

# TOWARDS ROBUST FIDELITY FOR EVALUATING EXPLAINABILITY OF GRAPH NEURAL NETWORKS

Xu Zheng<sup>1\*</sup>, Farhad Shirani<sup>1\*</sup>, Tianchun Wang<sup>2</sup>, Wei Cheng<sup>3</sup>, Zhuomin Chen<sup>1</sup>,  
Haifeng Chen<sup>3</sup>, Hua Wei<sup>4</sup>, Dongsheng Luo<sup>1</sup>✉

<sup>1</sup>School of Computing and Information Sciences, Florida International University, US

<sup>2</sup>College Information Sciences and Technology, The Pennsylvania State University, US

<sup>3</sup>NEC Laboratories America, US

<sup>4</sup>School of Computing and Augmented Intelligence, Arizona State University, US

{xzhen019, fshirani, zchen051, dluo}@fiu.edu

tkw5356@psu.edu

{weicheng, haifeng}@nec-labs.com

hua.wei@asu.edu

## ABSTRACT

Graph Neural Networks (GNNs) are neural models that leverage the dependency structure in graphical data via message passing among the graph nodes. GNNs have emerged as pivotal architectures in analyzing graph-structured data, and their expansive application in sensitive domains requires a comprehensive understanding of their decision-making processes — necessitating a framework for GNN explainability. An explanation function for GNNs takes a pre-trained GNN along with a graph as input, to produce a ‘sufficient statistic’ subgraph with respect to the graph label. A main challenge in studying GNN explainability is to provide fidelity measures that evaluate the performance of these explanation functions. This paper studies this foundational challenge, spotlighting the inherent limitations of prevailing fidelity metrics, including  $Fid_+$ ,  $Fid_-$ , and  $Fid_\Delta$ . Specifically, a formal, information-theoretic definition of explainability is introduced and it is shown that existing metrics often fail to align with this definition across various statistical scenarios. The reason is due to potential distribution shifts when subgraphs are removed in computing these fidelity measures. Subsequently, a robust class of fidelity measures are introduced, and it is shown analytically that they are resilient to distribution shift issues and are applicable in a wide range of scenarios. Extensive empirical analysis on both synthetic and real datasets are provided to illustrate that the proposed metrics are more coherent with gold standard metrics. The source code is available at <https://trustai4s-lab.github.io/fidelity>.

## 1 INTRODUCTION

Graph Neural Networks (GNNs) have become a cornerstone for processing graph-structured data, achieving impressive outcomes across various domains such as node classification and link prediction (Kipf & Welling, 2017; Hamilton et al., 2017; Veličković et al., 2018; Scarselli et al., 2008). However, with their proliferation in sensitive sectors like healthcare and fraud detection, the demand for understanding their decision-making processes has grown significantly (Zhang et al., 2022a; Wu et al., 2022; Li et al., 2022). To address this challenge, explanation techniques have been proposed for GNNs, most commonly focusing on identifying a subgraph that dominates the model’s prediction in a post-hoc sense (Ying et al., 2019; Luo et al., 2020; Yuan et al., 2021).

In the design and study of explainable GNNs, both model design and choice of evaluation metrics are important. While most efforts have primarily been made to develop new network architectures and optimization objectives to achieve more accurate explanations (Ying et al., 2019; Luo et al., 2020;

\*These authors contributed equally to this work.

✉Corresponding author

Yuan et al., 2021), in this paper, we underscore the critical importance of choosing the right evaluation metrics for the achieved explanations. In an ideal scenario, quantitative evaluation of an explanation subgraph can be achieved by comparing it with a gold standard or ground truth explanation (Ying et al., 2019). Yet, in real-world applications, such ground truth explanation subgraphs are a rarity, often making direct comparisons impracticable. In lieu of this, surrogate fidelity metrics, namely  $Fid_+$ ,  $Fid_-$ , and  $Fid$ , have been included to gauge the faithfulness of explanation subgraphs. At its core, the intuition driving such fidelity metrics is straightforward: if a subgraph is discriminative to the model, the prediction should change significantly when it is removed from the input. Otherwise, the prediction should be maintained. Hence,  $Fid_+$  is defined as the difference in accuracy (or predicted probability) between the original prediction and the new predictions of non-explanation subgraph which is obtained by masking out the explanation subgraph (Pope et al., 2019; Bajaj et al., 2021), and  $Fid_-$  measures the difference between predictions of the original graph and explanation subgraph (Yuan et al., 2022). As prevailing standards, these Fidelity metrics and their variants have been widely used in existing popular platforms, such as GraphFramEx Amara et al. (2022), GraphXAI (Agarwal et al., 2023), GNNX-BENCH (Kosan et al., 2023), and DIG (Liu et al., 2021).

Although intuitively correct, we argue that the aforementioned Fidelity metrics come with significant drawbacks due to the impractical assumption that the to-be-explained model can make accurate predictions of the explanation subgraph ( $Fid_-$ ) or non-explanation subgraph ( $Fid_+$ ). This does not hold in a wide range of real-world scenarios, because when edges are removed, the resultant subgraphs might be Out Of Distribution (OOD) (Fang et al., 2023c;b). For example, in MUTAG dataset (Debnath et al., 1991), each graph is a molecule with nodes representing atoms and edges describing the chemical bonds. The functional group  $OP_2$  is considered the dominating subgraph that causes a molecule to be mutagenic. The explanation subgraph only consists of 2 edges, which is much smaller than whole molecular graphs. Such disparities in properties introduce distribution shifts, putting the Fidelity metrics on shaky grounds, because of the violation of a key assumption in machine learning: the training and test data come from the same distribution (Hooker et al., 2019).

To build an evaluation foundation for explainable AI (XAI) in the graph domain, in this paper, we investigate robust fidelity measurements for evaluating the correctness of explanations. There are several non-trivial challenges associated with this problem. First, the to-be-explained GNN model is usually evaluated as a black-box model, which cannot be re-trained to ensure the generalization capacity (Hooker et al., 2019). Second, the evaluation method is required to be stable and ideally deterministic. As a result, complex parametric methods, such as adversarial perturbations (Hase et al., 2021; Hsieh et al., 2021), are not suitable as the results are affected by randomly initiated parameters.

We provide an information theoretic framework for GNN explainability. We define an explanation graph as a subgraph of the input that satisfies two conditions: 1) the explanation graph is an (almost) sufficient statistic of the input graph with respect to the output label, and 2) the size of the explanation graph is small compared to the input graph. Based on this, we provide two notions of explainability, namely, explainability of a classification task and explainability of a classifier for the task (Definitions 2 and 3). We quantify the relation between these two notions for low-error classifiers and classification tasks (Theorem 1). Computing the resulting information theoretic fidelity measure requires knowledge of the underlying graph statistics, which is not possible in most real-world scenarios of interest. We propose a generalized class of surrogate fidelity measures that are robust to distribution shift issues in a wide range of scenarios (Proposition A.4). Our contributions are summarized as follows.

- We pioneer in spotlighting the inherent shortcomings of widely accepted evaluation methodologies in the explainable graph learning domain.
- Grounded in solid theoretical underpinnings, we introduce novel evaluation metrics that are resilient to distribution shifts, enhancing their applicability in real-world contexts. This metric notably approximates evaluations conducted with ground truth explanations more closely.
- Through rigorous empirical analyses on a diverse mix of synthetic and real datasets, we validate that our approach resonates well with gold standard benchmarks.

## 2 PRELIMINARIES

**Random Graphs:** We parameterize a (random) labeled graph by a tuple  $(V; E; Y; X; A)$ , where i)  $V = \{v_1; v_2; \dots; v_n\}$  is the vertex set ii)  $E \subseteq V \times V$  is the edge set, iii)  $Y$  is the graph class label

<sup>1</sup>We use node and vertex interchangeably.

taking values from finite set of classes  $\mathcal{Y}$ .  $X \in \mathbb{R}^{n \times d}$  is the feature matrix, where the  $i$ -th row of  $X$ , denoted by  $x_i \in \mathbb{R}^{1 \times d}$ , is the  $d$ -dimensional feature vector associated with node  $i \in [n]$ , and  $A \in \mathbb{R}^{n \times n}$  is the adjacency matrix. The graph parameters  $(Y; A; X)$  are generated according to the joint probability measure  $P_{Y;A;X}$ . Note that the adjacency matrix determines the edge set where  $A_{ij} = 1$  if  $(v_i; v_j) \in E$ , and  $A_{ij} = 0$ , otherwise. We write  $|E|$  and  $|E_j|$  interchangeably to denote the number of edges of  $G$ . Throughout this paper, we use lower-case letters, such as  $x$ ; and  $a$ , to represent realizations of the random objects  $X; Y; X$  and  $A$ , respectively. Given a labeled graph  $G = (V; E; Y; X; A)$ , we denote the corresponding graph without labels and parameterize it by  $(V; E; X; A)$ . The induced distribution  $\bar{G}$  is represented as  $\bar{P}_{\bar{G}}$ , and its support by  $\bar{\mathcal{G}}$ .

Graph and Node Classification Tasks: In the classification tasks under consideration, we are given:

- A set of labeled training data  $\bar{\mathcal{D}} = \{(\bar{G}_i; Y_i) | Y_i \in \mathcal{Y}; i \in [T]\}$ , where  $(\bar{G}_i; Y_i)$  corresponds to the  $i$ -th graph and its associated class label. The pairs  $(\bar{G}_i; Y_i); i \in [T]$  are generated independently according to an identical joint distribution induced by  $P_{Y;X;A}$ .
- A classification function (GNN model)  $f(\cdot)$  trained to classify an unlabeled input graph  $\bar{G}$  into its class  $Y$ . It takes  $\bar{G}$  as input and outputs a probability distribution  $\bar{P}$  on alphabet  $\mathcal{Y}$ . The reconstructed label  $\hat{Y}$  is produced randomly based on  $\bar{P}$ .

In node classification tasks, each graph  $\bar{G}$  denotes a  $k$ -hop sub-graph centered around node  $v$ , with a GNN model  $f$  trained to predict the label for node  $v$  based on the node representation  $x_v$  learned from  $G$ , whereas in graph classification tasks,  $\bar{G}$  is a random graph whose distribution is determined by the (general) joint distribution  $P_{Y;A;Z}$ , with the GNN model  $f(\cdot)$  trained to predict the label for graph  $G$  based on the learned representation  $\bar{G}$ . Formally we define a classifier as follows.

Definition 1 (Classifier). For a classification task with underlying distribution  $P_{Y;X;A}$ , a classifier is a function  $f: \bar{\mathcal{G}} \rightarrow \mathcal{Y}$ . For a given  $\epsilon > 0$ , the classifier is called  $\epsilon$ -accurate if  $P(\hat{Y} \in Y) \geq 1 - \epsilon$ , where  $\hat{Y}$  is produced according to probability distribution  $\bar{P}$ .

## 2.1 INFORMATION THEORETIC MEASURES FOR QUANTIFYING EXPLAINABILITY

In this section, we introduce two different but related notions of explainability, namely, explainability of a classification task, and explainability of a classifier for a given task. The interrelations between these two notions are quantified in the subsequent sections.

Given a classification task with underlying distribution  $P_{Y;X;A}$ , an explanation function for the task is a mapping  $f: \bar{\mathcal{G}} \rightarrow (V_{exp}; E_{exp})$  which takes an unlabeled graph  $\bar{G} = (V; E; X; A)$  as input and outputs a subset of nodes  $V_{exp} \subseteq V$  and subset of edges  $E_{exp} \subseteq E$ . Loosely speaking, a good explanation  $(V_{exp}; E_{exp})$  is a subgraph which is an (almost) sufficient statistic  $\bar{G}$  with respect to the true label  $Y$ , i.e.,  $I(Y; \bar{G} | 1_{(V_{exp}; E_{exp})}) \approx 0$ , where we have defined

$$I(Y; \bar{G} | 1_{(V_{exp}; E_{exp})}) = \sum_{\bar{g}_{exp}} P_{\bar{G}}(\bar{g}_{exp}) \sum_{y \in \mathcal{Y}} P_{Y; \bar{G}}(y; \bar{g}_{exp} | \bar{G}) \log \frac{P_{Y; \bar{G}}(y; \bar{g}_{exp} | \bar{G})}{P_Y(y | \bar{g}_{exp}) P_{\bar{G}}(\bar{g}_{exp} | \bar{G})}.$$

In practice, a desirable explanation function is one whose output size is significantly smaller than the original input size, i.e.  $E_G(j(V_{exp}; E_{exp})) \ll E_G(j\bar{G})$ . This is formalized below.

Definition 2 (Explainable Classification Task). Consider a classification task with underlying distribution  $P_{Y;X;A}$ . An explanation function for this task is a mapping  $\bar{G} \rightarrow 2^V \times 2^E$ . For a given pair of parameters  $\epsilon \in [0; 1]$  and  $s \in \mathbb{N}$ , the task is called  $(\epsilon; s)$ -explainable if there exists an explanation function  $f: \bar{\mathcal{G}} \rightarrow (V_{exp}; E_{exp})$  such that:

$$i) I(Y; \bar{G} | 1_{(V_{exp}; E_{exp})}) \leq \epsilon; \quad ii) E_G(jE_{exp}) \leq s.$$

A similar notion of explainability can be provided for a given classifier as follows.

