
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Pereira, Tamara; Nascimento, Erik; Resck, Lucas E.; Mesquita, Diego; Souza, Amauri
Distill n' Explain

Published in:
Proceedings of The 26th International Conference on Artificial Intelligence and Statistics (AISTATS) 2023

Published: 01/01/2023

Document Version
Publisher's PDF, also known as Version of record

Published under the following license:
CC BY

Please cite the original version:
Pereira, T., Nascimento, E., Resck, L. E., Mesquita, D., & Souza, A. (2023). Distill n' Explain: explaining graph neural networks using simple surrogates. In F. Ruiz, J. Dy, & J.-W. van de Meent (Eds.), *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics (AISTATS) 2023* (pp. 6199-6214). (Proceedings of Machine Learning Research; Vol. 206). JMLR.
<https://proceedings.mlr.press/v206/pereira23a.html>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

Distill n’ Explain: explaining graph neural networks using simple surrogates

Tamara Pereira¹, Erik Nascimento¹, Lucas E. Resck², Diego Mesquita², Amauri Souza^{1,3}

¹Federal Institute of Ceará ²Getulio Vargas Foundation ³Aalto university

Abstract

Explaining node predictions in graph neural networks (GNNs) often boils down to finding graph substructures that preserve predictions. Finding these structures usually implies back-propagating through the GNN, bonding the complexity (e.g., number of layers) of the GNN to the cost of explaining it. This naturally begs the question: *Can we break this bond by explaining a simpler surrogate GNN?* To answer the question, we propose *Distill n’ Explain* (DnX). First, DnX learns a surrogate GNN via *knowledge distillation*. Then, DnX extracts node or edge-level explanations by solving a simple convex program. We also propose FastDnX, a faster version of DnX that leverages the linear decomposition of our surrogate model. Experiments show that DnX and FastDnX often outperform state-of-the-art GNN explainers while being orders of magnitude faster. Additionally, we support our empirical findings with theoretical results linking the quality of the surrogate model (i.e., distillation error) to the faithfulness of explanations.

1 Introduction

Graph neural networks (GNNs) (Gori et al., 2005; Scarselli et al., 2009) have become the pillars of representation learning on graphs. Typical GNNs resort to message passing on input graphs to extract meaningful node/graph representations for the task at hand. Despite the success of GNNs in many domains (Stokes et al., 2020; Gilmer et al., 2017; Ying et al., 2018; Sanchez-Gonzalez et al., 2020), their architectural design often results in models with limited interpretability. This naturally makes it hard to diagnose scenarios in which GNNs are fooled by confounding effects or align poorly with expert knowledge.

To mitigate this lack of interpretability, a popular strategy is to use post-hoc explanation methods (Ribeiro et al., 2016; Lundberg and Lee, 2017; Slack et al., 2021; Han et al., 2022; Huang et al., 2022). The idea is to increase model transparency by highlighting input/model elements that are particularly important for predictions, helping users to understand what is happening under the hood.

There has been a recent outbreak of methods for explaining GNNs (Yuan et al., 2022). Although GNN explanations can come in different flavors (Ying et al., 2019; Yuan et al., 2021; Wang et al., 2021; Lucic et al., 2022; Yuan et al., 2020), they usually take the form of (minimal) substructures of input graphs that are highly influential to the prediction we want to explain. The seminal work of Ying et al. (2019, GNNExplainer) proposes learning a *soft* mask to weigh graph edges. To find meaningful masks, GNNExplainer maximizes the mutual information between the GNN predictions given the original graph and the masked one. To alleviate the burden of optimizing again whenever we want to explain a different node, Luo et al. (2020, PGExplainer) propose using node embeddings to parameterize the masks, i.e., amortizing the inference. Nonetheless, GNNExplainer and PGExplainer impose strong assumptions on our access to the GNN we are trying to explain. The former assumes we are able to back-propagate through the GNN. The latter further assumes that we can access hidden activations of the GNN. Vu and Thai (2020, PGMEExplainer) relieve these assumptions by approximating the local behavior of the GNN with a probabilistic graphical model (PGM) over components, which can be used to rank the relevance of nodes and edges. On the other hand, getting explanations from PGMEExplainer involves learning the structure of a PGM, and may not scale well.

In this work, we adopt the same black-box setting of Vu and Thai (2020) but severely cut down on computational cost by extracting explanations from a *global* surrogate model. In particular, we propose *Distill n’ Explain* (DnX). DnX uses knowledge distillation to learn a simple GNN Ψ , e.g. simple graph convolution (Wu et al., 2019, SGC), that mimics the behavior of the GNN Φ we want to explain. Then, it solves a simple convex program to find a mask that weighs the influence of each node in the output of Ψ . We also propose FastDnX, a variant of DnX that leverages the

linear nature of our surrogate to speed up the explanation procedure. Notably, we only require evaluations of Φ to learn the surrogate Ψ and, after Ψ is fixed, we can use it to explain any node-level prediction. To back up the intuition that explaining a surrogate instead of the original GNN is a sensible idea, we provide a theoretical result linking the distillation quality to the faithfulness of our explanations.

Experiments on eight popular node classification benchmarks show that DnX and FastDnX often outperform GNN-, PG-, and PGM-Explainers. We also demonstrate that both DnX and FastDnX are much faster than the competitors. Remarkably, FastDnX presents a speedup of up to $65K\times$ over GNNExplainer. Finally, we discuss the limitations of current benchmarks and show that explainers capable of leveraging simple inductive biases can ace them.

Our contributions are three-fold:

1. we propose a new framework for GNN explanations that treats GNNs as black-box functions and hinges on explaining a simple surrogate model obtained through knowledge distillation;
2. we provide theoretical bounds on the quality of explanations based on these surrogates, linking the error in the distillation procedure to the faithfulness of the explanation;
3. we carry out extensive experiments, showing that our methods outperform the prior art while running orders of magnitude faster.

2 Background

Notation. We define a graph $\mathcal{G} = (V, E)$, with a set of nodes $V = \{1, \dots, n\}$ and a set of edges $E \subseteq V \times V$. We denote the adjacency matrix of \mathcal{G} by $A \in \mathbb{R}^{n \times n}$, i.e., A_{ij} is one if $(i, j) \in E$ and zero otherwise. Let D be the diagonal degree matrix of \mathcal{G} , i.e., $D_{ii} = \sum_j A_{ij}$. We also define the *normalized* adjacency matrix with added self-loops as $\tilde{A} = (D + I_n)^{-1/2}(A + I_n)(D + I_n)^{-1/2}$, where I_n is the n -dimensional identity matrix. Furthermore, let $X \in \mathbb{R}^{n \times d}$ be a matrix of d -dimensional node features. Throughout this work, we often represent a graph \mathcal{G} using the pair (A, X) .

Graph neural networks (GNNs). We consider the general framework of message-passing GNNs (Gilmer et al., 2017). Typical GNNs interleave aggregation and update steps at each layer. Specifically, for each node v at layer ℓ , the aggregation is a nonlinear function of the $(\ell - 1)$ -layer representations of v 's neighbors. The update step computes a new representation for v based on its representation at layer $\ell - 1$ and the aggregated messages (output of the aggregation step). Here we cover two specific GNN architectures: graph convolutional networks (Kipf and Welling,

2017, GCNs) and simplified graph convolutions (Wu et al., 2019, SGC). The former is arguably the most popular GNN in the literature and is used profusely throughout our experiments. The latter is a linear graph model, which will be an asset to our explanation method. For a more thorough overview of GNNs, we refer the reader to Hamilton (2020).

Graph convolutional networks combine local filtering operations (i.e., graph convolutions) and non-linear activation functions (most commonly ReLU) at each layer. Denoting the weights of the ℓ -th GCN layer by $W^{(\ell)}$ and the element-wise activation function by σ , we can recursively write the output of the ℓ -th layer $H^{(\ell)}$ as:

$$H^{(\ell)} = \sigma\left(\tilde{A}H^{(\ell-1)}W^{(\ell)}\right), \quad (1)$$

where $H^{(0)} = X$. To obtain node-level predictions, we propagate the final embeddings — after an arbitrary number of layers — through a modified convolution with a row-wise softmax instead of σ , i.e., $\hat{Y} = \text{softmax}(\tilde{A}H^{(\ell)}W^{(\ell+1)})$. In practice, it is also common to apply multilayer perceptron on top of the final embeddings.

SGC can be viewed as a simplification of the GCN model. Wu et al. (2019) derive SGC by removing the nonlinear activation functions in GCNs. Consequently, the chained linear transformations become redundant and we can use a single parameter matrix Θ . Thus, node predictions from an L -layer SGC are:

$$\hat{Y} = \text{softmax}(\tilde{A}^L X \Theta). \quad (2)$$

Interestingly, Wu et al. (2019) showed that SGC often performs similarly to or better than GCN in a variety of node classification tasks. On top of that, training SGCs is computationally more efficient than training GCNs, and SGC has significantly fewer parameters.

3 DnX: Distill n' Explain

We now introduce DnX — a new post-hoc explanation method for GNNs. DnX comprises two steps: knowledge distillation and explanation extraction. During the former, we use a linear GNN Ψ to approximate the predictions from the GNN Φ we want to explain. In the second step, we extract explanations directly from Ψ (instead of Φ). We hypothesize that, as long as Ψ is a good approximation of Φ , substructures highly influential to the output of Φ should also be relevant to Ψ . Therefore, explanations of our surrogate should also explain well the original GNN. To obtain explanations, we exploit the linear nature of Ψ and propose two simple procedures. The first consists of solving a convex program. The second ranks nodes based on a simple decomposition of predictions into additive terms.

