta_i(\boldsymbol{\pi}_{-i}, \boldsymbol{\pi}'_{-i}) = \mathbb{E}[\mathbf{y}_i(\mathbf{t}_i = \boldsymbol{\pi}_i, \mathbf{p}_i = \phi_e(\boldsymbol{\pi}_{-i}, \mathcal{G}, \mathbf{X}, \mathbf{Z})) | \mathbf{c}_i] - \mathbb{E}[\mathbf{y}_i(\mathbf{t}_i = \boldsymbol{\pi}_i, \mathbf{p}_i = \phi_e(\boldsymbol{\pi}'_{-i}, \mathcal{G}, \mathbf{X}, \mathbf{Z})) | \mathbf{c}_i]. \tag{2}$$

Causal identification. Now, we discuss the identification of peer effects that involves expressing counterfactual outcomes in terms of observational and/or interventional distributions.

Next, we make two commonly adopted assumptions in network interference settings.

Assumption 1 (Pre-treatment network). The network \mathcal{G} along with node attributes \mathbf{X} and edge attributes \mathbf{Z} are measured before treatment assignments $\mathbf{t} = \pi$ and treatments are not mutable.

Assumption 2 (Neighborhood Interference). The counterfactual outcome of a unit depends only on its immediate neighborhood treatments, i.e., $\mathbf{y}_i(\mathbf{t}_i = \pi_i, \mathbf{t}_{-i} = \pi_{-i})|\mathbf{c}_i = \mathbf{y}_i(\mathbf{t}_i = \pi_i, \mathbf{t}_{-i}^{\mathcal{N}_i} = \pi_{-i}^{\mathcal{N}_i})|\mathbf{c}_i = \mathbf{y}_i(\mathbf{t}_i = \pi_i, \mathbf{p}_i = \phi_e(\boldsymbol{\pi}_{-i}^{\mathcal{N}_i}, \mathcal{G}, \mathbf{X}, \mathbf{Z}))|\mathbf{c}_i$, where $\mathcal{N}_i = \{j: (v_i, v_j) \in \mathcal{E}\}$, $\mathbf{t}_{-i}^{\mathcal{N}_i} = \mathbf{t}_{-i} \cap \{\mathbf{t}_j: j \in \mathcal{N}_i\}$, and $\boldsymbol{\pi}_{-i}^{\mathcal{N}_i} = \boldsymbol{\pi}_{-i} \cap \{\pi_j: j \in \mathcal{N}_i\}$ denote neighborhood set, treatments, and assignments, respectively.

Assumption 1 is a general assumption in experimental and observational studies, and Assumption 2 is a common simplifying assumption that presumes network influence is mediated by immediate neighbors but our work could be extended to consider interference from multiple-hop neighborhoods. For ease of exposition, we drop the superscript \mathcal{N}_i in neighborhood treatments and assignments.

For causal identification, we assume unconfoundedness, similar to previous work (Ma et al., 2022; Wu et al., 2025):

Assumption 3 (Unconfoundedness). For all unit treatment $\pi_i \in \{0,1\}$ and peer treatment assignments $\pi_{-i} \in \{0,1\}^{n-1}$, there exists a feature mapping function $\phi_f \in \Phi_f$ and an exposure mapping function $\phi_e \in \Phi_e$ such that the counterfactual outcome is independent of unit treatment and peer exposure conditions given the context $\mathbf{c}_i = \phi_f(v_i, \mathcal{G}, \mathbf{X}, \mathbf{Z})$, i.e., $\mathbf{y}_i(\mathbf{t}_i = \pi_i, \mathbf{p}_i = \phi_e(\pi_{-i}, \mathcal{G}, \mathbf{X}, \mathbf{Z})) \perp \{\mathbf{t}_i, \mathbf{p}_i\} | \mathbf{c}_i$.

Assumption 3 implies that the observed network context is sufficient for controlling for confounding, and there are functions able to represent it compactly. Under this assumption, it is still possible to learn a feature mapping and exposure mapping functions that do not approximate the true functions which leads to a misspecification error. Therefore, it is important to learn an expressive function (e.g., a GNN) that is able to capture a wide range of possible functions. We also assume the standard *consistency* (Assumption 4) and *positivity* (Assumption 5), described in more detail in Appendix A.3. Next, we present the causal identification conditions and formally define the problem of exposure mapping function learning in the context of peer effect estimation.

Proposition 1. With Assumptions 1-5, the HPE δ_i in Eq. A.3 can be estimated from experimental or observational data as

$$\delta_i(\boldsymbol{\pi}_{-i}, \boldsymbol{\pi}'_{-i}) = \mathbb{E}[\mathbf{y}_i | \mathbf{t}_i = \pi_i, \mathbf{p}_i = \phi_e(\boldsymbol{\pi}_{-i}, \mathcal{G}, \mathbf{X}, \mathbf{Z}), \mathbf{c}_i] - \mathbb{E}[\mathbf{y}_i | \mathbf{t}_i = \pi_i, \mathbf{p}_i = \phi_e(\boldsymbol{\pi}'_{-i}, \mathcal{G}, \mathbf{X}, \mathbf{Z}), \mathbf{c}_i]. \tag{3}$$

The proof presented in Appendix A.3 stems from consistency and unconfoundedness assumptions.

Problem 1 (Exposure mapping function learning). Given network contexts $\{\mathcal{G}, \mathbf{X}, \mathbf{Z}\}$, treatments \mathbf{t} , and outcomes \mathbf{y} of n units, estimate the feature and exposure mapping functions $\hat{\phi_f}$ and $\hat{\phi_e}$ along with counterfactual outcome model \hat{f}_y such that mean squared error between true heterogeneous peer effect (HPE) δ_i and estimated HPE $\hat{\delta}_i$, i.e., $\frac{1}{n}\sum_{i=1}^n(\delta_i-\hat{\delta}_i)^2$, is minimized, where $\hat{\delta}_i=\hat{f}_y(\pi_i,\hat{\rho}_i,\hat{\mathbf{c}}_i)-\hat{f}_y(\pi_i,\hat{\rho}_i',\hat{\mathbf{c}}_i)$ with $\hat{\rho}_i=\hat{\phi_e}(\pi_{-i},\mathcal{G},\mathbf{X},\mathbf{Z}), \hat{\rho}_i'=\hat{\phi_e}(\pi_{-i}',\mathcal{G},\mathbf{X},\mathbf{Z})$, and $\hat{\mathbf{c}}_i=\hat{\phi}_f(v_i,\mathcal{G},\mathbf{X},\mathbf{Z})$.

The true HPE is unknown, but due Proposition 1, the factual outcomes can be utilized to jointly estimate $\hat{\phi_f}$, $\hat{\phi_e}$, and \hat{f}_{Y_i} as discussed in the next section.

3 EGONETGNN: LEARNING EXPOSURE MAPPING FUNCTION WITH GNNS

Figure 2 shows an overview of the proposed EGONETGNN model to simultaneously learn exposure mapping function $\hat{\phi}_e$, feature mapping function $\hat{\phi}_f$, and counterfactual outcome model \hat{f}_y for peer effect estimation. We aim to learn exposure mapping function $\hat{\phi}_e$ with three key properties: 1) expressiveness, 2) invariance, and 3) bounded and balanced representation. The expressiveness property ensures the peer exposure representation ρ_i returned by the function $\hat{\phi}_e$ is unique for

238

239

240 241

242

243

244

245

246

247

249

250

251

252

253

254

255

256

257

258

259 260

261

262 263

264

265

266

267

268

216 different relevant contexts, while the invariance property assures the representation ρ_i does not 217 vary due to irrelevant contexts. For example, in Figure 1, if the underlying peer influence depends 218 on clustering coefficients among treated, the function ϕ_e is expressive if it can capture the first 219 closed triad substructure. The standard message passing GNNs (e.g., GCN, GIN, etc) cannot capture 220 essential causal network motifs like closed triads (i.e., triangular motifs) (Chen et al., 2020). The 221 graph transformation and automated exposure mapping function learning in our EGONETGNN model 222 are designed to ensure that the peer exposure representation is at least as expressive as or superior to the approach of feature extraction by counting causal network motifs. In the above example, the 224 function ϕ_e is invariant to irrelevant contexts if the difference in other features like node attributes and edge weights do not change the learned representation ρ_i . To satisfy the property of bounded 225 representation, the learned representation ρ_i should be bounded, e.g., between 0 and 1, to reflect 226 no exposure and maximum exposure. Moreover, the representation ρ_i should have a substantial 227 coverage, which means it should be distributed across the possible range of exposure. Next, we 228 describe our feature mapping, exposure mapping, and counterfactual outcome model in detail. 229 230 3.1 ARCHITECTURE OF EGONETGNN 231 232 233 235 weight layer and an MLP followed by graph-level aggregation to learn peer exposure representation. 236

EGONETGNN first maps the attributed network to feature embedding using a MPGNN and extracts ego networks for each node v_i , incorporating peer treatments, node features, and edge attributes. It performs node-level aggregation to capture local context, which is processed through a masked

Feature mapping MPGNN. The feature mapping module aims to capture contexts that are potentially confounders or effect modifiers. Let Θ denote a multi-layer perceptron (MLP) and || denote a concatenation operator. The feature embedding c_i is obtained for l-th layer as:

$$c_i = \Theta_0(\boldsymbol{X}_i) || \boldsymbol{h}_i^l \text{ and } \boldsymbol{h}_i^l = \boldsymbol{h}_i^{l-1} + \sum_{j \in \mathcal{N}_i} \Theta_l \boldsymbol{h}_j^{l-1},$$
 (4)

where $h_i^0 = X_j || Z_{ij}$, and $h_i^0 = 0$ are initial conditions and N_i denote neighbors of node v_i . This MPGNN architecture incorporates edge attributes Z_{ij} while disentangling the hidden representation of the unit's own attributes $\Theta_0(X_i)$ from that of aggregated peer and edge attributes h_i^l .

Ego network construction. To learn an exposure mapping function that is as least as expressive as or superior to the approach of feature extraction by counting network motifs, we transform the node regression task to graph regression by extracting ego networks for each unit. In an ego network, the triangle structures involving an ego node are transformed as edges, which mitigates the limitation of GNNs to capture closed triad motifs. The ego network $\bar{\mathcal{G}}_i(\bar{\mathcal{V}}_i, \bar{\mathcal{E}}_i)$ is extracted from $\mathcal{G}(\mathcal{V}, \mathcal{E})$ for each node v_i such that node set \bar{V}_i consists neighbors of v_i , i.e., $\bar{V}_i = \{v_j : e_{ij} \in \mathcal{E} \land v_j \in \mathcal{V}\}$ and edge set $\bar{\mathcal{E}}_i$ consists edges between neighbors of v_i , i.e., $\bar{\mathcal{E}}_i = \{e_{jk} : e_{jk} \in \mathcal{E} \land v_j \in \bar{\mathcal{V}}_i \land v_k \in \bar{\mathcal{V}}_i\}.$

Feature encoder and node aggregation. Feature encoder module takes relevant peer feature embeddings and the distance between ego and peer feature embeddings, i.e., $c_{ij} = \Theta_{feat}(c_j || (c_i - c_i))$ $(c_j)^2$), to capture peer influence mechanisms involving peer attributes and feature similarity between ego and peers. Then, we transform an ego v_i 's edge attributes Z_{ij} to node attributes, i.e., $\bar{X}_j = Z_{ij}$, in the ego network $\bar{\mathcal{G}}_i(\bar{\mathcal{V}}_i,\bar{\mathcal{E}}_i)$ because the ego v_i itself is not present in the ego network. The node aggregation for each node v_j in the ego network $\bar{\mathcal{G}}_i$ considers neighbors' node attributes \bar{X}_k , feature encoding c_{ik} , edge attributes Z_{jk} , and peer treatments t_k , and is defined for l^{th} layer as follows:

$$\boldsymbol{h}_{j}^{l} = \boldsymbol{h}_{j}^{l-1} + \sum_{k \in \mathcal{N}_{j}} \boldsymbol{h}_{k}^{l-1}, \text{ with } \boldsymbol{h}_{k}^{0} = t_{k} || \bar{\boldsymbol{X}}_{k} || \boldsymbol{c}_{ik} || \boldsymbol{Z}_{jk} \text{ and } \boldsymbol{h}_{j}^{0} = 0.$$
 (5)

Masked weights and exposure encoder. Masked weights promotes representation that is invariant to irrelevant contexts and feeds the concatenation of node attributes and hidden state after L layers of node aggregation, i.e., $h_j^{agg} = \bar{\mathbf{X}}_j || c_{ij} || h_j^L$, through a masked fully connected layer as follows:

$$\boldsymbol{h}_{j}^{mask} = ReLU((\sigma(\mathbf{W}_{mask}) \odot \mathbf{W}_{agg}) \boldsymbol{h}_{j}^{agg} + \mathbf{b}_{agg}), \tag{6}$$

where ReLU and σ are a rectified linear unit and sigmoid activation functions, \odot indicates elementwise product, W_{mask} and W_{agg} are the weight matrices, and b_{agg} is the bias vector. The masked hidden representation h_i^{mask} is passed into an exposure encoder MLP to extract a low dimensional embedding. The goal of this module is to capture complex mechanisms based on the local neighborhood and reduce dimensionality. Formally, the output embedding h_j^{exp} is obtained as follows:

$$\boldsymbol{h}_{j}^{exp} = ReLU(\Theta_{exp}(ln(ReLU(\Theta_{enc}(\boldsymbol{h}_{j}^{mask})) + 1))), \tag{7}$$

 Θ_{enc} and Θ_{exp} are two MLPs and ln denotes log transformation that offers the benefit of rescaling features with large values that are significant in scale-free networks (e.g., online social networks) and introduces inductive bias to capture mechanisms involving ratios.