Definition 3 (Explainable Classifier). Consider a classification task with underlying distribution  $P_{Y;X;A}$  and a classifier  $f: \bar{\mathcal{G}} \rightarrow \mathcal{Y}$ . For a given pair of parameters  $\epsilon \in [0; 1]$  and  $s \in \mathbb{N}$ , the

<sup>2</sup> $I(Y; \bar{G} | 1_{(V_{exp}; E_{exp})})$  can be alternatively written as  $\sum_{\bar{g}_{exp}} P_{\bar{G}}(\bar{g}_{exp}) (I(Y; \bar{G} | \bar{g}_{exp}) - I(Y; \bar{G} | \bar{G}))$ , where  $P_{Y; \bar{G} | \bar{g}_{exp}} = P_{Y; \bar{G}} P_{\bar{G} | \bar{g}_{exp}}$ .

classifier  $f(\cdot)$  is called  $(s; \epsilon)$ -explainable if there exists an explanation function  $\phi(\cdot)$  such that

$$i) I(\phi; \bar{G} | \bar{G}) \leq \epsilon; \quad ii) E_G(jE_{\text{exp}}) \leq s;$$

where  $(\bar{G}) = (V_{\text{exp}}; E_{\text{exp}})$  and  $\phi$  is generated according to the probability distribution  $\mathbb{P}(\bar{G})$ . The explanation function  $\phi(\cdot)$  is called an  $(s; \epsilon)$  explanation for  $f(\cdot)$ .

### 3 FIDELITY MEASURES FOR EXPLAINABILITY

#### 3.1 EXPLAINABLE TASKS VS EXPLAINABLE CLASSIFIERS

It should be noted that the explainability of the classification task (Definition 2) does not imply nor is it implied by the explainability of a classifier for that task (Definition 3). For instance, the trivial classifier whose output is independent of input is explainable for any task, even if the task is not explainable itself. In this section, we characterize a quantitative relation between these two notions of explainability. To keep the analysis tractable, we introduce the following condition:

$$\text{Condition 1: } \mathbb{P}(\bar{g}; \bar{g}^0) \leq \mathbb{P}(\bar{g}^0) \quad (\bar{g}^0) = (\bar{g}): \quad (1)$$

Condition 1 holds for the ground-truth explanation in many of the widely studied datasets in the explainability literature such as BA-2 motifs, Tree-Cycles, Tree-Grid, and MUTAG datasets which are discussed in Section 5. An important consequence of Condition 1 is that it satisfies the condition, then  $I(\phi; \bar{G} | \bar{G}) = I(\phi; \bar{G} | \bar{G})$ . In the following, under the assumption of Condition 1, we show that if the classifier has a low error probability, then its explainability implies the explainability of the underlying task. Conversely, we show that if the task is explainable, and its associated Bayes error rate is small, then any classifier for the task with an error close to the Bayes error rate is explainable. The following provides the main result of this section.

**Theorem 1 (Sufficient Conditions for Equivalency of Task and Classifier Explainability).** Consider a classification task with underlying distribution  $\mathbb{P}_{Y;X;A}$ , parameters  $\epsilon; \delta \in [0; 1]$ , and an integer  $2 \leq N$ . Then,

1. If there exists a classifier  $f(\cdot)$  for this task which is  $(s; \epsilon)$ -accurate and  $(s; \delta)$ -explainable, with the explanation function satisfying Condition 1, then, the task is  $(s; \delta)$ -explainable, where

$$\delta \leq \epsilon + h_b(\epsilon) + \log(jYj - 1);$$

and  $h_b(p) = -p \log p - (1 - p) \log(1 - p)$ ;  $p \in [0; 1]$  denotes the binary entropy function. Particularly,  $\delta \geq 0$  as  $\epsilon \geq 0$ .

2. If the classification task is binary (i.e.  $|Y| = 2$ ), is  $(s; \delta)$ -explainable with an explanation function satisfying Condition 1, and has Bayes error rate equal to  $\epsilon$  then any  $(\epsilon + \delta)$ -accurate classifier  $h(\cdot)$  is  $(s; \delta)$ -explainable, where

$$\delta \leq \epsilon + \frac{2^{\epsilon} - 2^{-\epsilon}}{2}; \quad \epsilon_{\max}(\epsilon; \delta) = \frac{2^{\epsilon} - 2^{-\epsilon}}{2} + \epsilon_{\max}(\epsilon; \delta);$$

and  $\epsilon_{\max}(\epsilon)$  is defined in Proposition 1. Particularly,  $\delta \geq 0$  as  $\epsilon \geq 0$ .

The proof of Part 1) follows from the definition of explainability in Definitions 2 and 3, and Fano's inequality. The proof of Part 2) uses the following intermediate result, which proves the existence of an explainable classifier for any explainable task, whose error is close to that of the Bayes classifier.

**Proposition 1 (Existence of Accurate and Explainable Classifiers for Explainable Tasks).** Consider a classification task with underlying distribution  $\mathbb{P}_{Y;X;A}$ , parameters  $\epsilon; \delta \in [0; 1]$ , and an integer  $2 \leq N$ . Assume that the task is  $(s; \delta)$ -explainable with an explanation function satisfying Condition 1. Further assume that the classification task has a Bayes error rate equal to  $\epsilon$ . Then, there exists an  $(\epsilon + \delta)$ -accurate and  $(s; \delta)$ -explainable classifier  $h(\cdot)$ , such that

$$\epsilon_{\max}(\epsilon; \delta) \leq \epsilon + \delta; \quad (2)$$

where we have defined

$$e_{\max}(x; y) = \arg \max_z z e_{\text{RF}}(z) = h_b(x) + x \log(jYj - 1) + yg; \quad x; y \in [0; 1]$$

$$e_{\text{RF}}(z) = \frac{1}{1 - z} \ln \left( 1 + \frac{1}{1 - z} \right) \quad \text{for } z \in (0; 1);$$

In particular,  $e_{\max}(0; 1) = 0$  and  $e_{\max}(1; 0) = 0$ .

The complete proof of Proposition 1 is provided in Appendix A.1, and the proof of Theorem 1 is provided in Appendix A.2.

### 3.2 SURROGATE FIDELITY MEASURES FOR QUANTIFYING EXPLAINABILITY

Definitions 2 and 3 provide intuitive notions of explainability along with fidelity measures expressed as mutual information terms, however, in most practical applications it is not possible to quantitatively compute and analyze them. To elaborate, let us consider the mutual information term in condition i) of Definition 3. Estimating the mutual information term  $I(\psi; \bar{G} | \mathcal{G})$  is not practically feasible in most applications since  $\mathcal{G}$  has large alphabet size. To address this, prior works have considered alternative surrogate fidelity measures for evaluating the performance of explanation functions. At a high level, an ideal surrogate fidelity measure  $\text{Fid} : (f; \psi) \rightarrow \mathbb{R}^0$  must satisfy two properties: i) The surrogate fidelity value  $\text{Fid}(f; \psi)$  must be monotonic with the mutual information term  $I(\psi; \bar{G} | \mathcal{G})$  given in Definition 3, so that a 'good' explanation function with respect to the surrogate fidelity measure is 'good' under Definition 3 and vice versa, and ii) there must exist an empirical estimate of  $\text{Fid}(f; \psi)$  with sufficiently fast convergence guarantees so that the surrogate measure can be estimated accurately using a reasonably large set of observations. These conditions are formalized in the following two definitions.

**Definition 4 (Surrogate Fidelity Measure).** For a classification task with underlying distribution  $P_{Y;X;A}$ , a (surrogate) fidelity measure is a mapping  $\text{Fid} : (f; \psi) \rightarrow \mathbb{R}^0$ , which takes a pair consisting of a classification function  $f(\cdot)$  and explanation function  $\psi(\cdot)$  as input, and outputs a non-negative number. The fidelity measure is said to be well-behaved for a set of classifiers and explanation functions  $\mathcal{S}$  if for all pairs of explanation functions  $\psi_1; \psi_2 \in \mathcal{S}$  and classifiers  $f \in \mathcal{F}$ , we have:

$$I(\psi_1; \bar{G} | \mathcal{G}) \geq I(\psi_2; \bar{G} | \mathcal{G}) \implies \text{Fid}(f; \psi_2) \geq \text{Fid}(f; \psi_1); \quad (3)$$

The condition in equation 3 requires that better explanation functions in the sense of Definition 3 must have higher fidelity when evaluated using the surrogate measure.

**Definition 5 (Rate of Convergence Guarantee).** Let  $T_i = \{f(\bar{G}_i; Y_j) | j = 1, \dots, n\}$  be a sequence of sets of independent and identically distributed observations for a given classification problem. A fidelity measure  $\text{Fid}(f; \psi)$  is said to be empirically estimated with rate of convergence if there exists a sequence of functions  $\hat{\text{Fid}}_n : T_n \rightarrow \mathbb{R}^0$  such that for all  $\epsilon > 0$ , we have:

$$P(|\hat{\text{Fid}}_n(f; \psi) - \text{Fid}(f; \psi)| > \epsilon) = O(n^{-\alpha});$$

for all classifiers  $f$  and explanation functions  $\psi$ .

#### 3.2.1 THE OOD PROBLEM IN EVALUATING SURROGATE FIDELITY MEASURES

One class of surrogate fidelity measures has been considered in several recent works (e.g., (Pope et al., 2019; Yuan et al., 2022)) are the  $\text{Fid}_+$ ,  $\text{Fid}$ , and  $\text{Fid}^-$  measures:

$$\text{Fid}_+, E(\mathcal{P}(Y) | \mathcal{P}^+(Y)); \quad \text{Fid}, E(\mathcal{P}(Y) | \mathcal{P}(Y)); \quad \text{Fid}^-, \text{Fid}_+ - \text{Fid}; \quad (4)$$

where  $\mathcal{P}(\cdot)$  is the distribution given by  $\psi(\bar{G})$ ,  $\mathcal{P}^+(\cdot)$  is the distribution given by  $\psi(\bar{G}_i | (\bar{G}_i))$ ,  $\mathcal{P}^-(\cdot)$  is the distribution given by  $\psi(\cdot | (\bar{G}_i))$ ,  $\bar{G}_i | (\bar{G}_i)$  is the subgraph with edge set  $E_{\text{exp}}$  for

$(\bar{G}) = (V_{\text{exp}}; E_{\text{exp}})$ , and  $T$  is a set of independent observations  $\mathbf{y} = f(\bar{G}_i; Y_i); i = 1; 2; \dots; j; T$  jg. As mentioned in (Yuan et al., 2022), these fidelity measures can be empirically estimated by

$$\text{Fid}_{+} = \frac{1}{|T|} \sum_{i=1}^{|T|} \mathbb{P}(y_i) \mathbb{P}^{+}(y_i); \quad \text{Fid} = \frac{1}{|T|} \sum_{i=1}^{|T|} \mathbb{P}(y_i) \mathbb{P}(y_i); \quad \text{Fid}_{-} = \text{Fid}_{+} - \text{Fid} :$$

It should be noted that the rate of convergence of this empirical estimate is (e.g., using the Berry-Esseen theorem (Billingsley, 2017)). The following proposition shows that these measures are well-behaved for a class of deterministic classification tasks and completely explainable classifiers.