Following Vu and Thai (2020), we assume Φ is a black-box model that we can only probe to get outputs. More

specifically, we cannot access gradients of Φ , nor can we access inner layers to extract node embeddings.

3.1 Knowledge distillation

We use SGC (Wu et al., 2019) to approximate the predictions obtained with the GNN Φ . Formally, the surrogate model (SGC) Ψ receives the input graph $\mathcal{G} = (A, X)$ and provides class predictions $\hat{Y}^{(\Psi_\Theta)} = \text{softmax}(\tilde{A}^L X \Theta)$, where Θ is the matrix of model parameters, and L is a hyper-parameter.

The distillation process consists of adjusting the parameters of Ψ_Θ to match its predictions to those of the network Φ . We do so by minimizing the Kullback-Leibler divergence KL between the predictions of Φ and Ψ_Θ . Let $\hat{Y}_i^{(\Psi_\Theta)}$ and $\hat{Y}_i^{(\Phi)}$ denote the class predictions for node i from the Ψ_Θ and Φ models, respectively. We distill Φ into Ψ by solving:

$$\min_{\Theta} \left\{ \text{KL} \left(\hat{Y}^{(\Phi)}, \hat{Y}^{(\Psi_\Theta)} \right) := \sum_{i \in V} \sum_c \hat{Y}_{ic}^{(\Phi)} \log \frac{\hat{Y}_{ic}^{(\Phi)}}{\hat{Y}_{ic}^{(\Psi_\Theta)}} \right\}, \quad (3)$$

which is equivalent to the categorical cross-entropy between $\hat{Y}^{(\Phi)}$ and $\hat{Y}^{(\Psi_\Theta)}$. Note that minimizing this loss does not require back-propagating through the original GNN Φ , only through the surrogate Ψ . We also do not require any knowledge about Φ 's architecture.

3.2 Explanation extraction

To obtain an explanation to a given prediction $\hat{Y}_i^{(\Psi_\Theta)}$, we want to identify a subgraph of \mathcal{G} containing the nodes that influence the most that prediction. We denote an explanation \mathcal{E} as an n -dimensional vector of importance scores (higher equals more relevant), one for each node in the vertex set V . We introduce two strategies to compute \mathcal{E} .

Optimizing for \mathcal{E} . We can formulate the problem of finding the explanation \mathcal{E} by treating it as a vector of 0-1 weights, and minimizing the squared L_2 norm between the logits associated with $\hat{Y}_i^{(\Psi_\Theta)}$ and those from the graph with node features masked by \mathcal{E} :

$$\min_{\mathcal{E} \in \{0,1\}^n} \left\| \tilde{A}_i^L \text{diag}(\mathcal{E}) X \Theta - \tilde{A}_i^L X \Theta \right\|_2^2, \quad (4)$$

where \tilde{A}_i^L denotes the i -th row of the matrix \tilde{A}^L . Note that the formulation in Equation 4 has a major issue: it admits the trivial solution $\mathcal{E} = [1, 1, \dots, 1]$. To circumvent the issue and simultaneously avoid binary optimization, we replace the search space $\{0, 1\}^n$ by the $(n - 1)$ -simplex $\Delta = \{r \in \mathbb{R}^n : \sum_i r_i = 1, \forall_i r_i \geq 0\}$. Implementing this change and re-arranging computations, we wind up with:

$$\min_{\mathcal{E} \in \Delta} \left\| \tilde{A}_i^L (\text{diag}(\mathcal{E}) - I_n) X \Theta \right\|_2^2. \quad (5)$$

Note that nodes outside the L -hop neighborhood of node i do not affect how Ψ classifies it. Thus, we can mask all nodes at distance $\geq L + 1$ without altering the solution of Equation 5. For ease of implementation, we solve Equation 5 reparameterizing \mathcal{E} as a softmax-transformed vector.

Finding \mathcal{E} via linear decomposition. Let Z_i denote the logit vector associated with the prediction $\hat{Y}_i^{(\Psi_\Theta)}$. Due to the linear nature of Ψ , we can decompose Z_i into a sum of n terms, one for each node in V (plus the bias):

$$\tilde{A}_{i1}^L X_1 \Theta + \tilde{A}_{i2}^L X_2 \Theta + \dots + \tilde{A}_{in}^L X_n \Theta + b = Z_i. \quad (6)$$

Therefore, we can measure the contribution of each node to the prediction as its scalar projection onto $Z_i - b$:

$$\mathcal{E}_j := \tilde{A}_{ij}^L X_j \Theta (Z_i - b)^\top \quad (7)$$

When we use this strategy instead of solving Equation 5, we refer to our method as FastDnX.

4 Analysis

In this section, we discuss the theoretical and computational aspects of our method. We first provide theoretical results supporting the hypothesis that good explanations of a global surrogate Ψ also characterize good explanations of Φ — in terms of faithfulness. Then, we discuss the convexity of the optimization problem DnX solves to extract explanations. We delegate proofs to the Appendix A.

Let \mathcal{G}_u denote the subgraph of \mathcal{G} induced by the L -hop neighborhood around node u . We say an explanation \mathcal{E}_u for a node u is faithful with respect to Φ if: i) Φ outputs approximately the same predictions for u regardless of using \mathcal{E}_u to weigh the nodes of \mathcal{G}_u or not; and ii) the same holds under small perturbations of \mathcal{G}_u . We can define a perturbation \mathcal{G}'_u of \mathcal{G}_u by adding noise to u 's features or by randomly rewiring node u 's incident edges (Agarwal et al., 2022). In this work, we consider perturbations over node features. More precisely, this entails that $V(\mathcal{G}'_u) = V(\mathcal{G}_u)$, $E(\mathcal{G}'_u) = E(\mathcal{G}_u)$, and that features are corrupted by noise, i.e., $X'_i = X_i + \epsilon_i$ for $i \in V(\mathcal{G}_u)$ and $\epsilon_i \in \mathbb{R}^d$.

Definition 1 (Faithfulness). Given a set \mathcal{K} of perturbations of \mathcal{G}_u , an explanation \mathcal{E}_u is faithful to a model f if

$$\frac{1}{|\mathcal{K}| + 1} \sum_{\mathcal{G}'_u \in \mathcal{K} \cup \{\mathcal{G}_u\}} \|f(\mathcal{G}'_u) - f(t(\mathcal{G}'_u, \mathcal{E}_u))\|_2 \leq \delta,$$

where \mathcal{G}'_u is a possibly perturbed version of \mathcal{G}_u , t is a function that applies the explanation \mathcal{E}_u to the graph \mathcal{G}'_u , and δ is a small constant (Agarwal et al., 2022).

Lemma 1 provides an upper bound on the unfaithfulness of \mathcal{E}_u with respect to the surrogate model Ψ . Theorem 1 extends this result to obtain a bound for \mathcal{E}_u with respect to the model we originally want to explain, i.e., Φ .

Lemma 1 (Unfaithfulness with respect to Ψ). *Given a node u and a set \mathcal{K} of perturbations, the unfaithfulness of the explanation \mathcal{E}_u with respect to the prediction $Y_u^{(\Psi_\Theta)}$ of node u is bounded as follows:*

$$\frac{1}{|\mathcal{K}|+1} \sum_{\substack{\mathcal{G}'_u \in \\ \mathcal{K} \cup \{\mathcal{G}_u\}}} \|\Psi(\mathcal{G}'_u) - \Psi(t(\mathcal{G}'_u, \mathcal{E}_u))\|_2 \leq \gamma \left\| \frac{\Delta \tilde{A}_u^L}{\mathcal{E}_u} \right\|_2,$$

where \mathcal{G}'_u is a possibly perturbed version of \mathcal{G}_u , t is a function that applies the explanation \mathcal{E}_u to the graph \mathcal{G}'_u , γ is a constant that depends on the model weights Θ , node features X , and perturbation ϵ . Furthermore, $\frac{\Delta \tilde{A}_u^L}{\mathcal{E}_u}$ is the u -th row of the difference of the powered, normalized adjacency matrix \tilde{A}^L before and after applying the explanation \mathcal{E}_u .

Sketch of the proof. We first show that

$$\|\Psi(\mathcal{G}_u) - \Psi(t(\mathcal{G}_u, \mathcal{E}_u))\|_2 \leq \|(X\Theta)^\top\|_2 \left\| \tilde{A}_u^L - \tilde{E}_u^L \right\|_2$$

by using Lipschitz continuity of the softmax function and the compatibility property of the L_2 matrix norm. We repeat for $\mathcal{G}'_u \in \mathcal{K}$, take the mean in $\mathcal{K} \cup \{\mathcal{G}_u\}$ and isolate $\left\| \frac{\Delta \tilde{A}_u^L}{\mathcal{E}_u} \right\|_2 = \left\| \tilde{A}_u^L - \tilde{E}_u^L \right\|_2$. The complete proof is available in Appendix A. \square

Theorem 1 (Unfaithfulness with respect to Φ). *Under the same assumptions of Lemma 1 and assuming the L_2 distillation error is bounded by α , the unfaithfulness of the explanation \mathcal{E}_u for the original model Φ 's node u prediction is bounded as follows:*

$$\frac{1}{|\mathcal{K}|+1} \sum_{\substack{\mathcal{G}'_u \in \\ \mathcal{K} \cup \{\mathcal{G}_u\}}} \|\Phi(\mathcal{G}'_u) - \Phi(t(\mathcal{G}'_u, \mathcal{E}_u))\|_2 \leq \gamma \left\| \frac{\Delta \tilde{A}_u^L}{\mathcal{E}_u} \right\|_2 + 2\alpha.$$

Note that Theorem 1 establishes a bound on faithfulness that depends directly on the distillation error α . Importantly, when Ψ is a perfect approximation of Φ , we retrieve upper-bound on the RHS of Lemma 1.