Graph readout. Finally, the peer exposure embedding ρ_i for node v_i is obtained by aggregating the representation h_j^{exp} for all $v_j \in \bar{\mathcal{V}}_i$ on the entire ego network $\bar{\mathcal{G}}_i(\bar{\mathcal{V}}_i,\bar{\mathcal{E}}_i)$ as $\rho_i = \sum_j (t_j \times h_j^{exp})/\sum_j h_j^{exp}||1 - e^{-\sum_j (t_j \times h_j^{exp})}$. We consider two aggregations such that the peer exposure embedding is bounded between 0 and 1, with 0 being the case of no peer exposure. The first aggregation captures proportion similar to the fraction of treated peers, but we weight each peer by $h_j^{exp}/\sum_j h_j^{exp}$ learned by the preceding layer. The second aggregation captures scale and is analogous to the number of treated peers, except that each peer is weighted by h_j^{exp} .

3.2 END-TO-END LEARNING OF EGONETGNN

The resulting peer exposure embeddings (ρ_i) and the feature embeddings (c_i) from the above module along with unit treatment (π_i) are passed to a counterfactual outcome model $f_y(\mathbf{t}_i = \pi_i, \mathbf{p}_i = \boldsymbol{\rho}_i, \mathbf{c}_i = \boldsymbol{c}_i)$ to obtain conditional counterfactual outcome $\mathbb{E}[\mathbf{y}_i(\mathbf{t}_i = \pi_i, \mathbf{p}_i = \boldsymbol{\rho}_i) | \mathbf{c}_i = \boldsymbol{c}_i] = \mathbb{E}[\mathbf{y}_i|\mathbf{t}_i = \pi_i, \mathbf{p}_i = \boldsymbol{\rho}_i, \mathbf{c}_i = \boldsymbol{c}_i]$ (Eq. 3). We adapt the Treatment Agnostic Representation Network (TARNet) and Counterfactual Regression (CFR) models (Shalit et al., 2017) as the counterfactual outcome model \hat{f}_y . The TARNet architecture consists of a single embedding MLP and two prediction heads to estimate counterfactual outcomes with unit treatment $\mathbf{t}_i = 1$ and $\mathbf{t}_i = 0$, i.e.,

$$\mathbf{h}_{i}^{emb} = \Theta_{emb}(\mathbf{c}_{i})||\mathbf{\rho}_{i}, \quad \hat{\mathbf{y}}_{i}(0) = \Theta_{y_{0}}(\mathbf{h}_{i}^{emb}), \quad \hat{\mathbf{y}}_{i}(1) = \Theta_{y_{1}}(\mathbf{h}_{i}^{emb}).$$
 (8)

The CFR architecture is similar except for an autoencoder to produce the embeddings, i.e., $h_i^{emb} = \Theta_{emb}(c_i||\rho_i)$ and $h_i^{out} = \Theta_{dec}(h_i^{emb})$. The CFR or TARNet model $\hat{f}_y(\pi_i, \rho_i, c_i)$ predicts outcome $\hat{y}_i = \hat{y}_i(1)$ if $\pi_i = 1$ and $\hat{y}_i = \hat{y}_i(0)$ if $\pi_i = 0$. The unit-level factual prediction loss \mathcal{L}_{y_i} is defined as

$$\mathcal{L}_{y_i} = loss(y_i, \hat{f}_y(\mathbf{t}_i = \pi_i, \hat{\mathbf{p}}_i = \hat{\phi}_e(\boldsymbol{\pi}_{-i}, \mathcal{G}, \boldsymbol{X}, \boldsymbol{Z}; \Theta_e), \hat{\mathbf{c}}_i = \hat{\phi}_f(v_i, \mathcal{G}, \boldsymbol{X}, \boldsymbol{Z}; \Theta_f); \Theta_y)), \tag{9}$$

where loss is an appropriate loss function (e.g., square error loss) based on data type of the outcome and $\Theta = \{\Theta_e, \Theta_f, \Theta_y\}$ are learning parameters to be optimized for exposure mapping function $\hat{\phi_e}$, feature mapping function $\hat{\phi_f}$, and counterfactual outcome model $\hat{f_y}$, respectively.

Balance loss. The CFR architecture uses autoencoder reconstruction loss and the Integral Probability Metric (IPM) (Shalit et al., 2017) measure of distance between treatment and control groups using Wasserstein (Cuturi and Doucet, 2014; Arjovsky et al., 2017), jointly referred to as balance loss, i.e.,

$$\mathcal{L}_{bal} = \mathbf{1}_{\lambda_{bal} > 0} \times \frac{1}{n} \sum_{i} (\mathbf{h}_{i}^{out} - \mathbf{c}_{i} || \boldsymbol{\rho}_{i})^{2} + \lambda_{bal} \times IPM(\{\mathbf{h}_{i}^{emb} : t_{i} = 1\}, \{\mathbf{h}_{i}^{emb} : t_{i} = 0\}), \quad (10)$$

where $\lambda_{bal} \geq 0$ is a hyperparameter and IPM(.) balances the distribution $\mathbb{P}(\mathbf{c}, \mathbf{p}|\mathbf{t} = 0)$ and $\mathbb{P}(\mathbf{c}, \mathbf{p}|\mathbf{t} = 0)$, where $\mathbb{P}(\mathbf{c}, \mathbf{p}|\mathbf{t})$ is equivalent to $\mathbb{P}(\mathbf{p}|\mathbf{t})\mathbb{P}(\mathbf{c}|\mathbf{p},\mathbf{t})$. Intuitively, \mathcal{L}_{bal} balances peer exposure distribution \mathbf{p} between treatment groups and covariate distribution \mathbf{c} across peer exposure conditions and treatment groups while maintaining expressiveness due to the autoencoder component.

For the end-to-end learning of $\hat{\phi}_e$, $\hat{\phi}_f$, and \hat{f}_{Y_i} , we introduce three custom loss functions designed for EGONETGNN: coverage loss, sparsity loss, and entropy loss. These custom loss functions serve as priors to make the learned exposure mapping function stable and reliable.

Coverage loss. We use a prior that encourages the bounded peer exposure embedding to have substantial coverage. This loss function checks how far the learned peer embedding distribution is from a continuous uniform distribution between 0 and 1, i.e., $L_{cov} = (mean(\rho) - 0.5)^2 + (var(\rho) - \frac{1}{12})^2 + (range(\rho) - 1)^2$. Here, we consider mean squared error of mean, variance, and range of learned embedding ρ against corresponding value of the uniform distribution.

Entropy loss and sparsity loss. Entropy loss encourages mask weights, i.e., $p := \sigma(\mathbf{W}_{mask})$ to take values toward 0 or 1 and sparsity loss pushes for a few weights with high values. Formally, we define entropy loss and sparsity loss as $\mathcal{L}_{ent} = mean(-plog(p) - (1-p)log(1-p))$ and $\mathcal{L}_{sp} = mean(p)$.

Overall loss. We obtain the overall loss function \mathcal{L} as

$$\mathcal{L} = \frac{1}{n} \sum_{i} \mathcal{L}_{y_i} + \mathcal{L}_{bal} + \lambda_{cov} \times \mathcal{L}_{cov} + \lambda_{ent} \times \mathcal{L}_{ent} + \lambda_{sp} \times \mathcal{L}_{sp} + \lambda_{L1} \times ||\Theta_{gnn}||_{1}, \quad (11)$$

where λ_{cov} , λ_{ent} , and λ_{sp} are the hyperparameters and Θ_{gnn} denote overall parameters in $\hat{\phi}_f$ and $\hat{\phi}_e$, and the last term is L_1 loss to promote invariance to irrelevant contexts by preferring sparse weights.

Inference. The peer effect is obtained as $\hat{\delta}_i(\boldsymbol{\pi}_{-i}, \boldsymbol{\pi}'_{-i}) = \hat{f}(\pi_i, \boldsymbol{\pi}_{-i}, \mathcal{G}, \mathbf{X}, \mathbf{Z}) - \hat{f}(\pi_i, \boldsymbol{\pi}'_{-i}, \mathcal{G}, \mathbf{X}, \mathbf{Z}) = \hat{f}_y(\boldsymbol{\pi}_i, \boldsymbol{\rho}_i, \mathbf{c}_i) - \hat{f}_y(\boldsymbol{\pi}_i, \boldsymbol{\rho}'_i, \mathbf{c}_i)$, where \hat{f} is the end-to-end EGONETGNN.

3.3 THEORETICAL ANALYSES OF EGONETGNN

Expressiveness. We perform a theoretical analysis of the expressive power of graph neural networks (GNNs) in capturing the causal network motifs proposed in the Yuan et al. (2021) paper. Building on previous research regarding the capacity of GNNs to count substructures (Chen et al., 2020), we demonstrate that existing message-passing GNN methods are not expressive enough to capture all causal network motifs. In contrast, our method is expressive to capture relevant causal network motifs. We defer the detailed theoretical framework and results to Appendix A.4.

Time complexity. Our analysis of runtime complexity included in Appendix A.4.3 shows our method is, roughly on average, $\rho_{\mathcal{E}} \times avg(\mathbf{d})$ times more computationally expensive than standard MPGNNs, where $\rho_{\mathcal{E}}$ is the average edge density and $avg(\mathbf{d})$ is the average degree.

Misspecification errors. We extend Shalit et al. (2017)'s analyses of theoretical counterfactual prediction error bounds for the CFR model to study misspecification errors in the end-to-end EGONETGNN using the sequential error decomposition trick in Appendix A.4.4. By focusing on learning the expressive exposure mapping function, we are reducing its misspecification error directly.

4 EXPERIMENTS AND RESULTS

4.1 EXPERIMENTAL SETUP

Dataset. Similar to other works in causal inference, we rely on synthetic and semi-synthetic data. We consider three synthetic network models with a fixed number of nodes (N=3000) with different data generating parameters and edge densities: (1) the Watts Strogatz (WS) network (Watts and Strogatz, 1998), which models small-world phenomena, (2) the Barabási Albert (BA) network (Albert and Barabási, 2002), which models preferential attachment phenomena, and (3) the Stochastic Block Model (SBM) that models community structures. We control the density of edges for BA and WS networks and the number of communities in the SBM network. We also use a real-world social networks BlogCatalog and Flickr with more realistic topology and attributes to generate treatments and outcomes. We defer additional details on data generation to Appendix A.5.