**Proposition 2 (Well-Behavedness of Fidelity on Deterministic Tasks).** Consider a deterministic classification task for which the induced distribution  $\mathbb{P}_{\bar{G}}$  has support  $\bar{G}$  consisting of all graphs with  $n \geq N$  vertices, the graph edges are jointly independent, and  $Z \subseteq Z^n$ , where  $Z$  is a finite set. Further assume that the graph labels  $Y = 1(\mathbf{g}_{\text{exp}} \in \bar{G})$  for a fixed subgraph  $\mathbf{g}_{\text{exp}}$ , so that the task is deterministic. Let  $(\bar{G}) = 1(\mathbf{g}_{\text{exp}} \in \bar{G})$  be the 0-correct classifier. Let  $\mathcal{S} = \{p : \mathbb{P}(y \in \{0, 1\} | \mathbf{g}) \in \mathcal{S}\}$  be a class of explanation functions, where  $\mathbb{P}(y) = \mathbb{P}_{\text{exp}}(\mathbf{g}_{\text{exp}} \in \mathcal{S}(\mathbf{g})) = p$  and  $\mathbb{P}(p(\mathbf{g}) = 1 | \mathbf{g}_{\text{exp}} \in \mathcal{S}(\mathbf{g})) = 1 - p; p \in [0, 1]$ . The Fidelity measure is well-behaved for all explanation functions  $\mathcal{S}$  in

The proof is provided in Appendix A.3. Proposition 2 shows  $\text{Fid}_{+}$  is well-behaved in a specific set of scenarios, where the task is deterministic and the classifier is completely explainable. However, we argue that it is not well-behaved in a wide range of scenarios of interest which do not have these properties. This is due to the OOD issue mentioned in the introduction. To elaborate, for a good classifier, which has low probability of error, the distribution  $\mathbb{P}(\bar{G})$  must be close to  $\mathbb{P}_{Y|\bar{G}}(j\bar{G})$  on average, i.e.  $\mathbb{E}_{\bar{G}}(d_{TV}(f(\bar{G}); \mathbb{P}_{Y|\bar{G}}(j\bar{G})))$  should be small, where  $d_{TV}$  denotes the total variation. As a result,  $\mathbb{P}(Y)$  in equation 4 is close to  $\mathbb{P}_{Y|\bar{G}}(Y | \bar{G})$  on average. However, this is not necessarily true for the  $\mathbb{P}^{+}(Y)$  and  $\mathbb{P}(Y)$  terms. The reason is that the assumption  $\mathbb{E}_{\bar{G}}(d_{TV}(f(\bar{G}); \mathbb{P}_{Y|\bar{G}}(j\bar{G}))) = 0$  only ensures that  $\bar{G}$  is close to  $\mathbb{P}_{Y|\bar{G}}(j\bar{G})$  for the typical realizations of  $\bar{G}$ . However,  $\bar{G} \neq (\bar{G})$  and  $(\bar{G})$  are not typical realizations. For instance, in many applications, it is very unlikely or impossible to observe the explanation graph in isolation, that is, to have  $(\bar{G})$ . As a result,  $\mathbb{P}^{+}(Y)$  and  $\mathbb{P}(Y)$  are not good approximations for  $\mathbb{P}_{Y|\bar{G}}(Y | \bar{G} = (\bar{G}))$  and  $\mathbb{P}_{Y|\bar{G}}(Y | (\bar{G}))$ , respectively, and  $\text{Fid}_{+}; \text{Fid}$  and  $\text{Fid}_{-}$  are not well-behaved. This is observed in our empirical analysis in Section 5, and the notion is analytically investigated in a toy-example in Appendix B.

### 3.2.2 A ROBUST CLASS OF SURROGATE FIDELITY MEASURES

Generally, in scenarios where  $(\bar{G})$  and  $\bar{G} \neq (\bar{G})$  are not typical with respect to the distribution of  $\bar{G}$ , the Fidelity measure may not be well-behaved. We address this by introducing a class of modified fidelity measures by modifying the definitions of  $\text{Fid}_{+}$  and  $\text{Fid}_{-}$  in equation 4. To this end, we define the stochastic graph sampling function  $\mathbb{E}_{\bar{G}} : \bar{G} \rightarrow \bar{G}$  with edge erasure probability  $\alpha \in [0, 1]$ . That is,  $\mathbb{E}_{\bar{G}}(\cdot)$  takes a graph  $\bar{G}$  as input, and outputs a sampled graph whose vertex set is the same as that of  $\bar{G}$ , and its edges are sampled from  $\bar{G}$  such that each edge is included with probability  $1 - \alpha$  and erased with probability  $\alpha$ , independently of all other edges. We introduce the following generalized class of surrogate fidelity measures, and show that they are robust to OOD issues in a wide range of scenarios:

$$\text{Fid}_{1; \alpha; \beta} = \mathbb{E}(\mathbb{P}(Y) \mathbb{P}^{1; \alpha; \beta}(Y)); \quad \text{Fid}_{2; \alpha; \beta} = \mathbb{E}(\mathbb{P}(Y) \mathbb{P}^{2; \alpha; \beta}(Y)); \quad (5)$$

$$\text{Fid}_{1; \alpha; \beta} = \text{Fid}_{1; \alpha} - \text{Fid}_{2; \alpha}; \quad (6)$$

where  $\alpha; \beta \in [0, 1]$ ,  $\mathbb{P}(\cdot)$  is the distribution given by  $f(\bar{G})$ ,  $\mathbb{P}^{1; \alpha; \beta}(\cdot)$  is the distribution given by  $f(\bar{G} \cap \mathbb{E}_{\alpha}(\bar{G}))$ ,  $\mathbb{P}^{2; \alpha; \beta}(\cdot)$  is the distribution given by  $f(\mathbb{E}_{\alpha}(\bar{G} \cap (\bar{G})) + (\bar{G}))$ .

Note that if  $\alpha = 1$  and  $\beta = 0$ , we recover the original fidelity measures, i.e.  $\text{Fid}_{1; 1; 0} = \text{Fid}_{+}$ ,  $\text{Fid}_{0; 0} = \text{Fid}$ , and  $\text{Fid}_{1; 0} = \text{Fid}_{-}$ . On the other hand, if  $\alpha = 0$  and  $\beta = 1$ , we have  $\bar{G} \cap \mathbb{E}_{0}(\bar{G}) = \mathbb{E}_{0}(\bar{G} \cap (\bar{G})) + (\bar{G}) = \bar{G}$ . Consequently, in this case there would be no OOD issue, however, the resulting fidelity measure is not informative since  $\text{Fid}_{0; 1}(\cdot) = 0$ , for all classifiers  $f$  and explanation functions. Loosely speaking,  $\mathbb{P}_{Y|\bar{G}}(y | \mathbf{g}); y \in \{0, 1\}$  is 'smooth' as a

function of  $\bar{g}$ , then OOD does not manifest and  $\beta_1; \beta_2 = 0$  yields a suitable fidelity measure (e.g., Proposition 2). On the other hand if  $\bar{g}$  is not smooth, then  $\beta_1 < 1$  and  $\beta_2 > 0$  would yield a suitable fidelity measure as this choice avoids the OOD problem. The generalized fidelity measures can be empirically estimated by

$$Fid_{1;+} = \frac{1}{jTj} \sum_{i=1}^j \frac{X_i}{jA_{jE_{exp}}(i)} \prod_{j \in E_{exp}; i: E_{exp}} \frac{P_{E_i}(y_i)}{P_{E_i|E}(y_i)}$$

$$Fid_{2;} = \frac{1}{jTj} \sum_{i=1}^j \frac{X_i}{jA_{jE_i}(i)} \prod_{j \in E_i; i: E_i} \frac{P_{E_i}(y_i)}{P_{E_i|E_{exp}}(y_i)}$$

$$Fid_{1;2;} = Fid_{1;+} + Fid_{2;+}$$

where  $\beta > 0, B_{j;}; \beta \in [0, 1]; \beta > 0$  denotes the interval  $[\beta, \beta + 1]$ , the set  $A_{j;}; \beta \in [0, 1]; \beta > 0$  is the set of  $\beta$ -typical binary sequences of length  $j$  with respect to the Bernoulli distribution with parameter  $\beta$ , i.e.,  $A_{j;}; \beta = \{x \in \{0, 1\}^j : \sum_{i=1}^j 1(x_i = 1) \in [j\beta, j(\beta + 1)]\}$ , the distribution  $P_{E_i}(y_i)$  is the probability of  $y_i$  under the distribution  $\bar{g}_{E_i}$ , where  $\bar{g}_{E_i}$  is the subgraph of  $\bar{G}$  with edge set restricted to  $E_i$ , and  $T = \{(\bar{G}_i; Y_i) : i \in [1, j]\}$  is the set of observations, where  $E_i$  is the edge set of  $\bar{G}_i; i \in [1, j]$ . Using the Chernoff bound and standard information theoretic arguments, it can be shown that for  $\beta_1; \beta_2$  and  $\beta = O(\frac{1}{jE_j})$ , these empirical estimates converge to their statistical counterparts with rate of convergence  $\frac{1}{2} \log \frac{1}{jTj} \rightarrow 1$  for large input graph and explanation sizes (e.g., (Csiszr & K6rner, 2011)).

We show that  $Fid_{1;2;}$  is well-behaved for a general class of tasks and classifiers, where the original Fidelity measure  $Fid_{1;0;}$  is not well-behaved. Specifically, we assume that there exists a set of motifs  $\bar{g}_y; y \in Y$ , such that

$$P(Y = y | \bar{G} = \bar{g}) \geq \begin{cases} 1 & \text{if } \bar{g}_y = \bar{g} \text{ and } \exists y^0 \in Y : \bar{g}_{y^0} \neq \bar{g}; \\ 0 & \text{if } \exists y^0 \in Y : \bar{g}_y \neq \bar{G}; \bar{g}_{y^0} = \bar{G}; \\ \frac{1}{jYj} & \text{otherwise.} \end{cases}$$

Furthermore, given  $\beta \in [0, 1]$ , we assume that the graph distribution  $P_{\bar{G}}$  and the trained classifier  $f(\cdot)$ , satisfy the following conditions. The graph has  $n$  vertices. There exists a set of input graphs, called atypical set, such that  $P_{\bar{G}}(\bar{G}) > 1 - \beta$ , and

$$P(f(\bar{g}) = f(\bar{g})) = \left( \frac{1}{d(\bar{G}; \bar{g}) + 1} \right); \tag{7}$$

where  $f(\cdot)$  is the optimal Bayes classifier, i.e.,  $f(\bar{g}) = y$  if  $\bar{g}_y = \bar{g}$ , and  $f(\bar{g}) = \arg \max_y P_Y(y)$ , otherwise<sup>3</sup>. The distance between the graph  $\bar{g}$  and the set of graphs  $\bar{G}$  is defined as  $d(\bar{G}; \bar{g}) = \min_{\bar{g}^0 \in \bar{G}} d(\bar{g}^0; \bar{g})$ , and the distance between two graphs is defined as their number of edge differences.

**Proposition 3 (Well-Behavedness of  $Fid_{1;2;}$  Fidelity Measure).** In the classification scenario described above, Consider the class of explanation functions which consists of stochastic mappings  $\bar{g} \rightarrow \bar{G}_{exp}$  where  $P(\bar{G}_{exp} = \bar{g}_y | \bar{G} = \bar{g}) = p$  for any  $\bar{g}$  such that  $\exists y : \bar{g}_y \in \bar{G}$  and  $\bar{G}_{exp} = \bar{g}_y$ , otherwise. Then,

$$Fid_{1;+} = \frac{P(A)}{1 - p + p \max_y P_Y(y)} + \frac{1}{k+1} \left( P(A^c) - 1 - P(B^0) \right)$$

$$Fid_{2;} = P(A) \left( 1 - \frac{k^2}{2^{2n^2}} \right) \frac{1}{k+1} - p$$

$$Fid_{1;2;} = P(A^c) - 1 - P(B^0) + \frac{P(A)}{1 - p + p \max_y P_Y(y)} + \left( 1 - \frac{1}{k+1} \right) + P(A) \left( 1 - \frac{k^2}{2^{2n^2}} \right) \frac{1}{k+1} - p;$$

<sup>3</sup>The assumption in equation 7 captures the infrequency of occurrence of atypical inputs in the training set, which increases chance of missclassification for those inputs and hence gives rise to OOD issues.

where  $\epsilon$  is the Bayes error rate,  $\frac{1}{2\beta^2}$ ,  $\frac{k}{2n^2}$ ,  $1 - \frac{k}{2n^2}$ ,  $k < s_1$ ,  $\min_y |g_y|$ ,  $A$  is the event that  $\exists y : \bar{g}_y \in \bar{G}$ , and  $B^0$  is the event that  $\sum_{i=1}^{n^2} X_i > k$ , where  $X_i$  are independent and identically distributed realizations of a Bernoulli variable with parameter  $\epsilon$ . Particularly, ask  $n \geq 1$  and  $\epsilon \geq 0$  such that  $k = o(n)$  and  $\epsilon = o(\frac{1}{k})$ ,  $Fid_{1; 2;}$  becomes monotonically increasing in  $\epsilon$ . Consequently, there exists a non-zero error threshold  $\epsilon_0 > 0$ , such that  $Fid_{1; 2;}$  is well-behaved for all  $\epsilon > \epsilon_0$ .

The proof is provided in Appendix A.4

## 4 RELATED WORK

**Explainable GNNs.** To interpret GNN models, previous works (Ying et al., 2019; Luo et al., 2020; Yuan et al., 2020; 2022; 2021; Lin et al., 2021; Wang & Shen, 2023; Miao et al., 2023; Fang et al., 2023c; Xie et al., 2022; Ma et al., 2022) tried different methods to offer interpretability. According to granularity, these explanation methods can generally be divided into two categories: i) instance-level explanation (Ying et al., 2019; Zhang et al., 2022b; Xie et al., 2022), which offers explanations for every instance by recognizing important substructures; and ii) model-level explanation (Yuan et al., 2020; Wang & Shen, 2023; Azzolin et al., 2023), which is designed for providing global decision rules learned by target GNN models. Within these methodologies, these methods can be classified as post-hoc explanations (Ying et al., 2019; Luo et al., 2020; Yuan et al., 2021) and self-explainable GNNs (Baldassarre & Azizpour, 2019; Dai & Wang, 2021; Miao et al., 2022), where the former uses an extra GNN model to elucidate the target GNN and the latter offers explanations while making predictions. For a comprehensive survey on this topic, please refer to (Yuan et al., 2022).