We note that Theorem 1 by Agarwal et al. (2022) covers an upper bound for the unfaithfulness of GNN explanation methods. However, they do not cover the case in which the explanation is a (weighted) subset of nodes in the L -hop neighborhood of u , as in our method.

For completeness, we also extend Lemma 1 and Theorem 1 to account for the (very often) probabilistic nature of the noise, i.e., for the case in which ϵ_i are random variables.

Lemma 2 (Probability bound on unfaithfulness w.r.t. Ψ). *Given a node u and a set \mathcal{K} of perturbations and assuming the perturbations are i.i.d. with distribution $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$, the unfaithfulness of the explanation \mathcal{E}_u*

with respect to the prediction $Y_u^{(\Psi_\Theta)}$ of node u is bounded in probability as follows:

$$\mathbb{P} \left(\frac{1}{|\mathcal{K}|+1} \sum_{\substack{\mathcal{G}'_u \in \\ \mathcal{K} \cup \{\mathcal{G}_u\}}} \|\Psi(\mathcal{G}'_u) - \Psi(t(\mathcal{G}'_u, \mathcal{E}_u))\|_2 \leq \xi \right) \geq \geq F_{\chi^2_{|\mathcal{K}|nd}} \left(\frac{\xi - \gamma_1 \left\| \frac{\Delta \tilde{A}_u^L}{\mathcal{E}_u} \right\|_2}{\gamma_2 \left\| \frac{\Delta \tilde{A}_u^L}{\mathcal{E}_u} \right\|_2 \sigma} - |\mathcal{K}| \right)$$

where γ_1 is a constant that depends on the model weights Θ and node features X , γ_2 is a constant that depends on the model weights Θ , and $F_{\chi^2_{|\mathcal{K}|nd}}$ is the c.d.f. of a chi-square r.v. with $|\mathcal{K}| \times n \times d$ degrees of freedom where (n, d) are the row- and column-wise dimensions of X .

Theorem 2 (Probability bound on unfaithfulness w.r.t. Φ). *Under the same assumptions of Lemma 2 and assuming the L_2 distillation error is bounded by α , the unfaithfulness of the explanation \mathcal{E}_u for the original model Φ 's node u prediction is bounded in probability as follows:*

$$\mathbb{P} \left(\frac{1}{|\mathcal{K}|+1} \sum_{\substack{\mathcal{G}'_u \in \\ \mathcal{K} \cup \{\mathcal{G}_u\}}} \|\Phi(\mathcal{G}'_u) - \Phi(t(\mathcal{G}'_u, \mathcal{E}_u))\|_2 \leq \xi \right) \geq \geq F_{\chi^2_{|\mathcal{K}|nd}} \left(\frac{\xi - \gamma_1 \left\| \frac{\Delta \tilde{A}_u^L}{\mathcal{E}_u} \right\|_2 - 2\alpha}{\gamma_2 \left\| \frac{\Delta \tilde{A}_u^L}{\mathcal{E}_u} \right\|_2 \sigma} - |\mathcal{K}| \right)$$

In Lemma 2 and Theorem 2, when the variance σ^2 approaches zero, ξ relinquishes its random nature and the probability in the RHS converges to one. We note that numerators in the RHS must be non-negative.

Recall DnX/FastDnX's pipeline involves two steps: model distillation (Equation 3) and explanation extraction (Equation 5). The former is done only once to learn the surrogate Ψ . The latter, however, must be executed for each node whose prediction we want to explain. Then, gauging the cost of the extraction step may become a genuine concern from a practical point of view, especially for DnX, which implies solving an optimization problem repeatedly. Fortunately, the loss landscape of our extraction problem depends only on the shape of Ψ , and not on the original GNN Φ as in GNNExplainer. Since Ψ is an SGC, Equation 5 is a convex program (Theorem 3) and we reach global optima using, e.g., gradient-based algorithms.

Theorem 3 (Convexity of DnX). *The optimization problem of Equation 5 is convex.*

5 Additional related works

Explanations for GNNs. The ever-increasing application of GNNs to support high-stake decisions on critical domains (Stokes et al., 2020; Jiménez-Luna et al., 2020; Derrow-Pinion et al., 2021) has recently boosted interest in explainability methods for graph models. Pope et al. (2019) first extended classical gradient-based explanation methods for GNNs. Importantly, Ying et al. (2019) introduced GNNExplainer and synthetic benchmarks that have been widely adopted to assess GNN explainers. Building on parameterized explainers by Luo et al. (2020), Wang et al. (2021) proposed ReFine to leverage both global information (e.g., class-wise knowledge) via pre-training and local one (i.e., instance specific patterns) using a fine-tuning process. Lucic et al. (2022); Bajaj et al. (2021) investigated counterfactual explanations for GNNs, aiming to find minimal perturbations to the input graph such that the prediction changes, e.g., using edge deletions. Feng et al. (2021) proposed measuring the contribution of different components of the input graph to the GNN prediction by decomposing the information generation and aggregation mechanism of GNNs. Recently, Zhang et al. (2022b) introduced a structure-aware scoring function derived from cooperative game theory to determine node importance. Explainability methods for GNNs have also been approached through the lens of causal inference (Lin et al., 2021, 2022). For a more comprehensive coverage of the literature, we refer the reader to Yuan et al. (2022).

Knowledge distillation. Since the pivotal work of Hinton et al. (2015), condensing the knowledge from a possibly complex *teacher* model into a simpler *student* surrogate has been an active research topic (e.g. Vadera et al., 2020; Malinin et al., 2020; Ryabinin et al., 2021; Zhou et al., 2022; Hen et al., 2021). Nonetheless, despite numerous works using distillation in image domains (e.g. Rebuffi et al., 2017; Douillard et al., 2021; Baek et al., 2022), the distillation of GNNs is still a blooming direction. Yang et al. (2020) proposed the first method for GNN distillation, using a structure-preserving module to explicitly factor in the topological structure embedded by the teacher. (Joshi et al., 2021) proposed using contrastive learning to implicitly align the node embeddings of the student and the teacher in a common representation space. Jing et al. (2021) combined the knowledge of complementary teacher networks into a single student using a dedicated convolutional operator and topological attribution maps. Zhang et al. (2022a) used an attention mechanism to weigh different teachers depending on the local topology of each node.

6 Experiments

In this section, we assess the performance of DnX and FastDnX on several popular benchmarks, including arti-

cial and real-world datasets. We have implemented experiments using PyTorch (Paszke et al., 2017) and Torch Geometric (Fey and Lenssen, 2019). Our code is available at <https://github.com/tamararruda/DnX>.

6.1 Experimental setup

Datasets. We consider six synthetic datasets broadly used for evaluating explanations of GNNs: BA-House-Shapes, BA-Community, BA-Grids, Tree-Cycles, Tree-Grids, and BA-Bottle-Shaped. These datasets are available in (Ying et al., 2019) and (Vu and Thai, 2020). Each dataset is a single graph with multiple copies of identical motifs connected to base subgraphs. These subgraphs either consists of random sample graphs from the Barabási-Albert (BA) model (Barabási and Albert, 1999) or 8-level balanced binary trees. An explanation associated with a motif-node must only include motif elements. Thus, base nodes denote information irrelevant to the prediction of any node.

We also use two real-world datasets: Bitcoin-Alpha and Bitcoin-OTC (Kumar et al., 2016, 2018). These datasets denote networks in which nodes correspond to user accounts that trade Bitcoin. A directed edge (u, v) (between users u and v) denotes the degree of reliability assigned by u to v , i.e., each edge has a score denoting the degree of trust. Appendix B provides more details regarding datasets.

Baselines. We compare DnX against three baseline explainers: GNNExplainer (Ying et al., 2019), PGExplainer (Luo et al., 2020), and PGMEExplainer (Vu and Thai, 2020). To ensure a valid comparison, we closely follow guidelines and the evaluation setup from the original works. We first generate explanations for a 3-layer GCN (Kipf and Welling, 2017) with ReLU activation. We also consider three additional architectures: graph isomorphism networks (GIN) (Xu et al., 2019), gated graph sequence neural networks (GATED) (Li et al., 2016) and auto-regressive moving average GNNs (ARMA) (Bianchi et al., 2022) This allows for evaluating the robustness and performance of explainers across GNNs of different complexities.

Implementation details. We use an 80/10/10% (train/val/test) split for all datasets. All GNNs have 3 layers and are trained for 1000 epochs, with early stopping if the validation accuracy does not improve in 100 consecutive epochs. We train all baseline GNNs using Adam (Kingma and Ba, 2015) with a learning rate of 0.01 with a weight decay of 5.0×10^{-4} . We show the performance of these GNNs on the benchmark datasets in the supplementary material. Importantly, we observe accuracy $\geq 95\%$ for most data/model combinations.

For the distillation phase in DnX, we use an SGC model with 3 layers. We use the predictions for all nodes to train the surrogate SGC. For the optimization, we use

Table 1: Performance (accuracy) of explanation methods for node-level explanations (i.e., explanations given as subsets of nodes) in the synthetic datasets. Blue and Green numbers denote the best and second-best methods, respectively. Standard deviations are taken over 10 runs of the explanation process, distillation is not included. Since FastDnX’s explanations are deterministic, we mark its variance with not applicable (NA). In most cases, FastDnX achieves the best performance.