Evaluation metrics. To evaluate the performance of heterogeneous peer effect (HPE) estimation, we use the *Precision in the Estimation of Heterogeneous Effects* (ϵ_{PEHE}) (Hill, 2011) metric defined as $\epsilon_{PEHE} = \sqrt{\frac{1}{n} \sum_i (\delta_i(\pi_{-i}, \pi'_{-i}) - \hat{\delta}_i(\pi_{-i}, \pi'_{-i}))^2}$, where $\delta_i(\pi_{-i}, \pi'_{-i})$ is true HPE and $\hat{\delta}_i(\pi_{-i}, \pi'_{-i})$ is the estimated HPE, where π'_{-i} denotes a counterfactual scenario where treatments of peers are flipped. ϵ_{PEHE} (lower better) measures the deviation of estimated HPEs from true HPEs. For each experimental result, we report mean and standard deviation of ϵ_{PEHE} for 5 different simulations.

Baselines. We compare EGONETGNN with state-of-the-art (SOTA) peer estimation methods. NetEst (Jiang and Sun, 2022) and TNet (Chen et al., 2024) use the fraction of treated peers as peer exposure but the estimator is based on adversarial learning and doubly robust method, respectively, for robustness. DWR (Zhao et al., 2024) learns attention weights based on attribute similarity and 1GNN-HSIC (Ma and Tresp, 2021) use GNNs to summarize peer treatments as heterogeneous contexts while using homogeneous exposure. We also use the recently proposed GNN- and autoencoder-based automated exposure mapping approach (AEMNet) (Mao et al., 2025) and GNN- and transformer-based CauGramer (Wu et al., 2025) as baselines for estimating peer effects in our setup. We also consider INE-TARNet (Adhikari and Zheleva, 2025) adapted for peer effect estimation as a baseline, although it was developed for direct effect estimation. We include GNN-TARNet-Motifs approach that considers manually extracted causal network motifs (Yuan et al., 2021) as peer exposure and TARNet as estimator (Shalit et al., 2017) as a strong baseline. We discuss hyperparameter tuning and model selection in Appendix A.6.

4.2 RESULTS

Next, we present results for experimental setups designed to answer four research questions (RQs). **RQ1.** How well do methods for peer effect estimation perform when peer exposure mechanisms

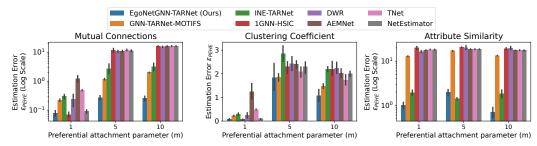


Figure 3: Peer effect estimation error for **Barabasi Albert** network when true peer exposure depends on mutual connections, clustering coefficient, and attribute similarity. Our method shows robust performance across different underlying peer influence mechanisms and edge densities (low to high).

Table 1: Mean and standard deviation of peer effect estimation error (ϵ_{PEHE}) for different methods in BlogCatalog (BC) dataset for four settings when true peer exposure mechanisms depend on clustering coefficients, connected components, mutual connections, and attribute similarity.

Mechanisms	Ours- TARNet	Ours- CFR	GNN- Motifs	INE- TARNet	1GNN- HSIC	DWR	AEMNet	TNet	NetEst	CauGramer
Clus. Coef.	$1.59_{\pm 0.4}$	$1.51_{\pm 0.3}$	$2.04_{\pm 0.7}$	$2.24_{\pm 0.6}$	$5.83_{\pm 3.7}$	$5.95_{\pm 1.8}$	$4.60_{\pm 2.2}$	$8.88_{\pm 9.5}$	$3.47_{\pm0.3}$	$6.31_{\pm 2.1}$
Con. Comp.	$2.98_{\pm 0.8}$	$2.77_{\pm 0.9}$	$4.41_{\pm 1.1}$	$3.93_{\pm 0.9}$	$6.83_{\pm 1.3}$	$6.84_{\pm 1.4}$	$8.05_{\pm 4.5}$	$11.60_{\pm 11.5}$	$7.38_{\pm0.8}$	$7.09_{\pm 1.0}$
Mut. Con.	$2.90_{\pm 1.1}$	$2.57_{\pm0.8}$	$3.83_{\pm 0.7}$	$3.37_{\pm 0.8}$	$9.29_{\pm 4.7}$	$8.21_{\pm 3.3}$	$11.71_{\pm 4.8}$	$13.15_{\pm 12.4}$	$6.62_{\pm 2.2}$	$8.24_{\pm 1.4}$
Attr. Sim.	$\frac{5.65}{\pm 0.7}$	$4.86 \scriptstyle{\pm 1.9}$	$6.09{\scriptstyle\pm0.2}$	$6.01_{\pm 2.0}$	$17.48_{\pm 9.1}$	$15.69_{\pm 6.7}$	$21.80_{\pm 9.3}$	$15.70_{\pm 7.7}$	$14.07_{\pm 5.8}$	$7.34_{\pm 3.8}$

Table 2: Mean and standard deviation HPE estimation error (ϵ_{PEHE} metric) for three variants of our method (original, without mask, and without feature encoder and mask) in the BlogCatalog (BC), Barabasi Albert (BA), and Watts Strogatz (WS) datasets for three true peer exposure mechanisms.

Mechanism Mutual Connections			Clustering Coefficient			Attribute Similarity			
Network	BC	BA	WS	BC	BA	WS	BC	BA	WS
Model Variants									
Ours	$2.90_{\pm 1.1}$	$0.20_{\pm0.1}$	$0.30_{\pm0.1}$	$1.59_{\pm0.4}$	$0.99_{\pm 0.9}$	$1.18{\scriptstyle\pm0.8}$	$5.65_{\pm0.7}$	$1.23_{\pm0.7}$	$1.09_{\pm 1.1}$
Ours (-mask)	$2.91_{\pm 1.5}$	$0.21_{\pm 0.1}$	$0.35_{\pm0.2}$	$1.98_{\pm 1.1}$	$1.01_{\pm 0.8}$	$1.20_{\pm 0.4}$	$6.17_{\pm 0.6}$	$2.37_{\pm 2.2}$	$1.29_{\pm 1.8}$
Ours (-feat&mask)	$2.90_{\pm0.7}$	$0.27_{\pm 0.2}$	$0.31_{\pm 0.1}$	$\boldsymbol{1.55}_{\pm0.5}$	$0.97_{\pm0.7}$	$1.91_{\pm 1.3}$	$5.99_{\pm0.8}$	$13.73_{\pm 2.8}$	$13.85_{\pm 4.0}$

Table 3: Evaluation of exposure representation, in terms of absolute correlation, in BlogCatalog data with no effect modification. The results for the learned peer exposure representation by our method is better (higher is better). We use the fraction of treated friends z_i as baseline and the dimension of $\hat{\rho}, \hat{\rho}' \in [0, 1]^{d=2}$ with highest correlation is shown.

Corr.	Clus. Coef.	Con. Comp.	Mut. Con.	Attr. Sim.	Corr.	Clus. Coef.	Con. Comp.	Mut. Con.	Attr. Sim.
$r(\hat{\rho}, \rho)$	$0.81_{\pm 0.1}$	$0.34_{\pm 0.3}$	$0.73_{\pm 0.2}$	$0.29_{\pm0.2}$	$r(\hat{\rho}', \rho')$	$0.85_{\pm 0.02}$	$0.30_{\pm 0.2}$	$0.74_{\pm 0.1}$	$0.50_{\pm 0.1}$
$r(z_i, \rho)$	$0.17_{\pm 0.1}$	$0.12_{\pm 0.1}$	$0.09_{\pm 0.03}$	$0.28_{\pm 0.2}$	$r(z_i', \rho')$	$0.41_{\pm 0.2}$	$0.14_{\pm 0.1}$	$0.09_{\pm0.1}$	$0.61_{\pm 0.1}$

depend on local neighborhood conditions? In this setup, we evaluate the performance of peer effect estimators when the underlying peer exposure mechanism is unknown. We generate treatments and outcomes such that there is confounding due to a subset of node attributes and mean peer attributes. For the outcome generation, we consider five mechanisms for true peer exposure conditions where peer exposure is given by 1) the clustering coefficient between the treated peers, 2) the number of connected components among treated peers, and weighted fraction of treated peers with weights as 3) the square root of number of mutual connections, 4) attribute similarity, and 5) tie strength. Here, the unit's treatment acts as an effect modifier, where the peer exposure is doubled if the unit is treated. Figure 3 shows peer effect estimation error (y-axis) when true peer exposure mechanisms depend on mutual connections, clustering coefficient, and attribute similarity in Barabasi Albert networks with three network generation parameters (x-axis) resulting different edge densities (low to high). The preferential attachment parameter m=1 produces a sparse star-topology network, lacking cycles or triangular structures. In this setting, all methods perform relatively well when peer exposure mechanisms depend on local structure because MPGNNs are expressive enough to capture star-shaped motifs. However, with increased edge density and more complex network topology, unlike our method, the baselines are not sufficiently expressive to capture underlying mechanisms and suffer significantly. The GNN-TARNet-Motifs (GTM) approach is expressive in capturing clustering coefficients, and both GTM and INE-TARNet approximate mutual connections. This is reflected in the performance, where GTM is competitive for the clustering coefficient peer exposure mechanism.

EGONETGNN-TARNet outperforms the baselines except for INE-TARNet, which is competitive in a setting with the peer exposure mechanism dependent on attribute similarity. Figure 3 and other results in Appendix A.7 show that for unknown peer exposure mechanisms, our method is as expressive as or superior to the strongest baseline with significantly better performance for denser networks.

RQ2. How reliable are the models for heterogeneous peer effect estimation in more realistic scenario? RQ2 investigates the performance of the models using more realistic semi-synthetic networks and node attributes. In addition to confounding and heterogeneous peer influence, there is more complex peer effect modification depending on whether the unit is treated and the values of the unit's attributes. Table 1 shows mean and standard deviation of peer effect estimation error (ϵ_{PEHE}) for different methods in BlogCatalog (BC) dataset for four settings when true peer exposure mechanisms depend on clustering coefficients, connected components, mutual connections, and attribute similarity. The results show the robustness of EGONETGNN in a more realistic setting, where EGONETGNN-CFR outperforms all baselines. EGONETGNN-TARNet is competitive with the best model and it also outperforms the baselines. In this setup, peer effects are heterogeneous due to the interaction of peer exposure conditions and effect modifiers, and our method is able to approximate it better than the baselines. Appendix A.8 presents additional experiments for this setup including results for the Flickr dataset (Table 4) which is more challenging for the baselines.

RQ3. How do the components of EGONETGNN contribute to its robustness in estimating peer effects? We conduct ablation studies to assess the contributions of masked weights and the feature encoder MLP. Table 2 displays the performance of three variants of EGONETGNN-TARNet (original, without the masked weights, and without the feature encoder and masked weights) across BlogCatalog (BC), Barabási Albert (BA), and Watts-Strogatz (WS) datasets. The results show that excluding masked weights can bias peer effect estimates due to the model's sensitivity to irrelevant contexts. Removing the feature encoder MLP limits EgoNetGNN's ability to capture mechanisms based on attribute similarity. However, certain peer exposure mechanisms relying on local structures perform better when irrelevant features are ignored. Overall, these findings demonstrate that the feature encoder MLP enhances expressiveness, while masked weights promote invariance to irrelevant contexts. Additionally, we analyze the EGONETGNN's sensitivity to the choices of peer exposure embedding dimension, coverage loss coefficient, and noisy networks in Appendix A.9.

RQ4. How well are underlying mechanisms captured by learned exposure mapping function? In Table 3, we directly compare the (absolute) Pearson correlation coefficient r (higher is better) between the learned peer exposure representation, $\hat{\rho}$ and $\hat{\rho}'$, and the actual peer exposure under four different mechanisms. Compared to commonly used fraction of treated friends baseline, learned peer exposures are informative of true peer exposures for mechanisms involving local structure.