**Evaluation Metrics for Explainable GNNs.** To comprehensively evaluate the explanations, both anecdotal evidence and quantitative measures, such as explanation accuracy, faithfulness, efficiency, completeness, consistency, are jointly adopted in the literature (Nauta et al., 2022). When ground-truth explanations are unavailable, fidelity measures and their variants are routinely adopted to evaluate the quality of explanations (Yuan et al., 2021). For example, DIG (Liu et al., 2021) and GraphFramEx (Amara et al., 2022) include  $bfid_{+}$  and  $Fid_{+}$ ; GraphXAI (Agarwal et al., 2023) adopts a variant that utilizes Kullback-Leibler (KL) divergence score to quantify the distance between predictions of the original graph and subgraph. GNNX-BENCH (Kosan et al., 2023) uses Fidelity to evaluate both factual and counterfactual explanations. However, the OOD problem of subgraph explanations is overlooked in these platforms. Concurrently to this work, OAR evaluates post-hoc explanation subgraphs via analyzing adversarial robustness (Fang et al., 2023a); GInX-Eval evaluates with a removing and re-tuning strategy (Amara et al., 2023).

## 5 EXPERIMENTS

In this section, we empirically verify the effectiveness of the generalized class of surrogate fidelity measures. We also conduct extensive studies to verify our theoretical claims. Four benchmark datasets with ground truth explanations are used for evaluation, with Tree-Circles, and Tree-Grid (Ying et al., 2019) for the node classification task, and BA-2motifs (Luo et al., 2020) and MUTAG (Debnath et al., 1991) for graph classification. We consider both GCN and GIN architectures (Ying et al., 2019; Xu et al., 2019) as the models to be explained. Detailed experimental setups, full experimental results, and extra experiments are presented in the Appendix.

### 5.1 QUANTITATIVE EVALUATION BY COMPARING TO THE GOLD STANDARD

In adopted datasets, motifs are included which determine node labels or graph labels. Thus, the relationships between graph examples and data labels are well-defined by humans. The correctness of an explanation can be evaluated by comparing it to the ground truth motif. Previous studies (Ying et al., 2019; Luo et al., 2020) usually model the evaluation as an edge classification problem. Specifically, edges in the ground-truth explanation are treated as labels, and importance weights given by the explainability method are viewed as prediction scores. Then, AUC scores are considered as the metric for correctness. In this section, we use a more tractable metric, edit distance (Gao et al., 2010) to compare achieved explanations with the ground-truth motifs as a Gold Standard metric.

Consider an input graph  $G$ , and let  $G_{i;exp}$  be the ground-truth explanation subgraph, i.e., the motif. We construct a set of explanation functions, with varying qualities, to evaluate the well-behavedness



Table 1: Spearman correlation coefficient between metric and gold standard edit distance.

	Dataset	Fid <sub>+</sub> #	Fid <sub>1;+</sub> #	Fid <sub>''</sub>	Fid <sub>2;''</sub>	Fid <sub>#</sub>	Fid <sub>1;2;#</sub>	AUC
GCN	Tree-Cycles	0.229	-0.990	-0.210	0.990	0.105	-0.990	-1.000
	Tree-Grid	0.095	-1.000	0.457	1.000	-0.781	-1.000	-1.000
	BA-2motifs	-0.924	-1.000	0.819	1.000	-0.990	-1.000	-1.000
	MUTAG	-0.190	-1.000	-0.276	1.000	-0.105	-1.000	-1.000
GIN	Tree-Cycles	0.200	-1.000	-0.229	1.000	-0.286	-1.000	-1.000
	Tree-Grid	1.000	1.000	-1.000	-0.993	1.000	1.000	-1.000
	BA-2motifs	-0.838	-1.000	0.905	1.000	-1.000	-1.000	-1.000
	MUTAG	-1.000	-1.000	0.886	1.000	-0.990	-1.000	-1.000

of the proposed delity measures. To elaborate, for a given  $\alpha_2 \in [0, 1]$ , we construct an explanation function  $f_{\alpha_1; \alpha_2}(\cdot)$  by random IID sampling of the explanation subgraph edges, and the non-explanation subgraph edges, with sampling rates  $\alpha_1$  and  $\alpha_2$ , respectively. That is, to construct  $f_{\alpha_1; \alpha_2}(\bar{G}_i)$ , we remove  $\alpha_1$  ratio of edges from ground-truth explanation  $\bar{G}_{i,exp}$  via random IID sampling, and randomly add  $\alpha_2$  ratio of edges from  $\bar{G}_i \setminus \bar{G}_{i,exp}$  to  $\bar{G}_{i,exp}$  by random IID sampling from the non-explanation subgraph. Clearly, the explanation function should receive a better delity score for smaller  $(\alpha_1; \alpha_2)$ . We sweep  $\alpha_1$  and  $\alpha_2$  in the range  $\{0; 0.1; 0.3; 0.5; 0.7; 0.9\}$ , and for each combination  $(\alpha_1; \alpha_2)$ , we randomly sample 10 candidate explanations. We adopt the proposed  $Fid_{1;+}$ ,  $Fid_{2;''}$ ,  $Fid_{1;2;#}$ , where we have taken  $\alpha_1 = 1 - \alpha_2 = 0.1$ , as well as their counterparts to evaluate their qualities. For each combination, we calculate the average metric scores.

As analyzed in previous works (Yuan et al., 2021), delity measurements ignore the size of the explanation. Thus, redundant explanations are usually with high and low  $Fid_{''}$  scores. In the extreme case, with the whole input graph as the explanation, delity measures achieve the trivial optimal scores. This limitation is inherent and cannot not solved with the proposed metrics. To fairly compare the proposed metrics with the original ones, for each given a delity measurement, we use the Spearman correlation coefficient (Myers et al., 2013) between it and the gold standard edit distance to quantitatively evaluate the quality of the metric. Then, we report the average correlation scores in Table 1. The number of sampling in our measurements are set to 50.

We have the following observations in Table 1. First,  $Fid_{1;+}$  consistently yields correlation scores near -1.0 with the edit distance. It signifies a robust inverse relationship between the two metrics. This is in agreement with the requirement in equation 3. In contrast, the original metric exhibits mixed results with half-positive and half-negative correlations. This inconsistency in underscores the potential superiority and consistency of our proposed  $Fid_{1;+}$  in aligning closely with the edit distance across various datasets. Moreover, the proposed  $Fid_{2;''}$  is strongly positively related to gold-standard edit distance compared to the original  $Fid_{''}$ . we have similar observations in  $Fid_{1;2;#}$  and  $Fid_{#}$ . Third, we observe that the AUC score of edge classification, which is used in previous papers, is perfectly aligned with the gold standard edit distance, which verifies the correctness of using AUC as the metric. Last, all delity measurements fail to evaluate the GIN classifier on the Tree-Grid dataset. The potential reason is that GIN was designed for graph classification and its generalization performance on node classification is limited in Tree-Grid dataset.

## 6 CONCLUSION

In this paper, we have comprehensively explored the limitations intrinsic to prevalent evaluation methodologies in the realm of explainable graph learning, emphasizing the pitfalls of conventional delity metrics, notably their vulnerability to distribution shifts. By delving deep into the theoretical aspects, we have innovated a novel set of evaluation metrics grounded in information theory, offering resilience to such shifts and promising enhanced authenticity and applicability in real-world scenarios. These proposed metrics have been rigorously validated across varied datasets, demonstrating their superior alignment with gold standard benchmarks. Our endeavor contributes to the establishment of a more robust and reliable machine-learning evaluation framework for graph learning.

## ACKNOWLEDGMENTS

This project was partially supported by NSF grants IIS-2331908 and CCF-2241057. The views and conclusions contained in this paper are those of the authors and should not be interpreted as representing any funding agencies.

## REFERENCES

- Chirag Agarwal, Owen Queen, Himabindu Lakkaraju, and Marinka Zitnik. Evaluating explainability for graph neural networks. *Scientific Data* 10(144), 2023. URL <https://www.nature.com/articles/s41597-023-01974-x>.
- Kenza Amara, Zhitao Ying, Zitao Zhang, Zhichao Han, Yang Zhao, Yanan Shan, Ulrik Brandes, Sebastian Schemm, and Ce Zhang. Graphframex: Towards systematic evaluation of explainability methods for graph neural networks. *The First Learning on Graphs Conference*, 2022. URL <https://openreview.net/forum?id=rGVGf1T-dK>.
- Kenza Amara, Mennatallah El-Assady, and Rex Ying. GInx-eval: Towards in-distribution evaluation of graph neural network explanations. *XAI in Action: Past, Present, and Future Applications 2023*. URL <https://openreview.net/forum?id=w6Qnoy2RXG>.
- Steve Azzolin, Antonio Longa, Pietro Barbiero, Pietro Li, and Andrea Passerini. Global explainability of gnns via logic combination of learned concepts. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023.
- Mohit Bajaj, Lingyang Chu, Zi Yu Xue, Jian Pei, Lanjun Wang, Peter Cho-Ho Lam, and Yong Zhang. Robust counterfactual explanations on graph neural networks. *Advances in Neural Information Processing Systems* 34:5644–5655, 2021.
- Federico Baldassarre and Hossein Azizpour. Explainability techniques for graph convolutional networks. *arXiv preprint arXiv:1905.13686*, 2019.
- Patrick Billingsley. *Probability and measure*. John Wiley & Sons, 2017.
- Thomas M Cover. *Elements of information theory*. John Wiley & Sons, 1999.
- Imre Csizár andános Körner. *Information theory: coding theorems for discrete memoryless systems*. Cambridge University Press, 2011.
- Enyan Dai and Suhang Wang. Towards self-explainable graph neural networks. *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pages 302–311, 2021.
- Asim Kumar Debnath, Rosa L Lopez de Compadre, Gargi Debnath, Alan J Shusterman, and Corwin Hansch. Structure-activity relationship of mutagenic aromatic and heteroaromatic nitro compounds. correlation with molecular orbital energies and hydrophobicity. *Journal of medicinal chemistry*, 34(2):786–797, 1991.
- Luc Devroye, lász b Györ , and Gábor Lugosi. *A probabilistic theory of pattern recognition* volume 31. Springer Science & Business Media, 2013.
- Junfeng Fang, Wei Liu, Yuan Gao, Zemin Liu, An Zhang, Xiang Wang, and Xiangnan He. Evaluating post-hoc explanations for graph neural networks via robustness analysis. *Thirty-seventh Conference on Neural Information Processing Systems*, 2023a. URL <https://openreview.net/forum?id=eD534mPhAg>.
- Junfeng Fang, Wei Liu, An Zhang, Xiang Wang, Xiangnan He, Kun Wang, and Tat-Seng Chua. On regularization for explaining graph neural networks: An information theory perspective, 2023b. URL <https://openreview.net/forum?id=5rX7M4wa2R>.
- Junfeng Fang, Xiang Wang, An Zhang, Zemin Liu, Xiangnan He, and Tat-Seng Chua. Cooperative explanations of graph neural networks. *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*, pages 616–624, 2023c.
- Xinbo Gao, Bing Xiao, Dacheng Tao, and Xuelong Li. A survey of graph edit distance. *Pattern Analysis and applications* 13:113–129, 2010.
- Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 2017.