Model	Explainer	BA-House	BA-Community	BA-Grids	Tree-Cycles	Tree-Grids	BA-Bottle
GCN	GNNExplainer	77.5 ± 1.2	64.7 ± 1.0	89.2 ± 2.0	77.2 ± 9.0	71.1 ± 1.0	73.3 ± 3.0
	PGExplainer	95.0 ± 1.1	70.6 ± 2.0	86.2 ± 9.0	92.4 ± 5.2	76.7 ± 1.2	98.2 ± 3.0
	PGMExplainer	97.9 ± 0.9	92.2 ± 0.2	88.6 ± 0.9	94.1 ± 0.8	86.8 ± 2.0	97.5 ± 1.5
	DnX	97.7 ± 0.2	94.6 ± 0.1	89.8 ± 0.1	83.3 ± 0.4	80.2 ± 0.1	99.6 ± 0.1
	FastDnX	99.6 ± NA	95.4 ± NA	93.9 ± NA	87.3 ± NA	85.0 ± NA	99.8 ± NA
ARMA	GNNExplainer	80.9 ± 1.2	78.5 ± 1.0	87.3 ± 1.3	77.7 ± 1.0	79.3 ± 1.1	84.3 ± 1.3
	PGExplainer	91.4 ± 0.1	72.1 ± 0.1	83.8 ± 1.0	92.6 ± 2.1	85.1 ± 0.1	97.0 ± 1.1
	PGMExplainer	99.3 ± 0.2	67.5 ± 0.8	86.8 ± 0.3	95.0 ± 0.2	90.6 ± 0.3	99.7 ± 0.1
	DnX	98.1 ± 0.2	92.7 ± 0.2	90.8 ± 0.1	83.5 ± 0.4	79.6 ± 0.3	96.9 ± 0.2
	FastDnX	100.0 ± NA	95.2 ± NA	94.7 ± NA	87.1 ± NA	87.7 ± NA	99.9 ± NA
GATED	GNNExplainer	79.7 ± 1.0	68.8 ± 1.0	91.4 ± 3.0	85.2 ± 2.0	73.2 ± 4.0	70.0 ± 2.0
	PGExplainer	96.1 ± 4.1	70.9 ± 3.0	90.7 ± 1.0	91.7 ± 7.0	83.7 ± 1.5	98.7 ± 0.1
	PGMExplainer	98.6 ± 0.1	69.4 ± 0.5	86.8 ± 0.3	94.1 ± 0.2	90.1 ± 0.2	98.3 ± 0.2
	DnX	98.3 ± 0.1	91.1 ± 0.1	90.8 ± 0.1	85.0 ± 0.3	82.1 ± 0.2	98.0 ± 0.2
	FastDnX	99.6 ± NA	93.5 ± NA	94.0 ± NA	76.8 ± NA	86.8 ± NA	98.0 ± NA
GIN	PGMExplainer	60.2 ± 0.2	84.5 ± 0.3	68.4 ± 0.2	89.3 ± 0.2	85.0 ± 0.5	55.7 ± 0.4
	DnX	99.0 ± 0.1	94.0 ± 0.2	91.1 ± 0.1	84.1 ± 0.3	77.3 ± 0.2	95.3 ± 0.2
	FastDnX	99.6 ± NA	94.7 ± NA	93.9 ± NA	75.2 ± NA	76.5 ± NA	99.1 ± NA

AdamW (Loshchilov and Hutter, 2019) with a learning rate of 0.1 with a weight decay of 5.0×10^{-6} and 10000 epochs.

It is worth mentioning that PGExplainer and GNNExplainer — as described in the experimental section of their respective papers — output edge-level explanations, so their results are not immediately comparable to that of our methods and PGMExplainer. More specifically, the two former output importance scores for each edge. On the other hand, our methods and PGMExplainer output node importance scores. Therefore, we convert edge-level explanations to node-level ones by averaging over the scores of all edges incident in a node. For completeness, we provide additional results doing the reverse transformation (i.e., node- to edge-level explanations) in the Supplement.

6.2 Results

Table 1 compares the performance of DnX and FastDnX against previous art in terms of explanation accuracy, i.e., the number of nodes in method’s output that are also in the ground-truth explanations divided by the total number of nodes in the latter. Overall, FastDnX is the best-performing method for all network architectures (GCN, ARMA, GATED, and GIN) on all datasets but Tree-Cycles and Tree-Grids. For Tree-Grids, FastDnX places second for GCN, ARMA and GATED whereas PGMExplainer ob-

tains the highest accuracies. We also note that, while DnX is often better than GNNExplainer and PGExplainer, its performance bests FastDnX only in 12.5% of cases. GNN- and PGExplainer do not appear in the comparison for GIN since they require propagating edge masks, and Torch Geometric does not support edge features for GIN.

Table 2 reports the performance of all explainers on the Bitcoin-Alpha and Bitcoin-OTC datasets. Following previous work (Vu and Thai, 2020), we use average precision (AP) as evaluation metric, i.e., the percentage of top- k nodes obtained from each explainer that are correct, averaged over all nodes to be explained. While running the experiments, we noticed that the evaluation protocol employed by Vu and Thai (2020) obtains explanations for a 3-layer GCN but only considers 1-hop candidate nodes during the explanation phase. This implies that some potentially relevant nodes are discarded by design. Table 2 shows results for both 1-hop and 3-hop settings. DnX is the best-performing method, and its fast variant is the second-best across all experiments. For 3-hop candidate nodes, the absolute precision gap between DnX and the best baseline is at least 14% for Bitcoin-Alpha and 11% for Bitcoin-OTC. Overall, DnX outperforms GNNExplainer and PGMExplainer by a large margin. Note that the performance of PGMExplainer drops considerably when going from 1- to 3-hop. We report additional results in the Appendix.




Figure 1: Time comparison. The bar plots show the average time each method takes to explain a prediction from GCN. FastDnX is consistently the fastest method, often by a large margin. For the datasets with largest average degree (Bitcoin datasets), FastDnX is 4 orders of magnitude faster than PGMExplainer and 2 orders faster than the other methods.

Table 2: Performance (average precision) of node-level explanations for real-world datasets. Blue and Green numbers denote the best and second-best methods, respectively. DnX significantly outperforms the baselines (GNN-, PG-, and PGM-Explainers).

		Bitcoin-Alpha			Bitcoin-OTC		
GNN	Explainer	top 3	top 4	top 5	top 3	top 4	top 5
GCN (1-hop)	GNNEx	86.3	85.2	81.2	83.3	81.7	77.0
	PGEx	83.5	83.6	79.5	79.9	80.1	76.6
	PGMEx	87.3	85.7	84.8	83.3	81.7	80.8
	DnX	92.2	89.5	88.4	89.4	86.6	84.7
	FastDnX	89.4	87.8	86.8	87.7	85.1	83.4
GCN (3-hop)	GNNEx	80.1	74.9	70.9	82.4	79.6	70.6
	PGEx	81.5	78.1	69.5	78.5	74.5	67.4
	PGMEx	67.0	59.8	51.8	63.0	55.2	47.4
	DnX	95.8	91.9	87.9	94.8	91.4	86.3
	FastDnX	89.8	85.2	80.2	88.0	83.0	78.8

Time comparison. To demonstrate the computational efficiency of DnX/FastDnX, Figure 1 shows the time each method takes to explain a single GCN prediction. For a fair comparison, we also take into account the distillation step in DnX/FastDnX. In particular, we add a fraction – one over the total number of nodes we wish to explain – of the distillation time and add it to the time DnX and FastDnX actually take to generate an explanation. Notably, both DnX and FastDnX are consistently much faster than GNNExplainer and PGMExplainer. For instance, FastDnX is more than forty thousand times faster than PGMExplainer in Bitcoin-Alpha and Bitcoin-OTC.

Distillation results. For completeness, Table 3 shows the distillation accuracy achieved by our linear network Ψ when Φ is a GCN, for both the synthetic and the real datasets. Here, we measure accuracy using the predictions

of the model Φ as ground truth. For all cases, we observe accuracy superior to 86%. Table 3 also shows the time elapsed during the distillation step. Similar results are achieved when distilling ARMA, GATED and GIN models, these results are shown and described in the Appendix.

Table 3: Distillation accuracy and time for GCN. For all cases, accuracy $> 86\%$ and the distillation phase takes considerably less than 1 minute.

Dataset	Accuracy	Time (s)
BA-House	94.2 \pm 1.2	13.996
BA-Community	86.6 \pm 0.1	16.447
BA-Grids	99.9 \pm 0.1	2.721
Tree-Cycles	97.7 \pm 0.2	3.820
Tree-Grids	98.0 \pm 0.2	3.803
BA-Bottle	98.5 \pm 0.2	3.181
Bitcoin-Alpha	90.4 \pm 0.1	28.317
Bitcoin-OTC	89.1 \pm 0.2	32.414

Interestingly, although BA-community is the dataset with the lowest distillation accuracy (86.6%), DnX and FastDnX achieve significantly better results than the previous state-of-the-art (cf. Table 1). The rationale for these counter-intuitive results is that the distiller can differentiate between motif nodes and base nodes, and this is enough to get good explanations – since the evaluation set comprises motif nodes only. More concretely, the confusion matrix in Figure 2 reveals that, despite the low distillation accuracy, the surrogate model Ψ correctly predicts the base nodes (classes 1 and 5). Therefore, Ψ achieves high accuracy for the binary classification problem of distinguishing motif and base nodes, supporting our hypothesis.

Table 4: Performance (fidelity) of different methods for node-level explanations (i.e., explanations given as subsets of nodes) on synthetic datasets. The numbers in **Blue** and **Green** denote the best and second-best methods, respectively. The closer to zero, the better. DnX performs as well as or better than GNNEx, PGEx, and PGMEEx in 5 out of 6 datasets.