5 DISCUSSION, LIMITATIONS & FUTURE WORK

Our work motivates the problem of learning exposure mapping function for peer effect estimation and proposes EGONETGNN for addressing unknown peer influence mechanisms involving local neighborhood conditions. Our theoretical analysis and experimental results demonstrate increased expressiveness of EGONETGNN to capture complex local neighborhood exposure conditions. We have designed EGONETGNN to promote invariance to irrelevant contexts, and output a low-dimensional peer exposure embedding with bounded and balanced representation to partially mitigate issue of potential violation of the positivity assumption with continuous treatment or exposure. The empirical results have shown the effectiveness of EGONETGNN in many peer effect estimation settings.

Limitations & Future Work. Ensuring theoretical bounds for variance with complex GNNs for heterogeneous causal effect estimation is still a developing research area (Khatami et al., 2024) and important future direction, but it is not within the scope of our current work. This work can be extended to incorporate other network effects like direct effects and total effects. The increased expressiveness and robust peer effect estimates of our model come with the trade-off of a slightly longer runtime to process ego networks. Future work could consider relaxing the assumption of interference from immediate peers while addressing the scalability. Our work relies upon a reliable attributed network as input, but future research should consider capturing expressive representations in noisy networks. Appendix A.1 discusses societal impacts, scalability, and plausibility of assumptions.

REPRODUCIBILITY STATEMENT

To support reproducibility, we release the complete codebase and experimental procedures. For all the experiments, we have repeated them at least *five* times. We provide the details of the data generation process (Sec. 4.1 and Appendix A.5). Appendix starts with an anonymous repository link containing the full source code. We provide the details of the configurations and setups for replicating our results in Appendix A.6.

REFERENCES

- Shishir Adhikari and Elena Zheleva. 2025. Inferring Individual Direct Causal Effects Under Heterogeneous Peer Influence. *Machine Learning Journal* (2025).
- Réka Albert and Albert-László Barabási. 2002. Statistical mechanics of complex networks. *Reviews of modern physics* 74, 1 (2002), 47.
- Joshua D Angrist, Guido W Imbens, and Donald B Rubin. 1996. Identification of causal effects using instrumental variables. *Journal of the American statistical Association* 91, 434 (1996), 444–455.
- David Arbour, Dan Garant, and David Jensen. 2016. Inferring network effects from observational data. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 715–724.
- Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein generative adversarial networks. In *International conference on machine learning*. PMLR, 214–223.
- Peter M Aronow and Cyrus Samii. 2017. Estimating average causal effects under general interference, with application to a social network experiment. *The Annals of Applied Statistics* 11, 4 (2017), 1912–1947.
- Eric Auerbach, Jonathan Auerbach, and Max Tabord-Meehan. 2024. Exposure effects are not automatically useful for policymaking. *arXiv preprint arXiv:2401.06264* (2024).
- Falco J Bargagli-Stoffi, Costanza Tortú, and Laura Forastiere. 2025. Heterogeneous treatment and spillover effects under clustered network interference. *The annals of applied statistics* 19, 1 (2025), 28.
- Brian G Barkley, Michael G Hudgens, John D Clemens, Mohammad Ali, and Michael E Emch. 2020. Causal inference from observational studies with clustered interference, with application to a cholera vaccine study. *Annals of Applied Statistics* 14, 3 (2020), 1432–1448.
- Derrick Blakely, Jack Lanchantin, and Yanjun Qi. [n.d.]. Time and space complexity of graph convolutional networks. https://qdata.github.io/deep2Read/talks-mb2019/Derrick_201906_GCN_complexityAnalysis-writeup.pdf.
- David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of machine Learning research* 3, Jan (2003), 993–1022.
- Ruichu Cai, Zeqin Yang, Weilin Chen, Yuguang Yan, and Zhifeng Hao. 2023. Generalization bound for estimating causal effects from observational network data. In *CIKM*. 163–172.
- Weilin Chen, Ruichu Cai, Zeqin Yang, Jie Qiao, Yuguang Yan, Zijian Li, and Zhifeng Hao. 2024. Doubly Robust Causal Effect Estimation under Networked Interference via Targeted Learning. In *Forty-first International Conference on Machine Learning*.
- Zhengdao Chen, Lei Chen, Soledad Villar, and Joan Bruna. 2020. Can graph neural networks count substructures? *Advances in neural information processing systems* 33 (2020), 10383–10395.
- Marco Cuturi and Arnaud Doucet. 2014. Fast computation of Wasserstein barycenters. In *International conference on machine learning*. PMLR, 685–693.
 - Laura Forastiere, Edoardo M Airoldi, and Fabrizia Mealli. 2021. Identification and estimation of treatment and interference effects in observational studies on networks. *J. Amer. Statist. Assoc.* 116, 534 (2021), 901–918.

- Jennifer L Hill. 2011. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics* 20, 1 (2011), 217–240.
- Michael G Hudgens and M Elizabeth Halloran. 2008. Toward causal inference with interference. *J. Amer. Statist. Assoc.* 103, 482 (2008), 832–842.
 - Daniel Jiwoong Im, Kyunghyun Cho, and Narges Razavian. 2021. Causal effect variational autoencoder with uniform treatment. *arXiv preprint arXiv:2111.08656* (2021).
 - Song Jiang and Yizhou Sun. 2022. Estimating Causal Effects on Networked Observational Data via Representation Learning. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 852–861.
 - Seyedeh Baharan Khatami, Harsh Parikh, Haowei Chen, Sudeepa Roy, and Babak Salimi. 2024. Graph Machine Learning based Doubly Robust Estimator for Network Causal Effects. arXiv:2403.11332 [cs.LG] https://arxiv.org/abs/2403.11332
 - Thomas N Kipf and Max Welling. 2016. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations*.
 - Michael P Leung and Pantelis Loupos. 2022. Unconfoundedness with network interference. *arXiv* preprint arXiv:2211.07823 6 (2022).
 - Xiaofeng Lin, Guoxi Zhang, Xiaotian Lu, Han Bao, Koh Takeuchi, and Hisashi Kashima. 2023. Estimating Treatment Effects Under Heterogeneous Interference. In *Joint European Conference on ML and KDD*. Springer, 576–592.
 - Xiaofeng Lin, Guoxi Zhang, Xiaotian Lu, and Hisashi Kashima. 2024. Treatment Effect Estimation Under Unknown Interference. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 28–42.
 - Jing Ma, Mengting Wan, Longqi Yang, Jundong Li, Brent Hecht, and Jaime Teevan. 2022. Learning causal effects on hypergraphs. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1202–1212.
 - Yunpu Ma and Volker Tresp. 2021. Causal inference under networked interference and intervention policy enhancement. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 3700–3708.
 - Yunxin Mao, Haotian Wang, Yishuai Cai, Minglong Li, Ji Wang, and Wenjing Yang. 2025. Automated Exposure Mapping for Networked Interference. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 1–5.
 - Wang Miao, Xu Shi, Yilin Li, and Eric J Tchetgen Tchetgen. 2024. A confounding bridge approach for double negative control inference on causal effects. Statistical Theory and Related Fields 8, 4 (2024), 262–273.
 - Razieh Nabi, Joel Pfeiffer, Denis Charles, and Emre Kıcıman. 2022. Causal inference in the presence of interference in sponsored search advertising. *Frontiers in big Data* 5 (2022).
 - Elizabeth L Ogburn, Oleg Sofrygin, Ivan Diaz, and Mark J Van der Laan. 2022. Causal inference for social network data. *J. Amer. Statist. Assoc.* (2022), 1–15.
 - Eleonora Patacchini, Edoardo Rainone, and Yves Zenou. 2017. Heterogeneous peer effects in education. *Journal of Economic Behavior & Organization* 134 (2017), 190–227.
 - Judea Pearl. 2009. *Causality*. Cambridge university press.
- Zhaonan Qu, Ruoxuan Xiong, Jizhou Liu, and Guido Imbens. 2021. Efficient Treatment Effect Estimation in Observational Studies under Heterogeneous Partial Interference. *arXiv preprint arXiv:2107.12420* (2021).
 - Fredrik Sävje. 2024. Causal inference with misspecified exposure mappings: separating definitions and assumptions. *Biometrika* 111, 1 (2024), 1–15.

- Uri Shalit, Fredrik D Johansson, and David Sontag. 2017. Estimating individual treatment effect: generalization bounds and algorithms. In *International Conference on Machine Learning*. PMLR, 3076–3085.
- Claudia Shi, David Blei, and Victor Veitch. 2019. Adapting neural networks for the estimation of treatment effects. *Advances in neural information processing systems* 32 (2019).
- Eric J Tchetgen Tchetgen, Andrew Ying, Yifan Cui, Xu Shi, and Wang Miao. 2020. An introduction to proximal causal learning. *arXiv preprint arXiv:2009.10982* (2020).
- Christopher Tran and Elena Zheleva. 2022. Heterogeneous Peer Effects in the Linear Threshold Model. *Proceedings of the AAAI Conference on Artificial Intelligence* (2022).
- Johan Ugander, Brian Karrer, Lars Backstrom, and Jon Kleinberg. 2013. Graph cluster randomization: Network exposure to multiple universes. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*. 329–337.
- Duncan J Watts and Steven H Strogatz. 1998. Collective dynamics of 'small-world' networks. *nature* 393, 6684 (1998), 440–442.
- Anpeng Wu, Haiyi Qiu, Zhengming Chen, Zijian Li, Ruoxuan Xiong, Fei Wu, and Kun Zhang. 2025. Causal Graph Transformer for Treatment Effect Estimation Under Unknown Interference. In *The Thirteenth International Conference on Learning Representations*.
- Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2018. How Powerful are Graph Neural Networks?. In *International Conference on Learning Representations*.
- Yuan Yuan, Kristen Altenburger, and Farshad Kooti. 2021. Causal Network Motifs: Identifying Heterogeneous Spillover Effects in A/B Tests. In *Proceedings of the Web Conference* 2021. 3359–3370.
- Ziyu Zhao, Yuqi Bai, Ruoxuan Xiong, Qingyu Cao, Chao Ma, Ning Jiang, Fei Wu, and Kun Kuang. 2024. Learning Individual Treatment Effects under Heterogeneous Interference in Networks. *ACM Trans. Knowl. Discov. Data* 18, 8, Article 199 (Aug. 2024), 21 pages. https://doi.org/10.1145/3673761

A APPENDIX

Source code and documentation are available at: https://anonymous.4open.science/r/EgoNetGNN-8D5C/

A.1 DISCUSSION

Societal impacts. The implications of our work include identifying unit-level peer effects and discovering subpopulations with heterogeneous peer effects. The potential societal impacts could include the development of targeted interventions or the identification of policies that enhance desired outcomes in social networks.

Plausibility of neighborhood interference assumption. Neighborhood interference (Assumption 2 in Sec. 2) is a common simplifying assumption and can be realistic in situations where peer interference is mediated by immediate neighbors or diminishes quickly for non-immediate neighbors. However, there could be some situations where interference could occur between peers beyond immediate neighbors. If we assume such interference is mediated via immediate neighbors, then we could model it by stacking multiple exposure mapping function learning layers, where the subsequent layers would summarize the exposures of neighbors. Another alternative is to use the K-hop ego network with edge existence and/or hop distance as additional node features. The former approach may be more scalable than the latter one because the K-hop neighborhood can grow rapidly. Ideas from recent works to infer unknown interference structure (Wu et al., 2025; Lin et al., 2024) could be adopted in conjunction with our approach of learning expressive peer exposure representations. While we assume a reliable network structure is provided as input, our experiments with noisy networks reveal that EGONETGNN performs reliably well with imperfect data.