- Peter Hase, Harry Xie, and Mohit Bansal. The out-of-distribution problem in explainability and search methods for feature importance explanation. *Advances in neural information processing systems* 34:3650–3666, 2021.
- Sara Hooker, Dumitru Erhan, Pieter-Jan Kindermans, and Been Kim. A benchmark for interpretability methods in deep neural networks. *Advances in neural information processing systems* 32:3062–3072, 2019.
- Cheng-Yu Hsieh, Chih-Kuan Yeh, Xuanqing Liu, Pradeep Kumar Ravikumar, Seungyeon Kim, Sanjiv Kumar, and Cho-Jui Hsieh. Evaluations and methods for explanation through robustness analysis. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=4dXmpCDGNp7>.
- Neil Jethani, Adriel Saporta, and Rajesh Ranganath. Don't be fooled: label leakage in explanation methods and the importance of their quantitative evaluation. *International Conference on Artificial Intelligence and Statistics*, pp. 8925–8953. PMLR, 2023.
- Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=SJU4ayYgl>.
- Mert Kosan, Samidha Verma, Burouj Armgaan, Khushbu Pahwa, Ambuj Singh, Sourav Medya, and Sayan Ranu. Gnnx-bench: Unravelling the utility of perturbation-based gnn explainers through in-depth benchmarking. *arXiv preprint arXiv:2310.01794*, 2023.
- Vladimir A Kovalevsky. The problem of character recognition from the point of view of mathematical statistics. *Character Readers and Pattern Recognition*, pp. 3–30, 1968.
- Mark A Kramer. Nonlinear principal component analysis using autoassociative neural networks. *AIChE journal* 37(2):233–243, 1991.
- Yiqiao Li, Jianlong Zhou, Sunny Verma, and Fang Chen. A survey of explainable graph neural networks: Taxonomy and evaluation metrics. *arXiv preprint arXiv:2207.12599*, 2022.
- Wanyu Lin, Hao Lan, and Baochun Li. Generative causal explanations for graph neural networks. In *International Conference on Machine Learning*, pp. 6666–6679. PMLR, 2021.
- Meng Liu, Youzhi Luo, Limei Wang, Yaochen Xie, Hao Yuan, Shurui Gui, Haiyang Yu, Zhao Xu, Jingtun Zhang, Yi Liu, et al. Dig: A turnkey library for diving into graph deep learning research. *The Journal of Machine Learning Research* 22(1):10873–10881, 2021.
- Dongsheng Luo, Wei Cheng, Dongkuan Xu, Wenchao Yu, Bo Zong, Haifeng Chen, and Xiang Zhang. Parameterized explainer for graph neural networks. *Advances in neural information processing systems* 33:19620–19631, 2020.
- Jing Ma, Ruocheng Guo, Saumitra Mishra, Aidong Zhang, and Jundong Li. Clear: Generative counterfactual explanations on graphs. *Advances in Neural Information Processing Systems* 35:25895–25907, 2022.
- Siqi Miao, Mia Liu, and Pan Li. Interpretable and generalizable graph learning via stochastic attention mechanism. In *International Conference on Machine Learning*, pp. 15524–15543. PMLR, 2022.
- Siqi Miao, Yunan Luo, Mia Liu, and Pan Li. Interpretable geometric deep learning via learnable randomness injection. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023.
- Jerome L Myers, Arnold D Well, and Robert F Lorch. *Research design and statistical analysis*. Routledge, 2013.
- Meike Nauta, Jan Trienes, Shreyasi Pathak, Elisa Nguyen, Michelle Peters, Yasmin Schmitt, J Schöttler, Maurice van Keulen, and Christin Seifert. From anecdotal evidence to quantitative evaluation methods: A systematic review on evaluating explainability. *ACM Computing Surveys* 2022.

- Phillip E Pope, Soheil Kolouri, Mohammad Rostami, Charles E Martin, and Heiko Hoffmann. Explainability methods for graph convolutional neural networks. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10772–10781, 2019.
- Luyu Qiu, Yi Yang, Caleb Chen Cao, Jing Liu, Yueyuan Zheng, Hilary Hei Ting Ngai, Janet Hsiao, and Lei Chen. Resisting out-of-distribution data problem in perturbation of graph neural networks. *arXiv preprint arXiv:2107.14000*, 2021.
- Yuta Sakai and Ken-ichi Iwata. Sharp bounds on arimoto's conditional entropies between two distinct orders. In *2017 IEEE International Symposium on Information Theory (ISIT)*, pp. 2975–2979. IEEE, 2017.
- Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE transactions on neural networks and learning systems*, 20(1):61–80, 2008.
- D Tebbe and S Dwyer. Uncertainty and the probability of error (correspondence). *IEEE Transactions on Information theory* 14(3):516–518, 1968.
- Petar Velicković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Li, and Yoshua Bengio. Graph attention networks. *International Conference on Learning Representations* 2018. URL <https://openreview.net/forum?id=rJXMpikCZ>.
- Xiaoqi Wang and Han-Wei Shen. Gnninterpreter: A probabilistic generative model-level explanation for graph neural networks. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023.
- Bingzhe Wu, Jintang Li, Junchi Yu, Yatao Bian, Hengtong Zhang, Chaochao Chen, Chengbin Hou, Guoji Fu, Liang Chen, Tingyang Xu, et al. A survey of trustworthy graph learning: Reliability, explainability, and privacy protection. *arXiv preprint arXiv:2205.10014*, 2022.
- Yaochen Xie, Sumeet Katariya, Xianfeng Tang, Edward Huang, Nikhil Rao, Karthik Subbian, and Shuiwang Ji. Task-agnostic graph explanation. *Advances in Neural Information Processing Systems* 35:12027–12039, 2022.
- Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=ryGs6iA5Km>.
- Zhitao Ying, Dylan Bourgeois, Jiaxuan You, Marinka Zitnik, and Jure Leskovec. Gnnexplainer: Generating explanations for graph neural networks. *Advances in neural information processing systems* 32, 2019.
- Hao Yuan, Jiliang Tang, Xia Hu, and Shuiwang Ji. Xggn: Towards model-level explanations of graph neural networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 430–438, 2020.
- Hao Yuan, Haiyang Yu, Jie Wang, Kang Li, and Shuiwang Ji. On explainability of graph neural networks via subgraph explorations. *International Conference on Machine Learning*, pp. 12241–12252. PMLR, 2021.
- Hao Yuan, Haiyang Yu, Shurui Gui, and Shuiwang Ji. Explainability in graph neural networks: A taxonomic survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- He Zhang, Bang Wu, Xingliang Yuan, Shirui Pan, Hanghang Tong, and Jian Pei. Trustworthy graph neural networks: Aspects, methods and trends. *arXiv preprint arXiv:2205.07424*, 2022a.
- Shichang Zhang, Yozen Liu, Neil Shah, and Yizhou Sun. Gstarx: Explaining graph neural networks with structure-aware cooperative game theory. *Advances in Neural Information Processing Systems* 35:19810–19823, 2022b.

## A PROOFS

### A.1 PROOF OF PROPOSITION 1

Proof. Let  $(\phi)$  be the explanation corresponding to  $(s)$ -explainable task, and let  $(g)$ ,  $\arg \max_{y \in \mathcal{Y}} P_{Y|\bar{G}}(y|g)$  be the Bayes classifier. Define  $(\bar{g})$  as the deterministic classifier which outputs  $\arg \max_{y \in \mathcal{Y}} P_{Y|(\bar{G})}(y|(\bar{g}))$  for a given input  $g$ . By construction  $(\phi)$  is  $(s; 0)$ -explainable since  $(\phi; \bar{G})1_{(\bar{G})} = 0$ . It remains to show that equation 2 holds, where  $\epsilon$  is the accuracy of  $(\phi)$  and  $\epsilon^*$  is the accuracy of  $(\bar{g})$ . To this end, note that by the assumption of explainability of the task, and using the fact that under Condition 1 we have  $I(Y; \bar{G})1_{(\bar{G})} = I(Y; \bar{G}|(\bar{G}))$ , it follows that:

$$I(Y; \bar{G})1_{(\bar{G})} = I(Y; \bar{G}|(\bar{G})) = H(Y|(\bar{G})) - H(Y|\bar{G}) + \epsilon$$

By Fano's inequality, we have:

$$H(Y|\bar{G}) \leq h_b(\epsilon) + \log(|\mathcal{Y}| - 1) = H(Y|(\bar{G})) - h_b(\epsilon) + \log(|\mathcal{Y}| - 1) + \epsilon$$

On the other hand, by reverse Fano's inequality (Tebbe & Dwyer, 1968; Kovalevsky, 1968; Sakai & Iwata, 2017), we have:

$$H(Y|(\bar{G})) \leq 1 - (1 - \epsilon) \frac{1}{|\mathcal{Y}|} \leq 1 + \frac{1}{|\mathcal{Y}|} \ln \left( 1 + \frac{1}{|\mathcal{Y}|} \right)$$

$$(1 - \epsilon) \frac{1}{|\mathcal{Y}|} \leq \frac{1}{|\mathcal{Y}|} \ln \left( \frac{1}{1 - \epsilon} \right)$$

Consequently,

$$\epsilon_{RF}(\phi) \leq h_b(\epsilon) + \log(|\mathcal{Y}| - 1) + \epsilon$$

This completes the proof.  $\square$

### A.2 PROOF OF THEOREM 1

Proof. Proof of 1): From the assumption that  $(\phi)$  is  $(s; \epsilon)$ -explainable, we conclude that there exists an explanation  $(\phi)$  satisfying conditions i) and ii) in Definition 3 and Condition 1 in equation 1. Then,

$$I(Y; \bar{G})1_{(\bar{G})} = I(Y; \bar{G}|(\bar{G})) - I(Y; \phi; \bar{G}|(\bar{G}))$$

$$\stackrel{(a)}{=} I(\phi; \bar{G}|(\bar{G})) + I(Y; \bar{G}; \phi; (\bar{G}))$$

$$\stackrel{(b)}{=} I(Y; \bar{G}; \phi; (\bar{G}))$$

$$\stackrel{(c)}{=} H(Y|\phi; (\bar{G}))$$

$$\stackrel{(d)}{=} H(Y|\phi)$$

$$\stackrel{(e)}{=} h_b(\epsilon) + \log(|\mathcal{Y}| - 1);$$

where (a) follows from the chain rule of mutual information, (b) follows from the assumption that  $(\phi)$  is an  $(s; \epsilon)$  explanation for  $(\bar{g})$ , (c) follows from the information theoretic identity that  $I(X; Y) = H(X)$  for all random variables  $X, Y$ , (d) follows from the fact that conditioning reduces entropy, and (e) follows from Fano's inequality (Cover, 1999).

Proof of 2): Following the proof of Proposition 1, define  $(\bar{g})$  as the deterministic classifier which outputs  $\arg \max_{y \in \mathcal{Y}} P_{Y|(\bar{G})}(y|(\bar{g}))$  for a given input  $g$ . We have:

$$P(h(\bar{G}) \neq f(\bar{G})) = P(h(\bar{G}) \neq f(\bar{G})) + P(f(\bar{G}) \neq f(\bar{G})) \quad (8)$$

Define  $G = f \circ \bar{g} \circ 2P_{Y|\bar{G}}(1|g) - 1|g - 1$ , where  $g \in [0; 1]$ . We have,

$$P(h(\bar{G}) \neq f(\bar{G})) = P(\bar{G} \geq G) + P(\bar{G} \leq G)P(h(\bar{G}) \neq f(\bar{G})|G \geq G) \quad (9)$$

Furthermore, following (Devroye et al., 2013, Section 2.4), we have:

$$= \frac{1}{2} \frac{1}{2} E(j2P_{Y|\bar{G}}(1j\bar{g}) - 1j) + E(j2P_{Y|\bar{G}}(1j\bar{g}) - 1j) = 1 - 2^{-p} :$$

Consequently, by the Markov inequality, we have:

$$P(1 - j 2P_{Y|\bar{G}}(1j\bar{g}) - 1j \geq \frac{2}{9}) \leq \frac{2}{9} ; \quad 2 \in [0; 1]$$

So,  $P(\bar{G} \neq G) \leq \frac{2}{9}$ ;  $2 \in [0; 1]$ . Using equation 9, we get:

$$P(h(\bar{G}) \neq f(\bar{G})) \leq \frac{2}{9} + P(h(\bar{G}) \neq f(\bar{G}); \bar{G} \neq G); \quad (10)$$

On the other hand, using (Devroye et al., 2013, Theorem 2.2),

$$\begin{aligned} P(h(\bar{G}) \neq Y) &= \sum_{\bar{g}} P_{\bar{G}}(\bar{g}) P_{Y|\bar{G}}(1j\bar{g}) - \frac{1}{2} 1_{f(\bar{g}) \neq f(\bar{g})} \\ &+ 2 \sum_{\bar{g} \neq G} P_{\bar{G}}(\bar{g}) P_{Y|\bar{G}}(1j\bar{g}) - \frac{1}{2} 1_{f(\bar{g}) \neq f(\bar{g})} \\ &+ 2(1 - \frac{1}{2}) \sum_{\bar{g} \neq G} P_{\bar{G}}(\bar{g}) 1_{f(\bar{g}) \neq f(\bar{g})}; \end{aligned}$$

where in the last inequality, we have used the definition of Note that  $\sum_{\bar{g} \neq G} P_{\bar{G}}(\bar{g}) 1_{f(\bar{g}) \neq f(\bar{g})} = P(\bar{G} \neq G) P(f(\bar{G}) \neq f(\bar{G}) | \bar{G} \neq G)$ . Consequently,

$$\begin{aligned} P(h(\bar{G}) \neq Y) &\leq 2(1 - \frac{1}{2}) P(\bar{G} \neq G) P(f(\bar{G}) \neq f(\bar{G}) | \bar{G} \neq G) \\ &+ P(\bar{G} \neq G) P(f(\bar{G}) \neq f(\bar{G}) | \bar{G} \neq G) = \frac{P(h(\bar{G}) \neq Y)}{2(1 - \frac{1}{2})} = \frac{P(h(\bar{G}) \neq Y)}{2(1 - \frac{1}{2})}; \end{aligned}$$

Hence, from equation 10, we have:

$$P(h(\bar{G}) \neq f(\bar{G})) \leq \frac{2}{9} + \frac{P(h(\bar{G}) \neq Y)}{2(1 - \frac{1}{2})};$$

Similarly,

$$P(f(\bar{G}) \neq f(\bar{G})) \leq \frac{2}{9} + \frac{P(h(\bar{G}) \neq Y)}{2(1 - \frac{1}{2})};$$

From equation 8, we get:

$$P(h(\bar{G}) \neq f(\bar{G})) \leq \frac{4}{9} + \frac{P(h(\bar{G}) \neq Y)}{2(1 - \frac{1}{2})};$$

Minimizing the right-hand-side over  $2 \in [0; 1]$ , we get  $= \frac{2 \cdot 4 + \frac{4}{2} \cdot \frac{1}{2}}{9} = \frac{2 \cdot 4 + 2}{9}$  given that  $9 > e_{\max}(\frac{1}{2}; \frac{1}{2}) + \frac{1}{2}$ . Hence,

$$\begin{aligned} P(h(\bar{G}) \neq f(\bar{G})) &\leq \frac{2 \cdot 4 + 2}{9} = \frac{10}{9} \\ &= \frac{64 \cdot 2 + 16 + 2}{2 \cdot 4 \cdot 2 + 8 + 2} \\ &= \frac{(8 \cdot 2)^2}{2 \cdot 2 \cdot 2 + 8 + 2} \\ &= \frac{2 \cdot 2 + 2}{2} = 2; \end{aligned}$$

where we have defined  $\frac{1}{2} + \frac{1}{2}$ . Using the data processing inequality and Fano's inequality, we get:

$$I(\Psi; \bar{G} | 1(\bar{G})) \leq I(h(\bar{G}); \bar{G} | 1(\bar{G})) \leq h_b(\frac{1}{2}) + \log_2 Y_j;$$

This proof is completed by noting that  $e_{\max}(\frac{1}{2}; \frac{1}{2}) = \frac{1}{2}$ .