Explainer	BA-House	BA-Community	BA-Grids	Tree-Cycles	Tree-Grids	BA-Bottle
GNNEx	0.035	-0.276	0.015	-0.810	-0.120	-0.290
PGEx	0.035	-0.232	-0.194	-0.830	-0.175	0.142
PGMEEx	0.035	-0.290	0.015	-0.677	-0.005	0.025
DnX	0.035	-0.286	0.008	-0.230	-0.001	0.002
FastDnX	0.035	-0.272	-0.018	-0.240	0.000	0.050

Table 5: Performance (fidelity) of different methods for node-level explanations on real-world datasets. The numbers in **Blue** and **Green** denote the best and second-best method, respectively. The closer to zero the better. We show results for sparsity levels of 30%, 50% e 70%. In all cases, FastDnX or DnX are among the two best-performing methods.

Sparsity	Explainer	Bitcoin-Alpha	Bitcoin-OTC	Cora	Citeseer	Pubmed
30%	GNNEx	0.008	0.060	0.015	0.006	0.000
	PGEx	0.101	0.100	0.019	0.051	0.046
	PGMEEx	0.154	0.155	0.013	0.012	-
	DnX	0.028	0.020	0.007	0.006	0.015
	FastDnX	0.012	0.036	0.015	0.006	0.015
50%	GNNEx	0.148	0.040	0.014	0.003	0.006
	PGEx	0.102	0.107	0.014	0.027	0.025
	PGMEEx	0.102	0.118	0.011	-0.003	-
	DnX	0.012	0.018	0.000	0.009	0.010
	FastDnX	0.004	0.056	0.000	0.003	0.005
70%	GNNEx	-0.004	0.016	0.015	-0.009	-0.005
	PGEx	0.091	0.099	-0.003	-0.006	0.005
	PGMEEx	0.088	0.099	0.009	0.008	-
	DnX	0.000	0.000	-0.004	0.003	0.000
	FastDnX	0.004	-0.012	0.000	-0.003	0.010

Fidelity results. To further assess the quality of explanations, we consider a fidelity metric — we use *Fidelity* as in (Yuan et al., 2022). This metric measures how the GNN’s predictive performance (accuracy) fluctuates when we classify nodes based only on the subgraph induced by the explanations. When the fidelity is positive, there is a decrease in performance. When it is negative, using “only the explanation” yields better predictions on average. Tables 4 and 5 report fidelity for the synthetic and the real datasets, respectively. Note that we have considered three additional real-world datasets (citation networks): Cora, Citeseer, and Pubmed. Results obtained from DnX for the synthetic datasets are the best ones in 50% of the cases. It is interesting to observe that for Tree-Cycles and Tree-Grids, DnX/FastDnX are not the best performing ones wrt accuracy (Table 1), but are the best ones wrt fidelity (Table 4). For real datasets, in most cases, either DnX or FastDnX achieves the best results overall. Importantly, this corroborates the results we observed for the precision metric on Bitcoin-Alpha/OTC datasets. We note that it was infeasible to run PGMEExplainer on Pubmed as explaining one prediction with it can take up to an hour in our local hardware.

7 Discussion

Are benchmarks too simple? Given that DnX/FastDnX often achieve remarkable performance by explaining simple surrogates, a natural question arises: *are these popular benchmarks for GNN explanations too simple?* Since these benchmarks rely on model-agnostic ground-truth explanations, we now investigate inductive biases behind these explanations, and show that they can be easily captured.

Figure 3 reports the degree distribution of motif and base nodes for all synthetic datasets. Recall that, by design, ground-truth explanations are always given by motif nodes. Note also that support for the distributions motif and base nodes have almost no overlap for most datasets (except Tree-Cycles & Tree-Grids). Thus, any explainer capable of leveraging degree information would obtain high accuracy.

To make this more concrete, we propose a very simple baseline “explainer” that outputs explanations based on the normalized adjacency matrix. In particular, we define the importance of node j to the prediction of node i as the (i, j) -entry of \tilde{A}^L , with $L = 3$. With this simple baseline, we obtain the following accuracy values: 99.9%

		BA-Community							
		1	2	3	4	5	6	7	8
True labels	1	300	0	0	0	0	0	1	0
	2	0	149	15	0	0	0	0	0
	3	0	12	155	0	1	0	0	0
	4	0	6	58	3	0	1	0	0
	5	0	0	0	0	301	0	0	0
	6	0	3	0	0	0	154	10	0
	7	0	0	3	0	0	14	128	3
	8	0	0	0	0	0	9	58	16
		1	2	3	4	5	6	7	8
		Predicted labels							

Figure 2: Confusion matrix of the distillation process for the BA-Community dataset. Classes 1 and 5 correspond to base nodes. While the surrogate misclassifies many motif nodes, it is able to correctly predict almost all base ones.

(BA-House), 98.1% (BA-Community), 99.9% (BA-Grids), 95.9% (Tree-Cycles), 90.4% (Tree-Grids), and 99.9% (BA-Bottle). Notably, this baseline would rank 1 if included as an explanation method for GCNs in Table 1.

Faber et al. (2021) have also raised issues regarding these benchmarks, proposing alternative datasets as well. We have run FastDnX to explain a 2-layer GCN model for two of their proposed datasets (*Community* and *Negative evidence*), and obtained remarkably good accuracy results: 94.0% and 99.5%, respectively. Also, simply ranking nodes based on the entries of A^L ($L = 2$) achieves accuracy of 93.0% (*Community*) and 99.6% (*Neg. evidence*).

Limitations. While simple graph models (like SGC) have been shown to achieve good performance on node-level classification tasks, they fail to rival recent GNNs for graph-level prediction tasks (Huang et al., 2021; Wu et al., 2019). Naturally, we would not expect DnX and FastDnX to work well out-of-the-shelf to explain graph-level predictions. However, our methods could be easily extended to use more powerful linear GNNs that incorporate different types of diffusion operators (Rossi et al., 2020), or use long-range residual connections (Chen et al., 2020).

8 Conclusion

This work proposes *DnX* as a simple and intuitive two-step framework for post-hoc explanation of GNNs. First, we distill the GNN into a simpler and more interpretable one, that serves as a global surrogate. Then, we leverage the simple structure of the surrogate to extract explanations. Experiments show that (Fast)DnX outperforms the prior art on a variety of benchmarks. Remarkably, our simple design allows FastDnX to run at least $200\times$ faster than relevant

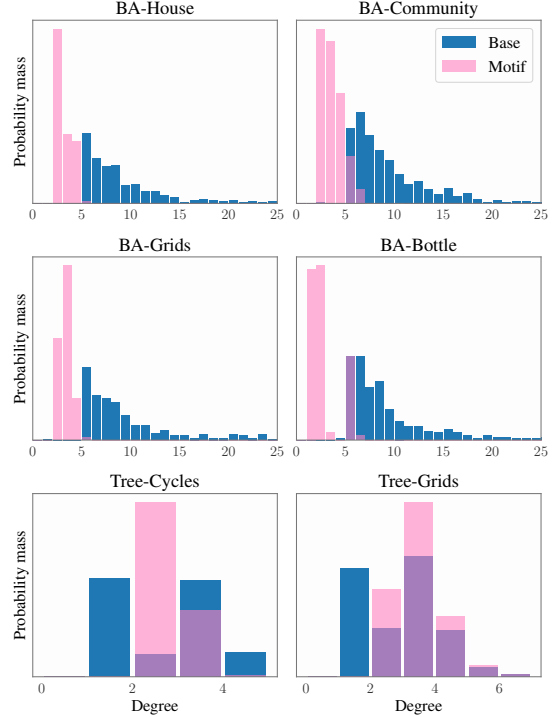


Figure 3: Degree distribution of motif and base nodes. While we can overall distinguish motif and base nodes from degree information on BA-based datasets, there is a significant overlap on Tree-Cycles and Tree-Grids.

baselines on real-world tasks. Additionally, we provide theoretical results that justify our framework and support our empirical findings. Besides advancing the current art, we hope this work will motivate other researchers to focus on developing compute-efficient explainability methods.

Acknowledgments

This work was supported by the Silicon Valley Community Foundation (SVCF) through the Ripple impact fund, the Fundação de Amparo à Pesquisa do Estado do Rio de Janeiro (FAPERJ), the Fundação Cearense de Apoio ao Desenvolvimento Científico e Tecnológico (FUNCAP), the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES), and the Getulio Vargas Foundation’s school of applied mathematics (FGV EMAp).