Plausibility of unconfoundedness assumption. Following existing work in the intersection of causal reasoning and representation learning (Shalit et al., 2017; Shi et al., 2019; Ma et al., 2022; Wu et al., 2025), we assume causal identification conditions are met and focus on expressive representation learning to mitigate model misspecification errors. Unconfoundness is a strong and untestable assumption and requires sufficiency of observed network contexts and expressiveness of their representation. While we assume the sufficiency of observed contexts, we make an effort to satisfy the expressiveness of representation by considering all network contexts, like node attributes, edge attributes, and network structure. If the presence of unobserved confounding cannot be ruled out, alternative causal identification approaches like proximal causal inference (Tchetgen et al., 2020) or double negative controls (Miao et al., 2024), front-door criteria (Pearl, 2009), and instrumental variables (Angrist et al., 1996) should be considered. Although a randomized experiment can remove unobserved confounding between unit treatments and the outcome, peer exposure conditions may not be randomized directly, and confounding could exist even for experiments unless the unconfoundedness assumption is made and observed network contexts are controlled for. So, an interesting future direction could be to explore alternative identification conditions.

Scalability. Although EGONETGNN is more expressive, it has additional computational costs. A few ways to address large runtime and/or memory usage could be sampling ego networks to reduce the training set or sampling the neighborhood within a K-hop ego network. In Appendix A.8 (Table 6), our experiments with a randomly augmented network show that the performance does not degrade significantly for our method with the removal of edges. From an implementation point of view, we can parallelize our framework easily to exploit the power of GPUs. More specifically, there are two components in our framework. The feature mapping GNN takes the entire network at once to learn an embedding with an L-layer GNN. Subsequently, EgoNetGNN batches B nodes with their neighbor nodes and a mapping of which edges belong to which node in the batch. This batching can be parallelized to improve the overall efficiency.

A.2 RELATED WORK

Research in causal inference under interference has focused on estimating three main causal effects of interest, referred to as network effects: direct effects induced by a unit's own treatment, peer effects induced by treatment of other units, and total effects induced by both the unit's and others' treatment (Hudgens and Halloran, 2008). These network effects are estimated as average effects (e.g., (Arbour et al., 2016; Ugander et al., 2013)) for the entire population or as heterogeneous

effects (e.g., (Forastiere et al., 2021; Bargagli-Stoffi et al., 2025)) for specific subpopulations or contexts. Our work focuses on heterogeneous peer effect estimation. Most methods for estimating heterogeneous or individual-level causal effects under interference, including peer effects, assume peer exposure is binary (Bargagli-Stoffi et al., 2025) or homogeneous, e.g., based on fraction of treated peers (Jiang and Sun, 2022; Ogburn et al., 2022; Cai et al., 2023; Chen et al., 2024). These methods assume a homogeneous or known exposure mapping function and focus on enhancing network effect estimation by adapting techniques like adversarial training (Jiang and Sun, 2022), propensity score reweighting (Cai et al., 2023), double machine learning (Khatami et al., 2024), doubly robust estimation (Leung and Loupos, 2022), targeted maximum likelihood estimate (Ogburn et al., 2022), and targeted learning (Chen et al., 2024).

Recent research has looked into more complex functions of peer exposure, allowing for heterogeneous peer influence, in which different peers can have varying degrees of influence. Some of these works refer to heterogeneous peer influence as heterogeneous interference (Qu et al., 2021; Zhao et al., 2024; Lin et al., 2023). Forastiere et al. (2021) considered peer exposure as a weighted fraction of treated peers using known edge attributes as weights. Lin et al. (2023) consider heterogeneity due to multiple entities types and Qu et al. (2021) considered heterogeneity due to known node attributes for defining peer exposure. Tran and Zheleva (2022) studied peer effect estimation with linear threshold peer exposure model but different unit-level threshold could be vary for different units capturing heterogeneous susceptibilities to the influence. Zhao et al. (2024) used attention weights derived based on the similarities of the units' covariates to determine peer exposure as the weighted sum of treated peers. Yuan et al. (2021) capture peer exposure with features based on counts of different causal network motifs, i.e., recurrent subgraphs in a unit's ego network with treatment assignments as attributes. Ma and Tresp (2021) consider homogeneous peer exposure based on fraction of treated peers but they summarize the covariates of treated peers using a graph neural network (GNN) to capture heterogeneous contexts involving treatment assignments. Unlike our work, none of these studies has explicitly studied the issue of automatically learning the exposure mapping functions to define peer exposure representation while capturing the underlying influence mechanisms.

Ma and Tresp (2021) learn heterogeneous contexts based on peer treatments but not the exposure mapping function or the peer exposure representation. Zhao et al. (2024) obtain single-dimension peer exposure embedding using a weighted sum of treated peers with attention weights derived from the cosine similarity of feature embeddings. Although Zhao et al. (2024) use attention weights to define peer exposure, they assume a specific exposure mapping function, and it cannot adapt according to the underlying peer influence mechanism. Adhikari and Zheleva (2025) use GNNs to learn peer exposure embedding by addressing unknown peer influence mechanisms, but their scope is limited to direct effect estimation, i.e., the effect of a unit's own treatment. Specifically, Adhikari and Zheleva (2025) learn a multi-dimensional peer exposure embedding using a weighted fraction of treated peers with feature embeddings and a second-order adjacency matrix as weights. Ma et al. (2022) employ similar method like Ma and Tresp (2021) for hypergraphs to model heterogeneity due to model group interactions. The idea is to learn a summary function and representation equivalent to the exposure mapping function and peer exposure using a hypergraph convolution network and attention mechanism. However, they assume the learned representation is expressive enough to capture the underlying influence mechanism. In this work, we do not make such an assumption and evaluate how well the learned peer exposure representation captures the underlying influence mechanisms.

Neural networks (NNs) (Shalit et al., 2017; Im et al., 2021; Shi et al., 2019) and, recently, graph neural networks (GNNs) (Jiang and Sun, 2022; Cai et al., 2023; Chen et al., 2024; Khatami et al., 2024) have been widely utilized for end-to-end learning of *feature mapping function* and *counterfactual outcome model* or *effect estimator*. A feature mapping function maps raw features to feature embedding to capture potential confounders and effect modifiers. A counterfactual outcome model (Shalit et al., 2017; Ma and Tresp, 2021) predicts counterfactual outcomes for different levels of treatment, while an effect estimator (Shi et al., 2019; Chen et al., 2024) directly learns the causal effect of interest. Only a few studies have considered learning the exposure mapping function (Mao et al., 2025) or peer exposure embedding (Adhikari and Zheleva, 2025; Zhao et al., 2024). Lin et al. (2024) consider a setting with an unknown network and interference structure and propose an approach to first infer network structure and represent peer exposure for direct effect estimation. Unlike their work, our settings focus on peer effect estimation with observed network structure but unknown peer exposure mechanisms that manifest due to local neighborhood contexts.

Sävje (2024) advocates for interpretable but possibly misspecified exposure mappings and characterizes causal estimation errors due to misspecified exposure mappings, but follow-up research (Auerbach et al., 2024) has highlighted the importance of capturing underlying interference mechanisms in policymaking. More recently, Mao et al. (2025) have explored the use of GNNs with autoencoders and clustering to learn discrete exposure conditions and their probabilities, aiming to estimate overall causal effects in networks. Similarly, Wu et al. (2025) utilize GNNs with Transformers to model unknown interference from K-hop neighborhood. Their identifiability assumption relies on capturing unit and peer covariates, while our identifiability assumption relies on capturing all attributed network contexts, including structure and edge attributes. These works use off-the-shelf message passing GNNs (like GCN and GIN) and lack expressiveness to capture mechanisms involving local neighborhood structure. Prior research (Xu et al., 2018; Chen et al., 2020) on the expressiveness of GNNs has shown that popular GNN architectures lack expressiveness to count subgraphs. On the other hand, counts of subgraphs like causal network motifs are rich features that could capture underlying influence mechanisms due to local neighborhood structure (Yuan et al., 2021). Counting such subgraphs can be computationally expensive, and they may not be able to capture every local structure. We design EGONETGNN to excel in counting attributed triangle subgraphs, enhancing its expressiveness to capture underlying mechanisms involving neighborhood contexts.

A.3 Causal Inference Assumptions and Identification of Peer Effects

A fundamental prerequisite for causal identification is the consistency assumption, which enables equivalence among counterfactual, interventional, and factual outcomes.

Assumption 4 (Consistency under interference). The underlying outcome generation is independent of the treatment assignment mechanisms (i.e., hypothetical, experimental, or natural). For a unit v_i , if $t_i = \pi_i$ and $t_{-i} = \pi_{-i}$, then $y_i(t_i = \pi_i, t_{-i} = \pi_{-i}) = y_i$.

Positivity is another standard assumption in causal inference that requires every unit v_i to have non-zero probability of being assigned every possible unit treatment and peer exposure conditions.

Assumption 5 (Positivity). There is a non-zero probability of unit treatment and peer exposure conditions for all possible contexts \mathbf{c}_i , i.e., $\mathbb{P}(\mathbf{t}_i, \mathbf{p}_i | \mathbf{c}_i) > 0$, for every level of \mathbf{t}_i and \mathbf{p}_i , where \mathbb{P} is the probability density function.

The proof of Proposition 1 is as follows.

Proof. Our causal estimand of interest (Eq. 2) is as follows:

$$\delta_i(\boldsymbol{\pi}_{-i}, \boldsymbol{\pi}'_{-i}) = \mathbb{E}[\mathbf{y}_i(\mathbf{t}_i = \pi_i, \mathbf{p}_i = \phi_e(\boldsymbol{\pi}_{-i}, \mathcal{G}, \mathbf{X}, \mathbf{Z})) | \mathbf{c}_i] - \mathbb{E}[\mathbf{y}_i(\mathbf{t}_i = \pi_i, \mathbf{p}_i = \phi_e(\boldsymbol{\pi}'_{-i}, \mathcal{G}, \mathbf{X}, \mathbf{Z})) | \mathbf{c}_i].$$

Due to unconfoundedness assumption (Assumption 3), unit treatment and peer exposure conditions are independent of counterfactual outcome conditioned on network contexts c_i . This allows us to rewrite the estimand as:

$$\delta_i(\boldsymbol{\pi}_{-i}, \boldsymbol{\pi}'_{-i}) = \mathbb{E}[\mathbf{y}_i(\mathbf{t}_i = \boldsymbol{\pi}_i, \mathbf{p}_i = \phi_e(\boldsymbol{\pi}_{-i}, \mathcal{G}, \mathbf{X}, \mathbf{Z})) | \mathbf{t}_i = \boldsymbol{\pi}_i, \mathbf{p}_i = \phi_e(\boldsymbol{\pi}_{-i}, \mathcal{G}, \mathbf{X}, \mathbf{Z})), \mathbf{c}_i] - \mathbb{E}[\mathbf{y}_i(\mathbf{t}_i = \boldsymbol{\pi}_i, \mathbf{p}_i = \phi_e(\boldsymbol{\pi}'_{-i}, \mathcal{G}, \mathbf{X}, \mathbf{Z})) | \mathbf{t}_i = \boldsymbol{\pi}_i, \mathbf{p}_i = \phi_e(\boldsymbol{\pi}_{-i}, \mathcal{G}, \mathbf{X}, \mathbf{Z})), \mathbf{c}_i].$$

Here, Assumption 1 ensures introducing new terms related to treatment and peer exposure in the conditional does not affect existing set of contexts because they are measured pre-treatment. Similarly, Assumption 2 makes the sufficiency of learned representation requirement in unconfoundedness assumption more plausible. Next, the consistency assumption allows replacing the counterfactual outcome with observed outcome, i.e.,

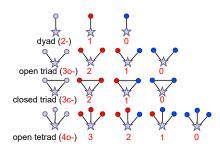
$$\delta_i(\boldsymbol{\pi}_{-i}, \boldsymbol{\pi}'_{-i}) = \mathbb{E}[\mathbf{y}_i | \mathbf{t}_i = \pi_i, \mathbf{p}_i = \phi_e(\boldsymbol{\pi}_{-i}, \mathcal{G}, \mathbf{X}, \mathbf{Z})), \mathbf{c}_i] - \mathbb{E}[\mathbf{y}_i | \mathbf{t}_i = \pi_i, \mathbf{p}_i = \phi_e(\boldsymbol{\pi}_{-i}, \mathcal{G}, \mathbf{X}, \mathbf{Z})), \mathbf{c}_i].$$

Assumption 1 also ensures consistency assumption is satisfied because the treatments are not mutable. This estimation above is tractable from observational or experimental data because of positivity assumption and the causal effects can be identified. $\hfill \Box$

A.4 THEORETICAL ANALYSES OF EGONETGNN

A.4.1 PRELIMINARIES

Causal network motifs. Yuan et al. (2021) proposed causal network motifs as important features to capture peer exposure accounting for local neighborhood conditions. Causal network motifs are



817

818

819

820

821

822

823

824

825

826

827

828

829 830 831

832

833

834

835

836

837

838

839

840

841

842

843

844

845

846

847

848

849

850

851

852 853

854

855

856

857

858

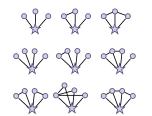
859

860 861

862

863

Figure 4: Example causal network motifs considered by Yuan et al. (2021). Stars represent ego nodes and circles represent their peers. The red circles indicate treated nodes and blue circles indicate control nodes. The gray shapes indicate nodes that could either be treated or control. Here, the characters in red indicate a particular causal network motif (e.g., 3c-2 indicate closed triad with 2 treated peers).