□

### A.3 PROOF OF PROPOSITION 2

Proof. We need to show that the condition in equation 3 is satisfied for any pair of explanations. We consider two cases:

Case 1:  $g_{\text{exp}}(\bar{g})$  for all  $\bar{g}$  for which  $g_{\text{exp}}(\bar{g})$ , that is, whenever the motif  $g_{\text{exp}}$  is in  $\bar{G}$ , then the explanation function  $H(\bar{G})$  outputs the motif along with potentially other irrelevant edges and vertices. In this case, it is straightforward to verify that  $\mathbb{P}(Y = 1) = 1$ ,  $\mathbb{P}^+(Y) = 1$ , and  $\mathbb{P}^-(Y) = 1$  for all  $y \in \{0, 1\}$ . So,  $\text{Fid}_+ = P(Y = 1)$ ,  $\text{Fid}_- = 0$ , and  $\text{Fid}_0 = P(Y = 1)$ . Furthermore,  $I(\bar{\psi}; \bar{G} | 1(\bar{G})) = 0$ .

Case 2: If there exists  $\bar{g}$  such that  $g_{\text{exp}}(\bar{g})$  but  $g_{\text{exp}} \notin \bar{G}$ , then for any  $\bar{g}$  such that  $g_{\text{exp}}(\bar{g})$  but  $g_{\text{exp}} \notin \bar{G}$ , we have  $\mathbb{P}(Y = 1) = 1$ ,  $\mathbb{P}^+(Y) = 1$ , and  $\mathbb{P}^-(Y) = 0$ . Otherwise, if  $g_{\text{exp}} \notin \bar{g}$  we have  $\mathbb{P}(Y = 1) = 0$  and  $\mathbb{P}^+(Y) = 1$ ,  $\mathbb{P}^-(Y) = 1$ , and  $\mathbb{P}^0(Y) = 1$ . So,  $\text{Fid}_+ = P(Y = 1) P(Y = 1; g_{\text{exp}} \notin \bar{G})$ ,  $\text{Fid}_- = 0$ , and  $\text{Fid}_0 = P(Y = 1) P(Y = 1; g_{\text{exp}} \notin \bar{G})$ . Furthermore,

$$I(\bar{\psi}; \bar{G} | 1(\bar{G})) = H(\bar{\psi} | 1(\bar{G}));$$

where we have used the fact that this is a deterministic classification task. Note that since the graph edges are pairwise independent, we have:

$$\begin{aligned} H(\bar{\psi} | 1(\bar{G})) &= P(g_{\text{exp}} \notin \bar{G}) H(\bar{\psi}) + \\ &P(g_{\text{exp}} \in \bar{G}) H(\bar{\psi} | 1(\bar{G}); g_{\text{exp}} \in \bar{G}); \end{aligned}$$

The latter is an increasing function  $H(g_{\text{exp}} \notin \bar{G})$ .

It can be observed that  $\text{Fid}_-$  is a decreasing and  $I(\bar{\psi}; \bar{G} | 1(\bar{G}))$  an increasing function of  $P(g_{\text{exp}} \notin \bar{G})$ . Consequently, the condition in equation 3 is satisfied. This completes the proof.  $\square$

### A.4 PROOF OF PROPOSITION 3

Proof. We have:

$$E(\mathbb{P}(Y)) \quad ;$$

where the left-hand-side follows from equation 7 since:

$$E(\mathbb{P}(Y)) = P(G)E(\mathbb{P}(Y) | G) = P(G)P(Y = f(\bar{G}) | G);$$

along with:

$$\begin{aligned} &P(G)P(Y = f(\bar{G}) | G) + P(G^c) \\ &= (1 - \alpha)P(Y = f(\bar{G}) | G) + \alpha P(Y = f(\bar{G}) | G) \quad \frac{1}{1 - \alpha}; \end{aligned}$$

Furthermore,

$$\begin{aligned} E(\mathbb{P}^{1:k}(Y)) &= P(G)E(\mathbb{P}^{1:k}(Y) | G) + P(G^c) \\ &\stackrel{(a)}{=} \frac{P(A)}{1} E(\mathbb{P}^{1:k}(Y) | G; A) + P(A^c) + \\ &\stackrel{(b)}{=} \frac{P(A)}{1} P(B^0) E(\mathbb{P}^{1:k}(Y) | G; A; B) + P(A^c) + 1 - P(B^0) + \\ &\stackrel{(c)}{=} \frac{P(A)}{1} E(\mathbb{P}^{1:k}(Y) | G; A; B) + P(A^c) + 1 - P(B^0) + \\ &\stackrel{(d)}{=} \frac{P(A)}{1} ((1 - p + p \max_y P_Y(y)) + 1 - \frac{1}{k+1}) + P(A^c) + 1 - P(B^0) + ; \end{aligned}$$

where  $A$  is the event that  $y \in \bar{G}$ ,  $B$  is the event that the sampling operation removes more than  $k$  edges, and  $B^0$  is the event that  $k$  consecutive realizations of a Bernoulli variable with parameter  $p$ , we get more than  $k$  ones. In (a) we have used the fact that  $P(A | G) = P(A)$ . In (b) we have used the fact that  $P(B | A; G) = P(B^0)$ . In (c), we have used  $P(B^0) = 1$ . In (d), we have used the

assumption in equation 7, and the fact that the explanation function does not output the motif with probability  $(1 - p)$ , in which case it is not subtracted in the definition of  $Fid_{1;+}$ , and it outputs the motif with probability  $p$  in which case the motif is not present  $\mathbb{P}(E_{1;+}(\bar{G}))$  since  $k < s_1$ . Additionally,

$$E(\mathbb{P}_{1;+}(Y)) = P(A)(1 - 2^{-\frac{k^2}{2n^2}})(\frac{1}{k+1})^p;$$

where we have used the fact that if occurs, and the explanation outputs the motif, which happens with probability  $p$ , then the optimal classifier outputs the correct label, and if does not occur, then  $(\cdot)$  agrees with  $(\cdot)$  with probability  $(\frac{1}{k+1})^p$ , and we have bounded  $P(B^0) = (1 - 2^{-\frac{k^2}{2n^2}})$  by Hoeffding's inequality. Consequently,

$$\begin{aligned} Fid_{1;+} &= \frac{P(A)}{1 - 2^{-\frac{k^2}{2n^2}}}((1 - p) + p \max_y P_Y(y)) + 1 - (\frac{1}{k+1})^p \\ P(A^c) &= 1 - P(B^0) \\ Fid_{2;+} &= P(A)(1 - 2^{-\frac{k^2}{2n^2}})(\frac{1}{k+1})^p \\ Fid_{1;2;+} &= P(A^c) - 1 - P(B^0) + \frac{P(A)}{1 - 2^{-\frac{k^2}{2n^2}}}((1 - p) + p \max_y P_Y(y)) + \\ &= 1 - (\frac{1}{k+1})^p + P(A)(1 - 2^{-\frac{k^2}{2n^2}})(\frac{1}{k+1})^p \end{aligned}$$

Note that the lower-bound of  $Fid_{1;2;+}$  is monotonically increasing in  $p$ . To prove that  $Fid_{1;2;+}$  is well-behaved as  $n \rightarrow \infty$  and  $k \rightarrow 0$ , it suffices to show that equation 3 holds. Note that:

$$\begin{aligned} I(Y; \bar{G} |_{(\bar{G})}) &= \sum_{g_{exp}} P_{(\bar{G})}(g_{exp}) \sum_{y; \bar{g}} P_{Y; \bar{G}}(y; \bar{g} |_{g_{exp}, \bar{G}}) \log \frac{P_{Y; \bar{G}}(y; \bar{g} |_{g_{exp}, \bar{G}})}{P_Y(y |_{g_{exp}, \bar{G}}) P_{\bar{G}}(\bar{g} |_{g_{exp}, \bar{G}})} \\ &= P_{(\bar{G})}(\cdot) I(Y; \bar{G}) + \sum_{y \in Y} P_{(\bar{G})}(g_y) I(Y; \bar{G} |_{g_y, \bar{G}}) \end{aligned} \quad (11)$$

From the proposition statement, we have:

$$P_{(\bar{G})}(\cdot) = P(\cdot |_{y: g_y, \bar{G}}) + (1 - p)P(\cdot |_{y: g_y, \bar{G}}) \quad (12)$$

$$P_{(\bar{G})}(g_y) = pP(\cdot |_{y^0: g_y^0, \bar{G}; y^0 = y}) + (1 - p)P(\cdot |_{y: g_y, \bar{G}}) \quad (13)$$

As a result, from equations 11-13, we have:

$$\begin{aligned} I(Y; \bar{G} |_{(\bar{G})}) &= P(\cdot |_{y: g_y, \bar{G}}) I(Y; \bar{G}) + \\ &= \sum_{y \in Y} P(\cdot |_{y^0: g_y^0, \bar{G}; y^0 = y})((1 - p)I(Y; \bar{G}) + pI(Y; \bar{G} |_{g_y, \bar{G}})); \end{aligned}$$

where we have used the fact that  $P(\cdot |_{y: g_y, \bar{G}}) = \sum_{y^0 \in Y} P(\cdot |_{y^0: g_y^0, \bar{G}; y^0 = y})$ . We note that  $I(Y; \bar{G}) \rightarrow H(Y)$  and  $I(Y; \bar{G} |_{g_y, \bar{G}}) \rightarrow 0$  as  $k \rightarrow 0$ . So, there exists an error threshold such that for all  $\epsilon$  and  $2 \leq Y$ , the term  $(1 - p)I(Y; \bar{G}) + pI(Y; \bar{G} |_{g_y, \bar{G}})$  is decreasing as a function of  $p$ . Since the lower-bound of  $Fid_{1;2;+}$  is monotonically increasing in  $p$ , and  $I(Y; \bar{G} |_{(\bar{G})})$  is monotonically decreasing in  $p$ , we conclude that equation 3 holds and  $Fid_{1;2;+}$  is well-behaved.  $\square$

## B EXAMPLES AND ANALYSIS

An Analytical Example with Distribution Shift for the  $Fid$  Measure: In this example, we provide an asymptotically deterministic classification scenario where for an errorless classifier there exists an explanation function which is asymptotically optimal in the sense of Definition 3, but has a worse fidelity score under  $Fid$  compared to a classifier  $\mathcal{C}_2$  which generates completely random outputs. We conclude that  $Fid$  is not well-behaved for this scenario. To elaborate, let us



consider a classification scenario where  $\mathcal{G}$  has  $2N$  vertices, and it can be decomposed as the union of two subgraphs  $\mathcal{G}_0$  and  $\mathcal{G}_{\text{exp}}$ . Further assume that  $\mathcal{G}_0$  is Erdős-Rényi with edge probability  $p \in (0, 1)$  and has  $2N$  vertices. The graph  $\mathcal{G}_{\text{exp}}$  is equal to  $C_n$  with probability  $q \in (0, 1)$ , and it is equal to an empty graph with probability  $1 - q$ , where  $C_n$  is then-cycle  $v_1 - v_2 - \dots - v_n - v_1$ . We call the graph  $\mathcal{G}$  atypical if it has less than  $\frac{n^2 p}{4}$  edges. Note that the expected number of edges in  $\mathcal{G}_0$  is  $\frac{n(n-1)}{2} p$ , and by the law of large numbers, as  $n \rightarrow \infty$ , the graph  $\mathcal{G}$  has more than  $\frac{n^2 p}{4}$  edges with probability one. Let us assume that  $Y = 1$  if  $\mathcal{G}$  is typical and  $C_n \subseteq \mathcal{G}$ , and  $Y = 0$ , otherwise. Let us consider the asymptotically optimal classifier  $\psi(\mathcal{G})$  which outputs  $\psi = 1$  if  $C_n \subseteq \mathcal{G}$  and  $\mathcal{G}$  is not atypical, and outputs  $\psi = 0$ , otherwise. Consider the explanation function

$$\psi_1(\mathcal{G}) = \begin{cases} C_n & \text{if } C_n \subseteq \mathcal{G} \\ \text{Otherwise} & \end{cases}$$

Clearly, we have  $\psi(\mathcal{G}) = \psi_1(\mathcal{G})$  as  $n \rightarrow \infty$ , and hence  $\psi_1$  is an  $(n, 0)$ -explanation function as  $n \rightarrow \infty$ . Note that  $\text{Fid}_+(\psi_1) = 1 - P(Y = 0)$ ,  $\text{Fid}_-(\psi_1) = 1 - P(Y = 0)$ , and  $\text{Fid}_0(\psi_1) = 0$  as  $n \rightarrow \infty$ . On the other hand, let  $\psi_2(\mathcal{G})$  be the explanation function which outputs a randomly and uniformly chosen subset of more than  $\frac{n^2 p}{4}$  edges from  $\mathcal{G}$ . Then  $\text{Fid}_-(\psi_2) = P(Y = 1) - P(Y = 0) = 2q - 1$  which is greater than 0 if  $q > 0.5$ . So, the trivial explanation  $\psi_2$  receives a higher score under  $\text{Fid}_-$  compared to the optimal explanation  $\psi_1$ .