References

- C. Agarwal, M. Zitnik, and H. Lakkaraju. Probing GNN Explainers: A Rigorous Theoretical and Empirical Analysis of GNN Explanation Methods. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2022.
- D. Baek, Y. Oh, S. Lee, J. Lee, and B. Ham. Decomposed knowledge distillation for class-incremental semantic segmentation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- M. Bajaj, L. Chu, Z. Y. Xue, J. Pei, L. Wang, P. C.-H. Lam, and Y. Zhang. Robust counterfactual explanations on graph neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- A. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439), 1999.
- F. Maria Bianchi, D. Grattarola, L. Livi, and C. Alippi. Graph neural networks with convolutional arma filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(7), 2022.
- M. Chen, Z. Wei, Z. Huang, B. Ding, and Y. Li. Simple and deep graph convolutional networks. In *International Conference on Machine Learning (ICML)*, 2020.
- A. Derrow-Pinion, J. She, D. Wong, O. Lange, T. Hester, L. Perez, M. Nunkesser, S. Lee, X. Guo, B. Wiltshire, P. W. Battaglia, V. Gupta, A. Li, Z. Xu, A. Sanchez-Gonzalez, Y. Li, and P. Velickovic. Eta prediction with graph neural networks in google maps. In *Conference on Information and Knowledge Management (CIKM)*, 2021.
- A. Douillard, Y. Chen, A. Dapogny, and M. Cord. PLOP: Learning without forgetting for continual semantic segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- L. Faber, A. K. Moghaddam, and R. Wattenhofer. When comparing to ground truth is wrong: On evaluating gnn explanation methods. In *Conference on Knowledge Discovery & Data Mining (KDD)*, 2021.
- Q. Feng, N. Liu, F. Yang, R. Tang, M. Du, and X. Hu. Degree: Decomposition based explanation for graph neural networks. In *International Conference on Learning Representations (ICLR)*, 2021.
- M. Fey and J. E. Lenssen. Fast graph representation learning with PyTorch Geometric. In *Workshop on Representation Learning on Graphs and Manifolds (ICLR)*, 2019.
- B. Gao and L. Pavel. On the Properties of the Softmax Function with Application in Game Theory and Reinforcement Learning. *ArXiv:1704.00805*, 2018.
- J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl. Neural message passing for quantum chemistry. In *International Conference on Machine Learning (ICML)*, 2017.
- M. Gori, G. Monfardini, and F. Scarselli. A new model for learning in graph domains. In *IEEE International Joint Conference on Neural Networks (IJCNN)*, 2005.
- W. L. Hamilton. Graph representation learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 14(3):1–159, 2020.
- T. Han, S. Srinivas, and H. Lakkaraju. Which explanation should i choose? a function approximation perspective to characterizing post hoc explanations. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- D. Hen, J. Mei, Y. Zhang, C. Wang, Z. Wang, Y. Feng, and C. Chen. Cross-layer distillation with semantic calibration. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2021.
- G. E. Hinton, O Vinyals, and J. Dean. Distilling the knowledge in a neural network. *Arxiv:1503.02531*, 2015.
- Q. Huang, H. He, A. Singh, S. Lim, and A. Benson. Combining label propagation and simple models outperforms graph neural networks. In *International Conference on Learning Representations (ICLR)*, 2021.
- Q. Huang, M. Yamada, Y. Tian, D. Singh, and Y. Chang. Graphlime: Local interpretable model explanations for graph neural networks. *IEEE Transactions on Knowledge and Data Engineering*, 2022.
- J. Jiménez-Luna, F. Grisoni, and G. Schneider. Drug discovery with explainable artificial intelligence. *Nature Machine Intelligence*, 2(10):573–584, 2020.
- Y. Jing, Y. Yang, X. Wang, M. Song, and D. Tao. Amalgamating knowledge from heterogeneous graph neural networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- C. Joshi, F. Liu, X. Xun, J. Lin, and C. Foo. On representation knowledge distillation for graph neural networks. *arXiv:2111.04964*, 2021.
- D. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.
- T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017.
- S. Kumar, F. Spezzano, V. Subrahmanian, and C. Faloutsos. Edge weight prediction in weighted signed networks. In *International Conference on Data Mining (ICDM)*, 2016.
- S. Kumar, B. Hooi, D. Makhija, M. Kumar, C. Faloutsos, and V. Subrahmanian. Rev2: Fraudulent user prediction in rating platforms. In *International Conference on Web Search and Data Mining (WSDM)*, 2018.
- Y. Li, D. Tarlow, M. Brockschmidt, and R. Zemel. Gated graph sequence neural networks. *International Conference on Learning Representations (ICLR)*, 2016.

-
- W. Lin, H. Lan, and B. Li. Generative causal explanations for graph neural networks. In *International Conference on Machine Learning (ICML)*, 2021.
- W. Lin, H. Lan, H. Wang, and B. Li. Orphicx: A causality-inspired latent variable model for interpreting graph neural networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- I. Loshchilov and F. Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations (ICLR)*, 2019.
- A. Lucic, M. Ter Hoeve, G. Tolomei, M. De Rijke, and F. Silvestri. Cf-gnnexplainer: Counterfactual explanations for graph neural networks. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2022.
- S. M. Lundberg and S.-I. Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- D. Luo, W. Cheng, D. Xu, W. Yu, B. Zong, H. Chen, and X. Zhang. Parameterized explainer for graph neural network. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- A. Malinin, B. Mlodozienec, and M. Gales. Ensemble distribution distillation. In *International Conference on Learning Representations (ICLR)*, 2020.
- A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in pytorch. In *Advances in Neural Information Processing Systems (NeurIPS - Workshop)*, 2017.
- P. E. Pope, S. Kolouri, M. Rostami, C. E. Martin, and H. Hoffmann. Explainability methods for graph convolutional neural networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- S.-A. Rebuffi, A. Kolesnikov, G. Sperl, and C. Lampert. iCaRL: Incremental classifier and representation learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- M. T. Ribeiro, S. Singh, and C. Guestrin. "Why should I trust you?": Explaining the predictions of any classifier. In *International Conference on Knowledge Discovery and Data Mining (KDD)*, 2016.
- E. Rossi, F. Frasca, B. Chamberlain, D. Eynard, M. M. Bronstein, and F. Monti. SIGN: scalable inception graph neural networks. *Workshop on Graph Representation Learning and Beyond (ICML)*, 2020.
- M. Ryabinin, A. Malinin, and M. Gales. Scaling ensemble distribution distillation to many classes with proxy targets. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- A. Sanchez-Gonzalez, J. Godwin, T. Pfaff, R. Ying, J. Leskovec, and P. Battaglia. Learning to simulate complex physics with graph networks. In *International Conference on Machine Learning (ICML)*, 2020.
- F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1), 2009.
- D. Slack, A. Hilgard, S. Singh, and H. Lakkaraju. Reliable post hoc explanations: Modeling uncertainty in explainability. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- J. M. Stokes, K. Yang, K. Swanson, W. Jin, A. Cubillos-Ruiz, N. M. Donghia, C. R. MacNair, S. French, L. A. Carfrae, Z. Bloom-Ackermann, V. M. Tran, A. Chiappino-Pepe, A. H. Badran, I. W. Andrews, E. J. Chory, G. M. Church, E. D. Brown, T. S. Jaakkola, R. Barzilay, and J. J. Collins. A deep learning approach to antibiotic discovery. *Cell*, 180(4), 2020.
- M. Vadera, B. Jalaian, and B. Marlin. Generalized bayesian posterior expectation distillation for deep neural networks. In *Uncertainty in Artificial Intelligence (UAI)*, 2020.
- R. van de Geijn and M. Myers. *Advanced Linear Algebra Foundations to Frontiers*. open EdX Publisher, Austin, Texas, 2022.
- M. Vu and M. T. Thai. Pgm-explainer: Probabilistic graphical model explanations for graph neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- X. Wang, Y. Wu, A. Zhang, X. He, and T. Chua. Towards multi-grained explainability for graph neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- F. Wu, A. Souza, T. Zhang, C. Fifty, T. Yu, and K. Weinberger. Simplifying graph convolutional networks. In *International Conference on Machine Learning (ICML)*, 2019.
- K. Xu, W. Hu, J. Leskovec, and S. Jegelka. How powerful are graph neural networks? *International Conference on Learning Representations (ICLR)*, 2019.
- Y. Yang, J. Qiu, M. Song, D. Tao, and X. Wang. Distilling knowledge from graph convolutional networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7074–7083, 2020.
- R. Ying, R. He, K. Chen, P. Eksombatchai, W. L. Hamilton, and J. Leskovec. Graph convolutional neural networks for web-scale recommender systems. In *International Conference on Knowledge Discovery & Data Mining (KDD)*, 2018.
- Z. Ying, D. Bourgeois, J. You, M. Zitnik, and J. Leskovec. Gnnexplainer: Generating explanations for graph neural

networks. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

H. Yuan, J. Tang, X. Hu, and S. Ji. XGNN: Towards model-level explanations of graph neural networks. In *International Conference on Knowledge Discovery & Data Mining (KDD)*, 2020.

H. Yuan, H. Yu, J. Wang, K. Li, and S. Ji. On explainability of graph neural networks via subgraph explorations. In *International Conference on Machine Learning (ICML)*, 2021.

H. Yuan, H. Yu, S. Gui, and S. Ji. Explainability in graph neural networks: A taxonomic survey. *IEEE transactions on pattern analysis and machine intelligence*, 2022.

C. Zhang, J. Liu, K. Dang, and W. Zhang. Multi-scale distillation from multiple graph neural networks. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2022a.

S. Zhang, N. Shah, Y. Liu, and Y. Sun. Explaining graph neural networks with structure-aware cooperative games. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022b.