Examples of higher-order network motifs network motifs that cannot that can be captured by EgoNetGNN

Examples of higher-order be capturedby EgoNetGNN

Figure 5: Examples of higher-order network motifs with four and five nodes. Stars represent ego nodes and circles represent their peers. The gray shapes indicate nodes with any treatment assignment. If the subgraph of a network motif, after removing edges connected to the ego node, forms a tree, then our model is expressive enough to capture the network motif and the corresponding causal network motifs. A network motif is a subgraph without any attributes, whereas a causal network motif is a subgraph that includes peer treatment assignments as attributes.

attributed subgraphs with peer treatments as attributes. Figure 4 shows four categories of causal network motifs: dyads, open triads, closed triads, and open tetrads. In the figure, stars represent ego nodes and circles represent their peers. The red circles indicate treated nodes and blue circles indicate control nodes. The gray shapes indicate nodes that could either be treated or control.

Message passing graph neural networks (MPGNNs). The message-passing graph neural network (MPGNN) is a generic GNN model that incorporates several standard GNN architectures and relies on local aggregations of information within graphs (Chen et al., 2020). For a graph $G(V, E, \mathbf{X}, \mathbf{Z})$, an MPGNN with L layers is defined iteratively with aggregate function AGG^l and update function U^l as follows:

$$h_i^l = U^l(h_i^{l-1}, AGG_{j \in \mathcal{N}_i}^l(\Theta^l(h_i^{l-1}, h_i^{l-1}, Z_{ij}))), \tag{12}$$

where \mathcal{N}_i denotes neighbors of unit v_i and Θ^l denote learnable parameters like multi-layer perceptron. To obtain the hidden state at the l^{th} layer, a local aggregation of the previous layer's hidden states $(h_i^{l-1} \text{ and } h_i^{l-1})$ and, optionally, edge attributes Z_{ij} is performed and then combined with h_i^{l-1} . The hidden states are initialized as node attributes, i.e., $h_i^0 = X_i$. Typically, in various GNN architectures, the update and aggregation functions are chosen as part of architecture design.

Expressiveness of MPGNNs in counting substructures. Here, we summarize the results obtained by Chen et al. (2020) that are relevant to our theoretical analysis. We list their findings after defining relevant concepts.

Definition 2 (Subgraph). A subgraph $G^{[S]}(V^{[S]}, E^{[S]})$) of a graph G(V, E) consists of subsets of its nodes, i.e., $V^{[S]} \subseteq V$ and edges, i.e., $E^{[S]} \subseteq E$.

Definition 3 (Induced subgraph). A induced subgraph $G^{[S']}(V^{[S']}, E^{[S']})$ of a graph G(V, E) consists of subset of its nodes, i.e., $V^{[S']} \subseteq V$ and all edges between nodes $V^{[S']}$, i.e., $E^{[S']} = E \cap V^{[S']}$.

All induced subgraphs are subgraphs but reverse is not true. For example, all causal network motfis are induced subgraphs (and subgraphs) of the original graph. An open triad motif is a subgraph, but not an induced subgraph, of a closed triad motif.

Definition 4 (Star-shaped pattern). A pattern $G^{[P]}(V^{[P]}, E^{[P]})$ is a star-shaped pattern if it can be represented by a tree structure.

Definition 5 (Connected pattern). A pattern $G^{[P]}(V^{[P]}, E^{[P]})$ is a connected pattern if it cannot be represented by a tree structure.

For example, a closed triad motif is a connected pattern and dyads, open triads, and open tetrads are star-shaped patterns.

Chen et al. (2020) obtain the following results on the expressiveness of MPGNNS for counting substructures.

Corollary 3.4. (Chen et al., 2020) MPGNNs cannot *induced-subgraph-count* any *connected pattern* with 3 or more nodes.

Theorem 3.5. (Chen et al., 2020) MPGNNs can perform *subgraph-count* of *star-shaped patterns*.

A.4.2 EXPRESSIVENESS OF EGONETGNN

Here, we demonstrate that standard MPGNNs lack the expressiveness to capture closed triad motifs, and our model addresses this limitation.

Without loss of generality, assume node attributes for each node v_i are $<1,T_i>$ and constant edge attributes <1>.

Definition 6 (Expressiveness in counting causal network motifs). Let \mathcal{G} be a space of graphs. A representation by an MPGNN f is expressive in counting causal network motif $G^{[P]}$ if, for all ego networks $G^{[1]}, G^{[2]} \in \mathcal{G}$, distinct counts, i.e., $C_I(G^{[1]}, G^{[P]}) \neq C_I(G^{[2]}, G^{[P]})$, get distinct representations, i.e., $f(G^{[1]}) \neq f(G^{[2]})$, where C_I returns induced-subgraph-count of pattern $G^{[P]}$.

Proposition 2 (Expressiveness of EGONETGNN). EGONETGNN is expressive enough to capture all dyad, open triad, closed triad, and open tetrad causal network motifs.

Proof. We proceed the proof by dividing the statement into following two claims.

Claim 1. EGONETGNN is as expressive as standard MPGNN in capturing dyad, open triad, and open tetrad causal network motifs.

Proof. The dyad, open triad, and open tetrad causal network motifs are star-shaped patterns, and these patterns can be counted by standard MPGNNs (Chen et al. (2020)'s Theorem 3.5.). Our model employs MPGNN (refer Eq. 5 and Figure 2) on a transformed graph, where all edges connected to the ego node are removed, and the corresponding edge attributes from the removed edges are included as node attributes in the transformed graph. We need to show that this transformation preserves the expressiveness to capture dyad, open triad, and open tetrad causal network motifs. The dyad, open triad, and open tetrad causal network motifs are transformed into subgraphs with isolated one, two, and three nodes, respectively, in the transformed ego network. MPGNN in the transformed graph can perform a subgraph count of patterns with k isolated nodes because they are subgraphs of star-shaped patterns with an empty set of edges. Furthermore, the addition of new attributes does not affect the expressiveness because these attributes are added as additional feature dimensions. Hence, our model is as expressive as standard MPGNN for capturing dyad, open triad, and open tetrad causal network motifs.

Claim 2. EGONETGNN also captures closed triad causal network motifs.

Proof. The closed triad causal network motifs are connected patterns of three nodes and these patterns cannot be counted by standard MPGNNs (Chen et al. (2020)'s Corollary 3.4.). Due to the construction of the ego network, all the edges with the ego node are removed, and the closed triads are transformed to dyads in the transformed ego network. These dyads can be counted by node aggregation (refer Eq. 5), which is an MPGNN employed in the ego network. Therefore, EGONETGNN captures closed triad causal network motifs. □

Higher-order causal network motifs and attributed causal network motifs. Here, we show how our model is superior to the approach of counting predetermined causal network motifs by discussing EGONETGNN's ability to capture relevant causal network motifs including higher-order and attributed causal network motifs. Proposition 2 showed our model is as expressive as the approach of counting predetermined causal network motifs considered by Yuan et al. (2021). In general, if the subgraph of a network motif, after removing edges connected to the ego node, forms a tree, then EGONETGNN is expressive enough to capture the network motif and the corresponding causal network motifs. Figure

5 depicts some examples of higher-order motifs with four and five nodes. EGONETGNN, with depths of L=2 and L=3 (refer Eq. 5), is expressive enough to capture most higher-order motifs with four and five nodes, respectively. Only if the network motifs consist of a cycle without the involvement of the ego node, then EGONETGNN is not expressive enough to capture it. Furthermore, compared to predetermined causal network motifs, EGONETGNN can accommodate motifs with additional node and edge attributes. Incorporating node and edge attributes will not reduce the expressiveness of counting original causal network motifs because these attributes are added as additional feature dimensions.

A.4.3 TIME COMPLEXITY OF EGONETGNN-TARNET

Typically, the complexity of a standard MPGNN (e.g. GCN), is $O(NLF^2 + L|E|F)$, where N,|E|, L, and F are the number of nodes, edges, GNN layers, and the dimensionality of feature embeddings, respectively (Blakely et al., [n. d.]). In our model, the feature mapping MPGNN (refer to Eq. 4) has the time complexity of $O(d_{\Theta}NF_x^2)$ for ego feature embedding module $\Theta_0(\mathbf{X_i})$, where d_{Θ} is the depth of MLP and F_x is the dimensionality of node feature embedding, and $O(Ld_{\Theta}|E|F^2 + L|E|F)$ for peer feature embedding and aggregation, where $F = F_x + F_z$ is the dimensionality of node and edge feature embeddings. For node aggregation (refer to Eq. 5), we extract ego network for each node and perform neighborhood aggregation. Therefore, the time complexity is $O(NL|\bar{E}_{max}|F)$, where $|\bar{E}_{max}|$ is the number of maximum edges in the ego network. For subsequent masking and exposure encoding MLP, the time complexity is $O(Nd_{MLP}|\bar{E}_{max}|F^2)$, where d_{MLP} is the depth considering overall MLPs.

Assuming a single-layer MPGNN with F << N < |E|, for simplicity, a standard MPGNN scales linearly with the number of edges, i.e., O(|E|) or $O(N \times avg(D))$, where avg(D) is the average degree. Similarly, for EGONETGNN the time complexity simplifies to $O(N \times |\bar{E}_{max}|)$. In the worst case, $|\bar{E}_{max}| = max(D)^2$, where max(D) is the maximum degree in the network G(V, E). However, since networks are generally sparse, the approximate runtime complexity for networks with uniform degree (e.g., Watts Strogatz network or Stochastic Block Model network) is $O(N \times P_e \times avg(D)^2)$, where P_e is density of edges. So, our method is approximately $P_e \times avg(D)$ times more computationally expensive than standard MPGNNs. On the other hand, the time complexity for counting predetermined causal network motifs with K nodes is $O(Nmax(D)^{K-1})$, assuming access to O(1) adjacency set and adjacency matrix. This approach scales poorly with higher-order motifs and EGONETGNN mitigates the problem by capturing most higher-order motifs with the same computational cost.

A.4.4 COUNTERFACTUAL OUTCOME PREDICTION ERROR BOUNDS FOR EGONETGNN

Our work utilizes Shalit et al. (2017)'s TARNet and CFR estimators, adapted to network settings, for estimating heterogeneous peer effects in both observational and experimental data. Their analysis shows the PEHE metric is bounded by factual (F), i.e., supervised learning and counterfactual (CF) prediction error, i.e., $\epsilon_{PEHE}(\hat{f}_y) \leq 2(\epsilon_{CF}(\hat{f}_y) + \epsilon_F(\hat{f}_y) - 2\sigma_y^2)$, where σ_y^2 is the variance of the outcome. These prediction errors or biases incorporate Sävje (2024)'s definition of exposure mapping specification errors along with feature representation errors and outcome prediction errors.