## C EVALUATION ALGORITHMS

In this section, we provide the pseudo-code for computing the proposed  $\text{Fid}_{1,+}$  and  $\text{Fid}_{2,-}$  in Alg. 1 and Alg. 2, respectively. Suppose that we have a set of input graphs  $\mathcal{G}_i$ . For each graph  $\mathcal{G}_i$ , the explanation subgraph to be evaluated is denoted by  $\mathcal{E}_i$ . The model to be explained is denoted by  $f(\cdot)$ . We have two hyper-parameters  $M$  and  $\gamma_1$  in computing  $\text{Fid}_{1,+}$ .  $M$  is the number of samples and  $\gamma_1$ , introduced in equation 5, is the ratio of edges sampled from explanation subgraph. For  $\text{Fid}_{2,-}$ , we have another hyper-parameter  $\gamma_2$  instead, which indicates the ratio of edges sampled from non-explanation subgraph.

---

### Algorithm 1 Computing $\text{Fid}_{1,+}$

---

```

1: Input: A set of input graphs and their subgraphs  $(\mathcal{G}_i; (\bar{\mathcal{G}}_i)_{i=1}^T)$ , a GNN model  $f(\cdot)$ , hyperparameters  $M$  and  $\gamma_1$ .
2: Output:  $\text{Fid}_{1,+}$  of  $f(\bar{\mathcal{G}}_i)_{i=1}^T$ .
3: for each pair  $\mathcal{G}_i; (\bar{\mathcal{G}}_i)$  do
4:   for  $m$  from 1 to  $M$  do
5:      $E_{1,m}(\bar{\mathcal{G}}_i)$  sample  $\gamma_1$  edges from  $(\bar{\mathcal{G}}_i)$ 
6:      $\bar{\mathcal{G}}_{i,m} = \bar{\mathcal{G}}_i \setminus E_{1,m}(\bar{\mathcal{G}}_i)$  # Compute the non-explanation subgraph
7:      $\text{Fid}_{1,+}[i; m] = f(\bar{\mathcal{G}}_i)_{y_i} - f(\bar{\mathcal{G}}_{i,m})_{y_i}$ 
8:   end for
9:    $\text{Fid}_{1,+}[i] = \frac{1}{M} \sum_{m=1}^M \text{Fid}_{1,+}[i; m]$ 
10: end for
11:  $\text{Fid}_{1,+} = \frac{1}{T} \sum_{i=1}^T \text{Fid}_{1,+}[i]$ 
12: Return  $\text{Fid}_{1,+}$ .

```

---

## D EXTENSION BEYOND GRAPHS

Some efforts have been conducted to address the distribution shift problem in evaluating general XAI methods (Hooker et al., 2019). Most of them focus on designing more robust perturbation-based explanatory methods (Hase et al., 2021; Hsieh et al., 2021; Qiu et al., 2021; Jethani et al., 2023). For example, ROAR retrains the model with a modified dataset to ensure the generalization of the classification model (Hooker et al., 2019). Robustness measurements include adversarial perturbations into metric (Hsieh et al., 2021). However, these methods are designed for grid data and it is non-trivial to adapt them to graphs. Moreover, these methods require access to the internal parameters of classification models and are not applicable to black-box models.

---

**Algorithm 2** Computing  $Fid_{2;}$ 

---

```

1: Input: A set of input graphs and their subgraphs  $(\bar{G}_i; (\bar{G}_i)_{g_{i=1}}^T)$ , a GNN model  $f(\cdot)$ , hyperparameters  $M$  and  $2$ .
2: Output:  $Fid_{2;}$  of  $f(\bar{G}_i)_{g_{i=1}}^T$ .
3: for each pair  $(\bar{G}_i; (\bar{G}_i))$  do
4:   for  $m$  from 1 to  $M$  do
5:      $\bar{G}_i^c = \bar{G}_i - (\bar{G}_i)$  # Compute the non-explanation subgraph
6:      $E_{2;}(\bar{G}_i^c)$  sample  $2$  edges from  $\bar{G}_i^c$ 
7:      $Fid_{2;}[i; m] = f(\bar{G}_i)_{y_i} - f(E_{2;}(\bar{G}_i^c) + (\bar{G}_i))_{y_i}$ 
8:   end for
9:    $Fid_{2;}[i] = \frac{1}{M} \sum_{m=1}^M Fid_{2;}[i; m]$ 
10: end for
11:  $Fid_{2;} = \frac{1}{T} \sum_{i=1}^T Fid_{2;}[i]$ 
12: Return  $Fid_{2;}$ .

```

---

In this study, we focus on the graph domain. The notion of explainability considered in this work is specific to sub-graph explanations. Recall that a graph is a subgraph of if its vertices can be 'aligned' with a subset of vertices  $g$  such that all edges between the two sets of vertices match. The subgraph explainability is key in the definition of mutual information  $I(Y; \bar{G}_j | \bar{G})$ , which is then used to define explainability in Definitions 2 and 3. As a basic foundation, unfortunately, the definition of mutual information is non-trivial to be extended to general domains. However, the underlying ideas considered in this paper are general and may be applicable to non-graphical tasks, which is an interesting avenue for future research.

## E DETAILED EXPERIMENTAL SETUP

We use a Linux machine with 8 NVIDIA A100 GPUs as the hardware platform, each with 40GB of memory. The software environment is with CUDA 11.3, Python 3.7.16, and Pytorch 1.12.1.

### E.1 DATASETS

We adopt two node classification datasets and two graph classification datasets with ground-truth motifs. Specifically, the Tree-Cycles dataset includes an 8-depth balanced binary tree as the base graph. 80 cycle motifs are then randomly attached to nodes from the base graph. (2) Tree-Grid is created in a similar way, except that the cycle motifs are replaced with 3-by-3 grids. (3) The BA-2motifs dataset consists of 1000 graphs. Half of them are obtained by attaching a 'house' motif to a 20-node Barabási-Albert random graph. Other graphs are with 5-node cycle motifs. Graphs are divided into 2 classes based on the type of attached motifs. For these datasets, we use a 10-dimensional all 1 vector as node attributes (Ying et al., 2019). We also use a real-life dataset, MUTAG, for graph classification, which consists of 4,337 molecular graphs. These graphs are labeled based on their mutagenic effects (Ying et al., 2019). As discussed in (Luo et al., 2020), chemical groups  $NO_2$  are used as ground truth motifs.

### E.2 CLASSIFICATION MODEL ARCHITECTURES

The GCN model architectures and hyperparameters of layers are the same as the previous work (Luo et al., 2020). Specifically, for the node classification, a two-stack GCN-Relu-BatchNorm block followed by a GCN-Relu block is used for embedding. A following linear layer is used for classification. For the graph classification, a three-stack GCN-Relu-BatchNorm block is used for compute node embedding. Global max and mean pooling are used to read out the graph representations. A linear layer is then used for classification. Similarly, we create GIN models by replacing the GCN layer with a Linear-Relu-Linear-Relu GIN layer. All the variables are initialized with the default setting in Pytorch. The models are trained with the Adam optimizer with an initial learning rate of  $1:0 \cdot 10^{-3}$ . For each dataset, we follow existing works (Luo et al., 2020; Ying et al., 2019) to split train/validation/test with 8 : 1 : 1 for all datasets. Each model is trained for 1000 epochs.

## F EXTRA EXPERIMENTAL STUDIES

### F.1 EFFECTS OF $\rho_1$ AND $\rho_2$ .

As shown in our theoretical analysis,  $\rho_1$  is the rate of removing edges from explanation subgraphs in  $Fid_{1,+}$  and  $\rho_2$  is the rate of retaining edges from non-explanation subgraphs in  $Fid_{2,-}$ . To empirically verify the effects of these two parameters, we use the GCN model and vary these two hyper-parameters in the range [0.1; 0.3; 0.5; 0.7; 0.9]. Results of Spearman correlation scores are shown in Figure 1. We observe that when  $\rho_1 = 0.1$  and  $\rho_2 = 0.9$ , the proposed  $Fid_{1,+}$  and  $Fid_{2,-}$  are strongly aligned with the gold standard edit distance. As  $\rho_1$  increases, the number of removing edges from explanation subgraphs increases, leading to more a severe distribution shifting problem. Thus, the Spearman correlation coefficient between  $Fid_{1,+}$  and edit distance increases. Specifically, in the Tree-Circles dataset, the correlation even becomes positive when  $\rho_1 > 0.2$ . The similar phenomena can be observed in  $Fid_{2,-}$ . As  $\rho_2$  decreases, a smaller number of edges will be added from non-explanation subgraphs, which leads to the distribution shift problem. As a result, the  $Fid_{2,-}$  cannot be reliably aligned with the gold standard metric.

(a) Tree-Circles (b) Tree-Grid (c) BA-2motifs (d) MUTAG

Figure 1: Parameter studies on the effects of  $\rho_1$  and  $\rho_2$ .

### F.2 CLASSIFICATIONS PERFORMANCE OF GNNs

We evaluate the accuracy performance of GNN models on training, validation, and test sets. The results are shown in Table 2. Both GCN and GIN achieve good performances in these datasets, with most test accuracy scores above 0.9. Following routinely adopted settings (Ying et al., 2019; Luo et al., 2020; Yuan et al., 2021), we can safely assume that both models can correctly use the informative components (motifs) in the input graphs to make predictions.

Table 2: Accuracy performance of GNN models.

Method	Accuracy	Node Classification		Graph Classification	
		Tree-Cycles	Tree-Grid	BA-2motifs	MUTAG
GCN	Training	0.99	0.92	1.00	0.82
	Validation	1.00	0.94	1.00	0.82
	Test	0.99	0.94	0.99	0.81
GIN	Training	1.00	0.98	0.99	0.86
	Validation	1.00	0.99	1.00	0.84
	Test	0.99	0.97	1.00	0.82

### F.3 DISTRIBUTION ANALYSIS

In Section 3.2.1, we argued that the previously studied  $Fid_{1,+}$  and  $Fid_{2,-}$  surrogate measures suffer from the OOD problem, since  $Fid_{1,+}$  and  $Fid_{2,-}$  rely on accurate classifier outputs for inputs  $(\bar{G})$  and  $\bar{G}(\bar{G})$ , respectively, and  $Fid_{1,+}$  depends on both. In this section, we empirically verify the aforementioned OOD problem. Specifically, we adopt an AutoEncoder (Kramer, 1991) to embed the adjacency matrix of a (sub-)graph to a 2-dimensional hidden representation. We choose the graph classification dataset, BA-2motifs, where all graphs have the same node size. The encoder network consists of two fully-connected layers, same as that of the decoder network. Cross Entropy is used as the reconstruction error to train the Autoencoder model. The original graphs, the ground-truth explanation subgraphs, and the non-explanation subgraphs are used for training.

The visualization results of these three types of (sub-)graphs are shown in Figure 2(a). We observe that both explanation subgraphs and non-explanation subgraphs have clear distribution shifts compared to the original graphs.

(a) Explanation & non-explanation subgraphs (b) Down-sampled subgraphs in  $\text{Fid}_{1,+}$  for  $\alpha = 0:1$ . (c) Up-sampled subgraphs in  $\text{Fid}_{2,-}$  for  $\beta = 0:1$ .