Y. Zhou, E. Nezhadarya, and J. Ba. Dataset distillation using neural feature regression. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

A Proofs

Proof of Lemma 1. For node u , it is known, by definition, that

$$\begin{aligned} & \|\Psi(\mathcal{G}_u) - \Psi(t(\mathcal{G}_u, \mathcal{E}_u))\|_2 \\ &= \left\| \sigma \left(\tilde{A}^L X \Theta \right)_u - \sigma \left(\tilde{E}^L X \Theta \right)_u \right\|_2 \\ &= \left\| \sigma \left[\left(\tilde{A}^L X \Theta \right)_u \right] - \sigma \left[\left(\tilde{E}^L X \Theta \right)_u \right] \right\|_2, \end{aligned}$$

if we call σ the softmax function and \tilde{E}^L the powered, normalized adjacency matrix \tilde{A}^L after applying the explanation \mathcal{E}_u . Because softmax is a Lipschitz continuous function with Lipschitz constant 1 with respect to norm $\|\cdot\|_2$ (Gao and Pavel, 2018) and considering induced matrix norm compatibility property (van de Geijn and Myers, 2022, Lemma 1.3.8.7),

$$\begin{aligned} & \|\Psi(\mathcal{G}_u) - \Psi(t(\mathcal{G}_u, \mathcal{E}_u))\|_2 \\ &= \left\| \sigma \left[\left(\tilde{A}^L X \Theta \right)_u \right] - \sigma \left[\left(\tilde{E}^L X \Theta \right)_u \right] \right\|_2 \\ &\leq \|(X \Theta)^\top\|_2 \left\| \tilde{A}_u^L - \tilde{E}_u^L \right\|_2. \end{aligned}$$

Similarly, for a perturbation \mathcal{G}'_u of \mathcal{G}_u given by Σ'_u being added to X ,

$$\begin{aligned} & \|\Psi(\mathcal{G}'_u) - \Psi(t(\mathcal{G}'_u, \mathcal{E}_u))\|_2 \\ &\leq \|[(X + \Sigma'_u) \Theta]^\top\|_2 \left\| \tilde{A}_u^L - \tilde{E}_u^L \right\|_2. \end{aligned}$$

Computing the mean over $\mathcal{K} \cup \{\mathcal{G}_u\}$:

$$\begin{aligned} & \frac{1}{|\mathcal{K}| + 1} \sum_{\mathcal{G}'_u \in \mathcal{K} \cup \{\mathcal{G}_u\}} \|\Psi(\mathcal{G}'_u) - \Psi(t(\mathcal{G}'_u, \mathcal{E}_u))\|_2 \\ &\leq \frac{1}{|\mathcal{K}| + 1} \left[\|(X \Theta)^\top\|_2 \left\| \tilde{A}_u^L - \tilde{E}_u^L \right\|_2 \right. \\ &\quad \left. + \sum_{\mathcal{G}'_u \in \mathcal{K}} \|[(X + \Sigma'_u) \Theta]^\top\|_2 \left\| \tilde{A}_u^L - \tilde{E}_u^L \right\|_2 \right] \\ &\leq \|\Theta^\top\|_2 \left\| \tilde{A}_u^L - \tilde{E}_u^L \right\|_2 \frac{1}{|\mathcal{K}| + 1} \left(\|X^\top\|_2 \right. \\ &\quad \left. + \sum_{\mathcal{G}'_u \in \mathcal{K}} \|(X + \Sigma'_u)^\top\|_2 \right) \\ &\leq \gamma_\Theta \left\| \Delta_{\mathcal{E}_u} \tilde{A}_u^L \right\|_2 \max \left(\{\|X^\top\|_2\} \cup \{\|(X + \Sigma'_u)^\top\|_2\}_{\mathcal{G}'_u \in \mathcal{K}} \right) \\ &= \gamma_{\Theta, X, \Sigma} \left\| \Delta_{\mathcal{E}_u} \tilde{A}_u^L \right\|_2. \end{aligned}$$

When Σ is limited, the constant γ may not depend on Σ . \square

Proof of Theorem 1. We know that

$$\begin{aligned} & \|\Phi(\mathcal{G}'_u) - \Phi(t(\mathcal{G}'_u, \mathcal{E}_u))\|_2 \\ &= \|\Phi(\mathcal{G}'_u) - \Psi(\mathcal{G}'_u) + \Psi(\mathcal{G}'_u) - \Psi(t(\mathcal{G}'_u, \mathcal{E}_u)) \\ &\quad + \Psi(t(\mathcal{G}'_u, \mathcal{E}_u)) - \Phi(t(\mathcal{G}'_u, \mathcal{E}_u))\|_2. \end{aligned}$$

So, by using triangle inequality and Lemma 1,

$$\begin{aligned} & \frac{1}{|\mathcal{K}| + 1} \sum_{\mathcal{G}'_u \in \mathcal{K} \cup \{\mathcal{G}_u\}} \|\Phi(\mathcal{G}'_u) - \Phi(t(\mathcal{G}'_u, \mathcal{E}_u))\|_2 \\ &\leq \frac{1}{|\mathcal{K}| + 1} \sum_{\mathcal{G}'_u \in \mathcal{K} \cup \{\mathcal{G}_u\}} \|\Phi(\mathcal{G}'_u) - \Psi(\mathcal{G}'_u)\|_2 \\ &\quad + \frac{1}{|\mathcal{K}| + 1} \sum_{\mathcal{G}'_u \in \mathcal{K} \cup \{\mathcal{G}_u\}} \|\Psi(\mathcal{G}'_u) - \Psi(t(\mathcal{G}'_u, \mathcal{E}_u))\|_2 \\ &\quad + \frac{1}{|\mathcal{K}| + 1} \sum_{\mathcal{G}'_u \in \mathcal{K} \cup \{\mathcal{G}_u\}} \|\Psi(t(\mathcal{G}'_u, \mathcal{E}_u)) - \Phi(t(\mathcal{G}'_u, \mathcal{E}_u))\|_2 \\ &\leq \gamma \left\| \Delta_{\mathcal{E}_u} \tilde{A}_u^L \right\|_2 + 2\alpha. \end{aligned}$$

\square

Proof of Lemma 2. From the proof of Lemma 1, we know that

$$\begin{aligned} & \frac{1}{|\mathcal{K}| + 1} \sum_{\mathcal{G}'_u \in \mathcal{K} \cup \{\mathcal{G}_u\}} \|\Psi(\mathcal{G}'_u) - \Psi(t(\mathcal{G}'_u, \mathcal{E}_u))\|_2 \leq \\ & \|\Theta^\top\|_2 \left\| \Delta_{\mathcal{E}_u} \tilde{A}_u^L \right\|_2 \frac{1}{|\mathcal{K}| + 1} \left(\|X^\top\|_2 + \sum_{\mathcal{G}'_u \in \mathcal{K}} \|(X + \Sigma'_u)^\top\|_2 \right) \end{aligned}$$

By using triangle inequality, we can write

$$\begin{aligned} & \frac{1}{|\mathcal{K}| + 1} \sum_{\mathcal{G}'_u \in \mathcal{K} \cup \{\mathcal{G}_u\}} \|\Psi(\mathcal{G}'_u) - \Psi(t(\mathcal{G}'_u, \mathcal{E}_u))\|_2 \\ &\leq \|\Theta^\top\|_2 \left\| \Delta_{\mathcal{E}_u} \tilde{A}_u^L \right\|_2 \|X^\top\|_2 \\ &\quad + \|\Theta^\top\|_2 \left\| \Delta_{\mathcal{E}_u} \tilde{A}_u^L \right\|_2 \frac{1}{|\mathcal{K}| + 1} \sum_{\mathcal{G}'_u \in \mathcal{K}} \|(X + \Sigma'_u)^\top\|_2. \end{aligned}$$

We can work in the perturbation summand:

$$\begin{aligned}
\sum_{\mathcal{G}'_u \in \mathcal{K}} \|(\Sigma'_u)^\top\|_2 &\leq \sum_{\mathcal{G}'_u \in \mathcal{K}} \|\Sigma'_u\|_F \\
&= \sum_{\mathcal{G}'_u \in \mathcal{K}} \sqrt{\sum_{i=1}^{nd} \epsilon_{u',i}^2} \\
&= \sigma \sum_{\mathcal{G}'_u \in \mathcal{K}} \sqrt{\sum_{i=1}^{nd} Z_{u',i}^2} \\
&\leq \sigma \sum_{\mathcal{G}'_u \in \mathcal{K}} \max\left(1, \sum_{i=1}^{nd} Z_{u',i}^2\right) \\
&\leq \sigma \sum_{\mathcal{G}'_u \in \mathcal{K}} \left(1 + \sum_{i=1}^{nd} Z_{u',i}^2\right) \\
&= \sigma |\mathcal{K}| + \sigma Q,
\end{aligned}$$

with

- $\epsilon_{j,i} \sim \mathcal{N}(0, \sigma^2)$;
- $Z_{j,i} \sim \mathcal{N}(0, 1)$;
- $Q \sim \chi_{|\mathcal{K}|nd}^2$;
- $(n, d) = \dim(X)$.

If α_p is the p -percentile of Q , then the probability of the last inequality is p :

$$\begin{aligned}
&\frac{1}{|\mathcal{K}| + 1} \sum_{\mathcal{G}'_u \in \mathcal{K} \cup \{\mathcal{G}_u\}} \|\Psi(\mathcal{G}'_u) - \Psi(t(\mathcal{G}'_u, \mathcal{E}_u))\|_2 \\
&\leq \|\Theta^\top\|_2 \left\| \frac{\Delta \tilde{A}_u^L}{\mathcal{E}_u} \right\|_2 \|X^\top\|_2 \\
&\quad + \|\Theta^\top\|_2 \left\| \frac{\Delta \tilde{A}_u^L}{\mathcal{E}_u} \right\|_2 \frac{1}{|\mathcal{K}| + 1} \sum_{\mathcal{G}'_u \in \mathcal{K}} \|(\Sigma'_u)^\top\|_2 \\
&\leq \gamma_1 \left\| \frac{\Delta \tilde{A}_u^L}{\mathcal{E}_u} \right\|_2 + \|\Theta^\top\|_2 \left\| \frac{\Delta \tilde{A}_u^L}{\mathcal{E}_u} \right\|_2 \frac{1}{|\mathcal{K}| + 1} \sigma(Q + |\mathcal{K}|) \\
&\leq \gamma_1 \left\| \frac{\Delta \tilde{A}_u^L}{\mathcal{E}_u} \right\|_2 + \|\Theta^\top\|_2 \left\| \frac{\Delta \tilde{A}_u^L}{\mathcal{E}_u} \right\|_2 \frac{1}{|\mathcal{K}| + 1} \sigma(\alpha_p + |\mathcal{K}|) \\
&= \gamma_1 \left\| \frac{\Delta \tilde{A}_u^L}{\mathcal{E}_u} \right\|_2 + \gamma_2 \left\| \frac{\Delta \tilde{A}_u^L}{\mathcal{E}_u} \right\|_2 \sigma(\alpha_p + |\mathcal{K}|)
\end{aligned}$$

Notice that, when σ approaches zero, the bound's probabilistic characteristic becomes negligible.