Moreover, Shalit et al. (2017) show that the bound for counterfactual prediction error (which cannot be measured in the real world) depends on the Integral Probability Metric (IPM) measure of distance between treatment and control group distribution, which implies $\epsilon_{PEHE}(\hat{f}_y) \leq 2(\epsilon_F^{t_i=1}(\hat{f}_y) + \epsilon_F^{t_i=0}(\hat{f}_y) + \alpha IPM(\{\boldsymbol{h}_i^{emb}:t_i=1\},\{\boldsymbol{h}_i^{emb}:t_i=0\}) - 2\sigma_Y^2)$, where $t_i=\pi_i$ denotes conditioning, $\boldsymbol{h}_i^{emb}=\Theta_{emb}(\hat{c}_i||\hat{\boldsymbol{\rho}}_i)$, and || denotes concatenation. To study how misspecification errors of EgoNetGNN propagate to the factual prediction error, we can substitute the oracle values and estimated values (denoted with hat) and further decompose the errors by using sequential error decomposition trick, i.e.,

$$\begin{split} \epsilon_F^{\mathbf{t}_i = \pi_i}(\hat{f}_y) &= \mathbb{E}[(\hat{\mathbf{y}}_i - \mathbf{y}_i)^2] \\ \hat{\mathbf{y}}_i - \mathbf{y}_i &= \hat{f}_y(\pi_i, \hat{\boldsymbol{\rho}}_i, \hat{\boldsymbol{c}}_i) - f_y(\pi_i, \boldsymbol{\rho}_i, \mathbf{c}_i) \end{split}$$

 $\hat{y_i} - y_i = \epsilon_y + \epsilon_e + \epsilon_f$, where ϵ_y captures error due to learned outcome prediction module using learned representations, i.e.,

$$\epsilon_y := \hat{f}_y(\pi_i, \hat{\boldsymbol{\rho}}_i, \hat{\mathbf{c}}_i) - f_y(\pi_i, \hat{\boldsymbol{\rho}}_i, \hat{\mathbf{c}}_i),$$

 ϵ_e captures error due to exposure mapping misspecification using learned feature representation but true outcome prediction module, i.e.,

$$\epsilon_e := f_y(\pi_i, \hat{\boldsymbol{\rho}}_i, \hat{\mathbf{c}}_i) - f_y(\pi_i, \boldsymbol{\rho}_i, \hat{\mathbf{c}}_i),$$

and, finally, ϵ_f captures error due to feature mapping misspecification but true exposure and outcome prediction function, i.e.,

$$\epsilon_f := f_y(\pi_i, \boldsymbol{\rho}_i, \hat{\mathbf{c}}_i) - f_y(\pi_i, \boldsymbol{\rho}_i, \mathbf{c}_i).$$

By plugging these decomposed errors in the factual prediction loss, we get,

$$\begin{split} \epsilon_F^{t_i = \pi_i}(\hat{f}_y) &= \mathbb{E}[(\epsilon_y + \epsilon_e + \epsilon_f)^2] \\ &= \mathbb{E}[\epsilon_y^2] + \mathbb{E}[\epsilon_e^2] + \mathbb{E}[\epsilon_f^2] + 2(\mathbb{E}[\epsilon_y \epsilon_e] + \mathbb{E}[\epsilon_e \epsilon_f] + \mathbb{E}[\epsilon_f \epsilon_y]). \end{split}$$

By automatically learning relevant exposure mapping function, we aim to directly minimize the error terms involving ϵ_e and the downstream error ϵ_y . Other estimators (e.g., Doubly robust or orthogonal learning after handling unknown exposure mapping function) can be employed in future work for more tight error bounds.

A.5 DATASET GENERATION

For the Barabasi Albert (BA) model, the preferential attachment parameter $m \in [1,5,10]$ is used to generate sparse to dense networks, where a new node connects to p_{ba} existing nodes to form the network. For the Watts Strogatz (WS) model, we set mean degree parameters $k \in \{0.002N, 0.005N, 0.01N\}$ with fixed rewiring probability of 0.5, similar to prior works (Yuan et al., 2021; Adhikari and Zheleva, 2025). For the Stochastic Block Model (SBM) model, we use the number of blocks parameters $b \in \{500, 200, 100\}$ with randomly generated edge probabilities within and across communities. We also use two real-world social networks BlogCatalog and Flickr with more realistic topology and attributes to generate treatments and outcomes. We use LDA (Blei et al., 2003) to reduce the dimensionality of raw features to 50.

Treatment model. The treatment assignments could depend on the unit's covariates as well as peer covariates and some edge attribute. We generate treatment T_i for a unit v_i as $T_i \sim \theta \left(a(\tau_c \mathbf{W}_T \times \frac{\sum_{j \in \mathcal{N}_i} \mathbf{X}^c_{j}}{\sum_{j \in \mathcal{N}_i} Z_{ij}^c} \right) + (1 - \tau_c) \mathbf{W}_T \cdot \mathbf{X}^c_{i}$, where θ denotes Bernoulli distribution, $a : \mathbb{R} \mapsto [0, 1]$ is an activation function, $\tau_c \in [0, 1]$ controls spillover influence from unit v_i 's peers, $\mathbf{X}^c \subset \mathbf{X}$ is a subset of node attributes, $Z^c \in \mathbf{Z}$ is an edge attribute, and \mathbf{W}_T is a weight matrix.

Outcome model. The outcomes depend on unit's treatment, peer treatments based on the local neighborhood condition, the confounders, and the effect modifiers. We generate outcome Y_i for a unit v_i as:

$$Y_{i} = (\delta_{exp} + \delta_{em} \times T_{i}) \times \phi_{e}(G, \mathbf{X}, \mathbf{Z}, T_{-i}) + (\tau_{d} + \tau_{em} \times \phi_{em}(G, \mathbf{X}, \mathbf{Z})) \times T_{i} + g(\mathbf{X}_{c}, Z_{c}, G) + \epsilon.$$

$$(13)$$

Here, the first term $(\delta_{exp} + \delta_{em} \times T_i) \times \phi_e(G, \mathbf{X}, \mathbf{Z}, T_{-i})$ captures peer effects, where $\phi_e(G, \mathbf{X}, \mathbf{Z}, T_{-i})$ captures true peer exposure that depends on local neighborhood condition (e.g., the number of mutual connections between treated peers and ego unit or attribute similarity) and δ_{exp} and δ_{em} are coefficients controlling magnitude/direction of peer effects. The term $g(\mathbf{X_c}, Z_c, G)$ captures confounding and $\epsilon \sim \mathcal{N}(0,1)$ is random noise. The remaining term captures direct effect due to unit's own treatment with effect modification by some contexts. For semi-synthetic data, to generate heterogeneous peer effects, we use additional effect modification due to a unit's covariates, i.e., $\delta_{em} \times T_i \times \phi_v(\mathbf{X}_{em})$, where $\mathbf{X}_{em} \subset \mathbf{X}$ and ϕ_v is a weighted mean function with randomly generated weights. Please refer to the source code in anonymous repository for detailed implementation of data generation.

A.6 ADDITIONAL EXPERIMENTAL SETTINGS

Model implementation, hyperparameters, and model selection. For the experiments, we choose $\lambda_{bal}=0.01$ for encouraging balanced representation and L1 loss regularization coefficient $\lambda_{L1}=1$ for encouraging invariance to irrelevant mechanism. We set the output embedding dimension of

exposure encoder MLP to 3 giving 6-dimensional peer exposure representation. We use 1 - layerdeep MPGNNs for feature and exposure mapping functions. Moreover, we perform grid search hyperparameter tuning by varying GNN learning rate $\{0.1, 0.04, 0.02, 0.01\}$, and setting TARNet learning rate to 0.01. We use Adam optimizer with weight decay of 10^{-5} and the learning rate is decayed by 50% after 50 epochs. A 20% held-out dataset is used for model selection, where model with lowest outcome prediction loss L_{Y_i} is chosen for reporting. We employ model checkpointing every other epoch to select the best performing model in a total of 100 epochs. Our implementation is similar to Adhikari and Zheleva (2025)'s INE-TARNet (also known as IDE-Net in original paper) in terms of MLP with residual network architecture, parameter tuning and model selection, and data generation.

The baselines INE-TARNet and GNN-TARNet-Motifs are also tuned similarly to our method by conducting grid search of the GNN's learning rate with $\{0.2, 0.02\}$ and variance smoothing regularization hyperparameter with $\{0.1, 1\}$, keeping TARNet's learning rate 0.02 and other hyperparameters default. DWR is calibrated for 5 epochs to balance representation. For other baselines, we use default hyperparameters.

Implementation of baselines. We use publicly available code shared for the baselines INE-TARNet (Adhikari and Zheleva, 2025), TNet (Chen et al., 2024), NetEstimator (Jiang and Sun, 2022), and CauGramer (Wu et al., 2025). We adapt the code provided by authors to extend it for peer effect estimation for AEMNet (Mao et al., 2025). We implement 1GNN-HSIC (Ma and Tresp, 2021) and DWR (Zhao et al., 2024) ourselves following the paper as closely as possible. GNN-TARNet-MOTIFS is available as a baseline of INE-TARNet.

Computational resources. All the experiments are performed in a machine with the following resources.

• CPU: AMD EPYC 7662 64-Core Processor (128 CPUs)

• Memory: 256 GB RAM

Operating system: Ubuntu 20.04.4 LTSGPU: NVIDIA RTX A5000 (24 GB)

• CUDA Version: 11.4

As discussed in Section A.4, the runtime of computation depends on the number of nodes and the number of edges in the ego networks along with the feature dimension. Here, we report execution time per iteration for training, evaluating, and checkpointing our model for synthetic and semi-synthetic network data. For the Barabasi Albert network with 3000 nodes, which is sparser, it takes approximately 2.1 seconds per iteration, whereas, for a Stochastic Block Model (SBM) with 3000 nodes, which is denser, it takes approximately 3.3 seconds per iteration. For the BlogCatalog network with 5196 nodes and 50-dimensional features, it takes around 5.7 seconds per iteration.

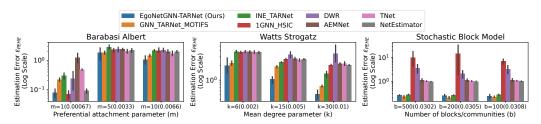
A.7 SYNTHETIC DATA EXPERIMENTS AND RESULTS

Figures 6 to 10 show the performance of our method and baselines for three synthetic networks when the underlying peer exposure mechanisms depend on clustering coefficient, connected components, number of mutual connections, tie strengths, and attribute similarity. The results discussed in the main paper apply to additional peer exposure mechanisms and data generation conditions.

A.8 Semi-synthetic Data Experiments and Results

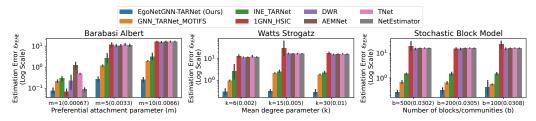
First, we present results for RQ2 for the Flickr dataset in Table 4. Either EGONETGNN-CFR or EGONETGNN-TARNet is still the best performing model in all settings. For mechanisms involving attribute similarity and clustering coefficient, EGONETGNN-TARNet is slightly better than EGONETGNN-CFR, most likely due to EGONETGNN-CFR's sensitivity to hyperparameter. INE-TARNet is the baseline with competitive performance.

Next, we utilize EGONETGNN's feature mapping MPGNN $\hat{\phi}_f$ and outcome prediction model \hat{f}_Y in the leading two baselines: GNN-TARNet-MOTIFS and INE-TARNet. The goal of this experiment is to ascertain the contribution of EGONETGNN-TARNet's exposure mapping function $\hat{\phi}_e$. Table 5



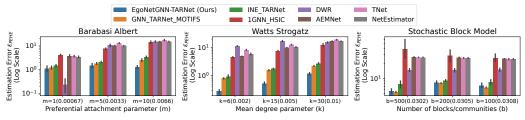
Network generating parameter and corresponding average edge densities (lower to higher)

Figure 6: Peer effect estimation error when true peer exposure depends on clustering coefficient among treated peers. Our method is better than or competitive to baseline using predetermined causal network motif counts when the underlying peer exposure mechanism can be explained by causal network motif counts.