Figure 2: Distribution analysis with visualization.

The proposed methods address the challenge by considering 'down/up-sampled subgraphs'. Specifically, in  $\text{Fid}_{1,+}$ , given a graph  $G$  and an explanation subgraph  $h(G)$ , as shown in Algorithm 1, we randomly remove  $\alpha$  ratio of edges in  $(\bar{G})$  from  $G$  and compare the prediction with the original one. For simplicity, we denote this as a 'down-sampled' subgraph. For  $\alpha = 0:1$ , Figure 2(b) shows the visualization results of these 'down-sampled' subgraphs. We observe that these subgraphs are in-distributed. Similarly, in  $\text{Fid}_{2,-}$ , we sample  $\beta$  ratio of edges from the non-explanation subgraph and add them to the explanation subgraph. We denote this graph as an 'up-sampled' subgraph. The visualization results with  $\beta = 0:1$  are shown in Figure 2(c), which also verifies the in-distribution property.

To quantitatively evaluate the distribution shifting problem, we first normalize the representation vector and use KL-divergence to measure the distance between up/down-sampled subgraphs and the original graphs. The KL divergence between down-sampled subgraphs and original graphs with respect to  $\alpha$  is shown in Figure 3(a). When  $\alpha = 1:0$ , the down-sampled subgraphs degenerate into non-explanation subgraphs, where the KL divergence is large. For a small value of  $\alpha$ , the KL divergence is small. In Figure 3(b), we show the KL divergence between up-sampled subgraphs and original graphs with respect to  $\beta$ . When  $\beta = 0:0$ , the up-sampled subgraphs become explanation subgraphs, which are out-of-distributed, as evidenced by a large KL divergence value. On the other hand, with a large value of  $\beta$ , for example 0.9, we are able to largely alleviate the OOD problem.

(a) Down-sampled subgraphs used in computing  $\text{Fid}_{1,+}$ . (b) Up-sampled subgraphs used in computing  $\text{Fid}_{2,-}$ .

Figure 3: KL divergence between down/up-sampled subgraphs and original graphs.

#### F.4 STABILITY ANALYSIS

In the proposed stability measurements, we sometimes need to compute  $\text{Fid}_{1,+}$  and  $\text{Fid}_{2,-}$ , respectively. To verify the stability of our measurements, we change the value of  $\alpha$  in the range

(10; 100) and keep  $\alpha_1 = 0.1$  and  $\alpha_2 = 0.9$ . For each value of  $M$ , we evaluate the quality of ground-truth explanations 10 times and report the mean as well as the standard deviation in Fig. 4. Specifically, for the default setting  $M = 50$ , we also report the comparison of running time between our methods and their counterparts in Table 3.

We have the following observations. First, as  $M$  increases, the proposed metrics are more stable. Second, our measurements are quite robust as the mean scores are stable even when  $M$  varies from 10 to 100. Thus, although our method is slower than the original fidelity measurements, in practice, a small number of samples are enough to get precise estimations.

Table 3: Running time performance on GT explanation

Model	Time(s)	Node Classification		Graph Classification	
		Tree-Circles	Tree-Grid	BA-2motifs	MUTAG
GCN	Ori.	1:06	3:77	2:98	56:91
	Ours	13:34	64:44	48:97	302:38
GIN	Ori.	0:63	2:71	2:12	50:99
	Ours	7:74	37:65	30:53	198:93

(a) Fid<sub>1;+</sub> on Tree-Circles(b) Fid<sub>2;</sub> on Tree-Circles(c) Fid<sub>1;+</sub> on Tree-Grid(d) Fid<sub>2;</sub> on Tree-Grid(e) Fid<sub>1;+</sub> on BA-2motifs(f) Fid<sub>2;</sub> on BA-2motifs(g) Fid<sub>1;+</sub> on MUTAG(h) Fid<sub>2;</sub> on MUTAGFigure 4: Parameter studies on  $M$ .

## F.5 QUALITATIVE EVALUATION

To qualitatively show the quality of different fidelity metrics, in this part, we visualize the fidelity scores with heat maps. We adopt the GCN model and follow the experimental setting in Sec.5.1. Results are shown in Fig. 5, 6, 7, and 8. We can observe several phenomena. First, with the same  $\alpha_2$ , a candidate explanation generated by a smaller edit distance and a better explanation quality. Thus, for  $Fid_+$ , the values should decrease as  $\alpha_1$  increases. However, as shown in Fig. 5(a), 6(a), 7(a), and 8(a), the original  $Fid_+$  fails to keep the monotonicity. For example, when looking at the first column of Fig. 7(a), we observe that  $Fid_+$  scores first go up and then go down with  $\alpha_1$  increasing and it achieves the highest score when there are 50% edges removed from the ground-truth explanations. The result shows  $Fid_+$  fails to measure the correctness of explanations. On the other hand, the new proposed  $Fid_{1;+}$  is highly monotonic to gold standard edit distance. We have similar observations with  $Fid$  and  $Fid_{1;2}$  on all these datasets.

(a)  $Fid_+$ (b)  $Fid$ (c)  $Fid$ (d)  $Fid_{1;+}$ (e)  $Fid_{2;}$ (f)  $Fid_{1;2;}$ 

Figure 5: Fidelity scores with GCN on Tree-Circles. The X-axis values the ratio of added edges from the non-explanation subgraph to the candidate explanation. The Y-axis values the ratio of edges removed from ground truth in the candidate explanation.

(a)  $Fid_+$ (b)  $Fid$ (c)  $Fid$ (d)  $Fid_{1;+}$ (e)  $Fid_{2;}$ (f)  $Fid_{1;2;}$ 

Figure 6: Fidelity scores with GCN on Tree-Grid.




Figure 7: Fidelity scores with GCN on BA-2motifs.




Figure 8: Fidelity scores with GCN on MUTAG.

## F.6 CASE STUDIES

In this section, we adopt case studies to show the effects of the OOD problem in original fidelity measurement and the robustness of the new proposed metrics. We use both graph classification datasets, BA-2motifs and MUTAG in this part. For each dataset, we choose a graph. We consider two candidate explanations: the gold standard one and an adversarial explanation. As shown in Table 4, the first row is the gold standard explanation and the second row is an adversarial explanation with an edit distance of 10, which doesn't contain the label information. Fidelity scores of these two explanations are close. Moreover,  $Fid$  and  $Fid$  mistakenly suggest that the adversarial explanation is better. These results further verify that the original fidelity metrics are fragile. On the other hand, the proposed fidelity metrics are more robust.  $Fid_{1;+}$ ,  $Fid_{2;}$ , and  $Fid_{1;2;}$  successfully give the correct measurements. We have a similar conclusion in the case study of MUTAG in Table 5.

Table 4: A case study on BA-2motifs dataset. The bold font indicates the better explanation subgraph suggested by fidelity measurements.

Explanation	Edit Dis.#	Fid <sub>+</sub>	Fid <sub>#</sub>	Fid <sub>''</sub>	Fid <sub>1;+</sub>	Fid <sub>2;#</sub>	Fid <sub>1;2;''</sub>
	0	-0.012	-0.036	0.024	0.169	0.002	0.167
	10	-0.010	-0.052	0.042	0.009	0.103	-0.095

Table 5: A case study on MUTAG dataset. The bold font indicates the better explanation subgraph suggested by fidelity measurements.

Explanation	Edit Dis.#	Fid <sub>+</sub>	Fid <sub>#</sub>	Fid <sub>''</sub>	Fid <sub>1;+</sub>	Fid <sub>2;#</sub>	Fid <sub>1;2;''</sub>
	0	0.619	0.627	-0.008	0.092	0.303	-0.211
	10	0.697	-0.030	0.727	0.056	0.352	-0.296

### F.7 ACCURACY BASED FIDELITY

As another variant, accuracy-based Fidelity measurement is also used to evaluate the performance of explanations (Yuan et al., 2022). Let  $\{(G_i; y_i)\}_{i=1}^T$  denote a set of  $T$  graphs and their labels.  $f(\cdot)$  is the GNN classifier to be explained. For the graph  $G_i$ ,  $(G_i)$  denotes the explanation subgraph output by an explainer  $(\cdot)$  acting on  $G_i$ . The accuracy-based fidelities are defined as follows.

$$\begin{aligned}
 \text{Fid}_+^{(\text{acc})} &= \frac{1}{T} \sum_{i=1}^T \mathbb{1}(\arg \max_x (G_i) = y_i) - \mathbb{1}(\arg \max_x (G_i - (G_i)) = y_i) \\
 \text{Fid}^{(\text{acc})} &= \frac{1}{T} \sum_{i=1}^T \mathbb{1}(\arg \max_x (G_i) = y_i) - \mathbb{1}(\arg \max_x ((G_i)) = y_i) \\
 \text{Fid}^{(\text{acc})} &= \text{Fid}_+^{(\text{acc})} - \text{Fid}^{(\text{acc})}
 \end{aligned}
 \tag{14}$$

In a similar way, we formulate our accuracy-based Fidelity measurements  $\text{Fid}_{1;+}^{(\text{acc})}$ ,  $\text{Fid}_{2;#}^{(\text{acc})}$  and  $\text{Fid}_{1;2;''}^{(\text{acc})}$ . We adopt the same setting as Sec. F.5 and show the visualization results in Fig 9,10,11, and 12. We observe consistent improvements achieved by our methods over their counterparts, indicating their effectiveness in the accuracy-based setting.






Figure 9: Accuracy-based Fidelity results of GCN on Tree-Circles.




Figure 10: Accuracy-based Fidelity results of GCN on Tree-Grid.

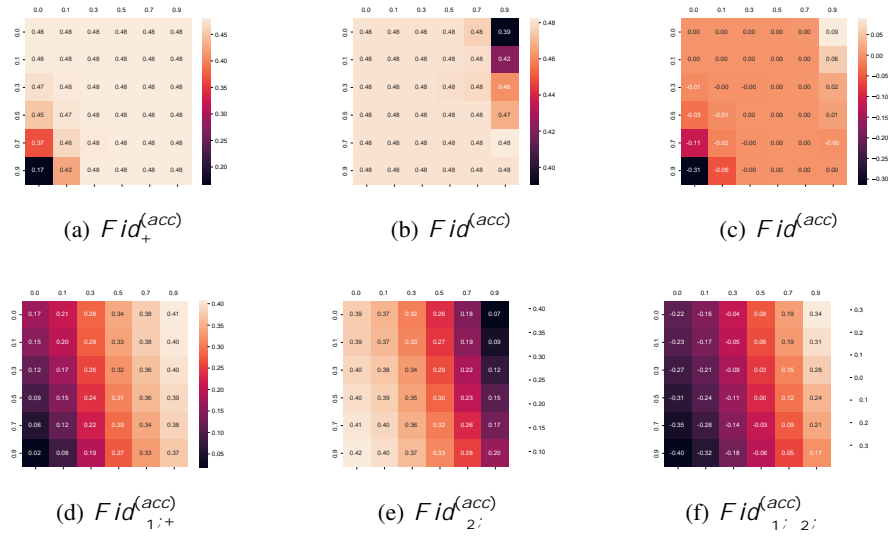


Figure 11: Accuracy-based Fidelity results of GCN on BA-2motifs.

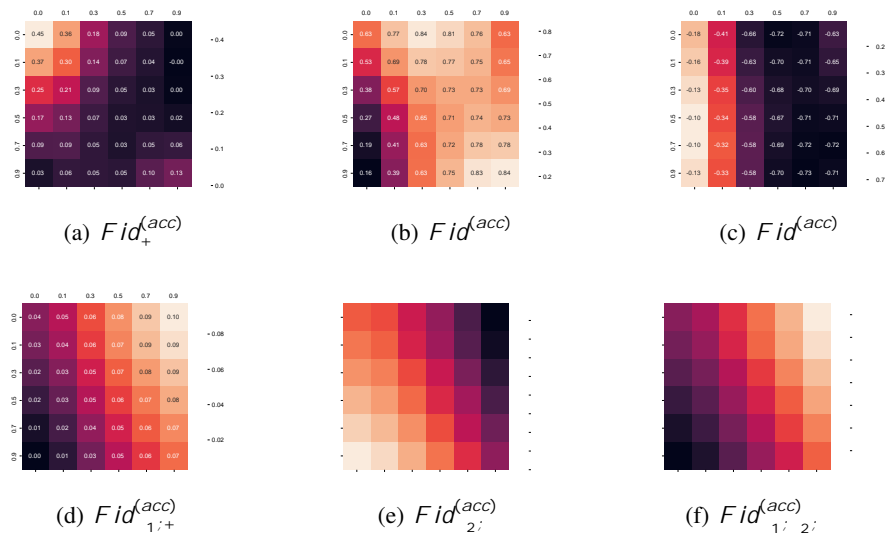


Figure 12: Accuracy-based Fidelity results of GCN on MUTAG.