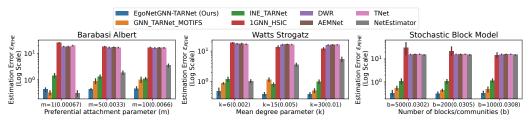
Network generating parameter and corresponding average edge densities (lower to higher)

Figure 7: Peer effect estimation error when true peer exposure depends on number of mutual connections with the ego. Our method significantly outperforms all baselines showing its capability to count closed triad network motifs (i.e., triangle substructures) in the ego network.



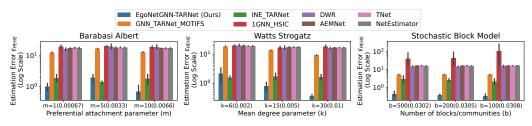
Network generating parameter and corresponding average edge densities (lower to higher)

Figure 8: Peer effect estimation error when true peer exposure depends on connected components among treated peers. Our method performs well compared to all baselines when underlying peer exposure mechanism cannot be explained totally with causal network motif structures only.



Network generating parameter and corresponding average edge densities (lower to higher)

Figure 9: Peer effect estimation error when true peer exposure depends on tie strengths between ego and treated peers. Our method consistently outperforms all baselines because it can incorporate edge attributes and learn if those attributes are relevant for underlying peer exposure mechanisms.



Network generating parameter and corresponding average edge densities (lower to higher)

Figure 10: Peer effect estimation error when true peer exposure depends on attribute similarity between ego and treated peers. Our method consistently outperforms all baselines because it can capture and learn if attribute similarity are relevant for underlying peer exposure mechanisms.

Table 4: Mean and standard deviation of peer effect estimation error (ϵ_{PEHE}) for different methods in BlogCatalog (BC) dataset for four settings when true peer exposure mechanisms depend on clustering coefficients, connected components, mutual connections, and attribute similarity. Either EGONETGNN-TARNet or EGONETGNN-CFR outperforms all other baselines across multiple settings.

Mechanisms	Ours- TARNet	Ours- CFR	GNN- Motifs	INE- TARNet	1GNN- HSIC	DWR	AEMNet	TNet	NetEst	CauGramer
	4.00	F 10	T 0.4	F 00	0.50	0.51	0.05	0.75		
Clus. Coef.	$4.93_{\pm 1.6}$	$\frac{5.12}{\pm 1.8}$	$5.34_{\pm 1.5}$	$5.26_{\pm 1.6}$	$9.56_{\pm 4.9}$	$9.51_{\pm 2.2}$	$8.05_{\pm 5.5}$	$9.75_{\pm 4.6}$	$7.57_{\pm 1.3}$	$7.84_{\pm 0.7}$
Con. Comp.	$1.83_{\pm 0.6}$	$1.40_{\pm 0.5}$	$2.80_{\pm 1.2}$	$1.85_{\pm 0.7}$	$3.36_{\pm0.8}$	$2.75_{\pm 0.6}$	$4.69_{\pm 1.7}$	$2.94_{\pm 0.9}$	$2.67_{\pm 0.5}$	$2.84_{\pm 0.6}$
Mut. Con.	$2.38_{\pm 1.3}$	$\boldsymbol{1.99}_{\pm 1.2}$	$2.55_{\pm 0.5}$	$2.36_{\pm 0.6}$	$4.03_{\pm 1.6}$	$3.57_{\pm 1.7}$	$10.95_{\pm 12.3}$	$10.96_{\pm 17.2}$	$4.24_{\pm 1.8}$	$4.34_{\pm 1.9}$
Attr. Sim.	$11.32_{\pm 6.6}$	$13.06_{\pm 12.7}$	$713.15_{\pm 10.8}$	$8\frac{12.17}{\pm 8.8}$	$16.94_{\pm 8.1}$	$18.03_{\pm 9.7}$	$17.43_{\pm 10.0}$	$23.09_{\pm 20.3}$	$16.87_{\pm 7.8}$	$20.38_{\pm 11.6}$

shows the mean and standard deviation of peer effect estimation error (ϵ_{PEHE}) for EGONETGNN and these baselines in BlogCatalog (BC) dataset for four settings when true peer exposure mechanisms depend on clustering coefficients, connected components, mutual connections, and attribute similarity. The results show our method still performs better than the baselines, verifying the contribution of the learned exposure mapping function.

Table 5: Mean and standard deviation of peer effect estimation error (ϵ_{PEHE}) for EGONETGNN and top baselines using EGONETGNN's feature mapping and outcome prediction in BlogCatalog (BC) dataset for four settings when true peer exposure mechanisms depend on clustering coefficients, connected components, mutual connections, and attribute similarity.

Method Mechanism	EgoNetGNN-TARNet	GNN-TARNet-MOTIFS	INE-TARNet
Clustering Coefficient Connected Components Mutual Connections Attribute Similarity	${\begin{aligned}&1.59_{\pm 0.4}\\&2.98_{\pm 0.8}\\&2.90_{\pm 1.1}\\&5.65_{\pm 0.7}\end{aligned}}$	$2.09_{\pm 1.2}$ $4.08_{\pm 1.0}$ $3.50_{\pm 0.7}$ $6.95_{\pm 0.9}$	$2.73_{\pm 0.6} \\ 4.52_{\pm 1.0} \\ 4.66_{\pm 2.1} \\ 5.86_{\pm 2.1}$

A.9 ABLATION STUDIES AND HYPERPARAMETER SENSITIVITY

Table 6 presents the performance of EgoNetGnn without balance loss, i.e., $\lambda_{bal}=0$, and with two different coefficients of balance loss, i.e., $\lambda_{bal}=0.01$ and $\lambda_{bal}=0.1$ for four settings when true peer exposure mechanisms depend on clustering coefficients, connected components, mutual connections, and attribute similarity. In general, using balance loss with a small coefficient results in a more robust performance. EgoNetGnn performs well for more complex peer influence mechanisms in the absence of balance loss. However, the performance for other mechanisms is comparatively poor in the absence of balance loss.

Table 7 shows the performance of EGONETGNN for different output dimension of peer exposure embedding ρ_i for four settings when true peer exposure mechanisms depend on clustering coefficients,

Table 6: Performance of EGONETGNN in BlogCatalog Data for different coefficients of balance loss for four settings when true peer exposure mechanisms depend on clustering coefficients, connected components, mutual connections, and attribute similarity.

λ_{bal} Mechanism	0.00	0.01	0.10
Clustering Coefficient Connected Components Mutual Connections Attribute Similarity	$\begin{array}{c} 1.96_{\pm 1.1} \\ \textbf{2.90}_{\pm 0.8} \\ 3.35_{\pm 0.7} \\ \textbf{5.54}_{\pm 0.6} \end{array}$	$\begin{array}{c} \underline{1.59}_{\pm 0.4} \\ \underline{2.98}_{\pm 0.8} \\ 2.90_{\pm 1.1} \\ \underline{5.65}_{\pm 0.7} \end{array}$	$\begin{array}{c} \textbf{1.33}_{\pm 0.3} \\ 3.08_{\pm 1.0} \\ \underline{2.92}_{\pm 0.7} \\ 5.77_{\pm 0.6} \end{array}$

Table 7: Performance of EGONETGNN in BlogCatalog Data for different output dimension of peer exposure embedding ρ_i for four settings when true peer exposure mechanisms depend on clustering coefficients, connected components, mutual connections, and attribute similarity.

Output Dimension Mechanism	2	6	10
Clustering Coefficient Connected Components Mutual Connections Attribute Similarity	$\begin{array}{c} \textbf{1.51}_{\pm 0.5} \\ 3.35_{\pm 0.4} \\ \underline{3.09}_{\pm 1.1} \\ 6.55_{\pm 1.8} \end{array}$	$\begin{array}{c} \underline{1.59}_{\pm 0.4} \\ 2.98_{\pm 0.8} \\ 2.90_{\pm 1.1} \\ \underline{5.65}_{\pm 0.7} \end{array}$	$\begin{array}{c} 1.95_{\pm 1.4} \\ \underline{3.03}_{\pm 0.8} \\ 3.17_{\pm 0.9} \\ 5.54_{\pm 0.7} \end{array}$

connected components, mutual connections, and attribute similarity. As seen in the results, lower-dimensional peer exposure embeddings could lose expressiveness, while higher dimensions could introduce variance due to irrelevant contexts or violations of positivity. Lower-dimensional peer exposure embedding has better performance for simpler peer exposure mechanism like clustering coefficient and higher-dimensional peer exposure embedding has better performance for complex peer exposure mechanism like attribute similarity.

Table 8 shows the performance of EgoNetGNN and top baselines in the BlogCatalog Data when the network is augmented to make it noisy by randomly removing or adding 10 and 20 percent of edges. We expect the models to perform inconsistently or worse with higher noise. The results show that for different noisy settings, our model is consistently better than the baselines. The results, however, do not show an obvious trend of higher degradation in performance with high noise. This may be because the augmentation by randomly adding or removing edges may still preserve the signal to capture underlying peer exposure mechanisms.

Table 8: Performance of EGONETGNN and top baselines in the BlogCatalog Data when the network is augmented to make it noisy by randomly removing or adding a certain percentage of edges.

Mechanism	Edge Augmentation Estimator	-20%	-10%	0%	10%	20%
Attribute Similarity	EgoNetGNN-TARNet GNN-TARNet-MOTIFS INE-TARNet	$5.51_{\pm 1.0}$ $7.19_{\pm 2.1}$ $5.97_{\pm 1.2}$	$6.03_{\pm 0.9} \ 6.74_{\pm 1.0} \ {f 5.88}_{\pm 1.6}$	$\begin{array}{c} \textbf{5.65}_{\pm 0.7} \\ 6.09_{\pm 0.2} \\ 6.01_{\pm 2.0} \end{array}$	$\begin{array}{c} \textbf{5.67}_{\pm 0.5} \\ 6.66_{\pm 0.6} \\ 5.86_{\pm 0.6} \end{array}$	$\begin{array}{c} \textbf{5.77}_{\pm 0.5} \\ 6.54_{\pm 1.1} \\ 6.50_{\pm 0.8} \end{array}$
Clustering Coefficient	EgoNetGNN-TARNet GNN-TARNet-MOTIFS INE-TARNet	$\begin{array}{c} \textbf{1.55}_{\pm 0.4} \\ 2.19_{\pm 1.1} \\ 1.79_{\pm 0.4} \end{array}$	$\begin{array}{c} \textbf{1.59}_{\pm 0.3} \\ 1.89_{\pm 0.6} \\ 1.80_{\pm 0.4} \end{array}$	$\begin{array}{c} \textbf{1.59}_{\pm 0.4} \\ 2.04_{\pm 0.7} \\ 2.24_{\pm 0.6} \end{array}$	$\begin{array}{c} \textbf{1.36}_{\pm 0.4} \\ \textbf{1.90}_{\pm 0.5} \\ \textbf{1.78}_{\pm 0.4} \end{array}$	$2.14_{\pm 1.2}$ $1.94_{\pm 0.7}$ $1.95_{\pm 0.5}$
Mutual Connections	EgoNetGNN-TARNet GNN-TARNet-MOTIFS INE-TARNet	$3.09_{\pm 0.3} \ 4.00_{\pm 1.2} \ 3.41_{\pm 0.6}$	$3.00_{\pm 0.5} \ 3.67_{\pm 0.6} \ 3.08_{\pm 0.7}$	$2.90_{\pm 1.1} \ 3.83_{\pm 0.7} \ 4.58_{\pm 2.0}$	$3.32_{\pm 0.9} \ 4.57_{\pm 2.3} \ 3.61_{\pm 1.5}$	$\begin{array}{c} \textbf{2.85}_{\pm 0.5} \\ 4.23_{\pm 2.5} \\ 3.60_{\pm 1.2} \end{array}$