Finally, if we ask the bound to be ξ , then the probability is

$$p = F_{\chi_{|\mathcal{K}|nd}^2} \left(\frac{\xi - \gamma_1 \left\| \frac{\Delta \tilde{A}_u^L}{\mathcal{E}_u} \right\|_2 - |\mathcal{K}|}{\gamma_2 \left\| \frac{\Delta \tilde{A}_u^L}{\mathcal{E}_u} \right\|_2 \sigma} \right),$$

$F_{\chi_{|\mathcal{K}|nd}^2}$ being the c.d.f. of Q . \square

Proof of Theorem 3. The objective function of the problem in Equation 5 can be written as

$$\begin{aligned}
f(\mathcal{E}) &= \left\| \tilde{A}_i^L \text{diag}(\mathcal{E}) X \Theta - \tilde{A}_i^L X \Theta \right\|_2^2 \\
&= \left(\tilde{A}_i^L \text{diag}(\mathcal{E}) X \Theta - \tilde{A}_i^L X \Theta \right) \cdot \\
&\quad \cdot \left(\tilde{A}_i^L \text{diag}(\mathcal{E}) X \Theta - \tilde{A}_i^L X \Theta \right)^\top \\
&= \mathcal{E}^\top \text{diag} \left[\left(\tilde{A}_i^L \right)^\top \right] X \Theta \Theta^\top X^\top \text{diag} \left[\left(\tilde{A}_i^L \right)^\top \right] \mathcal{E} \\
&\quad - 2 \mathcal{E}^\top \text{diag} \left[\left(\tilde{A}_i^L \right)^\top \right] X \Theta \Theta^\top X^\top \left(\tilde{A}_i^L \right)^\top \\
&\quad + \left\| \tilde{A}_i^L X \Theta \right\|_2^2 \\
&= \frac{1}{2} \mathcal{E}^\top Q \mathcal{E} + \mathcal{E}^\top c + \delta.
\end{aligned}$$

Note also that

$$Q = P^\top P \text{ with } P = \sqrt{2} \Theta^\top X^\top \text{diag} \left[\left(\tilde{A}_i^L \right)^\top \right],$$

thus Q is symmetric and positive semidefinite. Since both the objective function and feasible set are convex, the optimization problem is also convex. \square

B Datasets and implementation details

B.1 Datasets

Bitcoin-Alpha and Bitcoin-OTC are real-world networks comprising 3783 and 5881 nodes (user accounts), respectively. Users rate their trust in each other using a score between -10 (total distrust) and 10 (total trust). Then, user accounts are labeled as trusted or not based on how other users rate them. Accounts (nodes) have features such as average ratings. Ground-truth explanations for each node are provided by experts. The synthetic datasets are available in (Ying et al., 2019) and (Vu and Thai, 2020). Table 6 shows summary statistics for all datasets.

Table 6: Statistics of the datasets used in our experiments.

Dataset	nodes	edges	labels
BA-House	700	4110	4
BA-Community	1400	8920	8
BA-Grids	1020	5080	2
Tree-Cycles	871	1950	2
Tree-Grids	1231	3410	2
BA-Bottle	700	3948	4
Bitcoin-Alpha	3783	28248	2
Bitcoin-OTC	5881	42984	2

B.2 Implementation details

We ran all experiments using a laptop with an Intel Xeon 2.20 GHz CPU, 13 GB RAM DDR5, and RTX 3060ti 16GB GPU.

Table 7: Performance (mean and standard deviation of accuracy) in node classification tasks for the models to be explained.

	BA-House	BA-Community	BA-Grids	Tree-Cycles	Tree-Grids	BA-Bottle	Bitcoin-Alpha	Bitcoin-OTC
GCN	97.9 ± 1.6	85.6 ± 1.7	99.9 ± 0.2	97.8 ± 1.2	88.9 ± 1.8	99.0 ± 0.1	93.3 ± 0.4	93.2 ± 0.6
ARMA	98.1 ± 2.3	92.8 ± 2.3	99.5 ± 0.6	96.9 ± 1.5	92.4 ± 2.3	99.6 ± 0.9	93.6 ± 1.4	92.9 ± 1.1
GATED	98.0 ± 1.3	92.3 ± 2.7	99.9 ± 0.2	97.8 ± 2.2	94.4 ± 3.0	99.4 ± 0.6	94.4 ± 1.3	93.8 ± 0.7
GIN	95.6 ± 5.0	87.0 ± 1.5	99.4 ± 0.6	97.8 ± 1.1	91.4 ± 2.6	98.8 ± 1.0	93.4 ± 1.0	92.6 ± 1.1

The architecture of the GNNs to be explained in this work are: *i*) a GCN model with 3 layers (20 hidden units) followed by a two-layer MLP with 60 hidden units; *ii*) an ARMA model with 3 layers (20 hidden units each) followed by a two-layer MLP with 60 hidden units; *iii*) a GIN model with 3 layers (20 hidden units each); and *iv*) a GATED model with 3 layers (100 hidden units) followed by a two-layer MLP with 300 hidden units. We train all these GNNs using with a learning rate of 0.001. All models use ReLU activations in their hidden units. We also use the one-hot vector of the degree as node features.

For computational reasons, on Bitcoin datasets, we provide results for 500 test nodes whenever the candidate explanation set contemplates 3-hop neighborhood (matching the GNNs we want to explain). For 1-hop cases, we use the full set of 2000 nodes as in the original setup.

Node and edge-level explanations. DnX and FastDnX were originally designed to generate node-level explanations like PGM-Explainer but some baselines such as PG-explainer and GNNexplainer provide edge-level explanations. For a fair comparison with these baselines, we use a procedure to convert edge-level explanations to node-level ones (and vice-versa). To convert node scores to edge scores, we sum the scores of endpoints of each edge, i.e., an edge (u, v) gets score $s_{u,v} = s_u + s_v$ where s_u and s_v are node scores. For the reverse, we say $s_u = |\mathcal{N}_u|^{-1} \sum_{j \in \mathcal{N}_u} s_{u,j}$ where \mathcal{N}_u is the neighborhood of u .

C Additional experiments

GNN performance. Table 7 shows the classification accuracy of each model we explain in our experimental campaign. Means and standard deviations reflect the outcome of 10 repetitions. All classifiers achieve accuracy > 85%.

Distillation. Table 8 shows the extent to which the distiller network Ψ agrees with the GNN Φ . We measure agreement as accuracy, using the predictions of Φ as ground truth. The means and standard deviations reflect the outcome of 10 repetitions. Additionally, Table 8 shows the time elapsed during the distillation step. For all cases, we observe accuracy values over 88%.

Results for edge-level explanations. Table 9 and Table 10 complement Tables 1 and 2 in the main paper,

Table 8: Distillation accuracy and elapsed time for ARMA, GATED and GIN. For all cases, accuracy > 88% and learning Ψ takes less than a minute.

Model	Dataset	Accuracy	Time (s)
ARMA	BA-House	97.7 ± 0.01	7.254
	BA-Community	96.7 ± 0.03	22.848
	BA-Grids	100.0 ± 0.00	2.701
	Tree-Cycles	98.8 ± 0.03	6.331
	Tree-Grids	91.7 ± 0.15	4.152
	BA-Bottle	100.0 ± 0.00	6.312
	Bitcoin-Alpha	91.3 ± 0.01	28.613
	Bitcoin-OTC	91.2 ± 0.03	39.832
GATED	BA-House	98.0 ± 0.01	5.954
	BA-Community	92.1 ± 0.04	27.847
	BA-Grids	100.0 ± 0.00	2.679
	Tree-Cycles	99.4 ± 0.04	6.643
	Tree-Grids	90.9 ± 0.09	4.609
	BA-Bottle	100.0 ± 0.00	6.171
	Bitcoin-Alpha	91.3 ± 0.02	26.611
	Bitcoin-OTC	91.1 ± 0.01	42.626
GIN	BA-House	93.4 ± 0.46	11.461
	BA-Community	90.1 ± 0.13	24.204
	BA-Grids	100.0 ± 0.00	3.884
	Tree-Cycles	95.8 ± 0.02	3.335
	Tree-Grids	88.1 ± 0.06	5.198
	BA-Bottle	100.0 ± 0.01	6.970
	Bitcoin-Alpha	95.5 ± 0.03	44.359
	Bitcoin-OTC	89.8 ± 2.04	55.860

showing results for edge-level explanations. All experiments were repeated ten times. For all datasets, we measure performance in terms of AUC, following Luo et al. (2020). Both tables corroborate our findings from the main paper. In most cases, FastDnX is the best or second-best model for all models and datasets. For GCN and GATED, PGExplainer yields the best results. Overall, both DnX and FastDnX outperform GNNExplainer and PGM-Explainer. Remarkably FastDnX and DnX’s performance is steady, with small fluctuations depending on the model we are explaining. The same is not true for the competing methods, e.g., PGExplainer loses over 15% AUC for the BA-Community (cf., GCN and ARMA). Note also that FastDnX and DnX are the best models on the real-world datasets.

Table 9: Performance (mean and standard deviation of AUC) of explainer models on synthetic datasets for edge-level explanations. **Blue** and **Green** numbers denote the best and second-best methods, respectively.

Model	Explainer	BA-House	BA-Community	BA-Grids	Tree-Cycles	Tree-Grids	BA-Bottle
GCN	GNNExplainer	82.4 ± 6.0	71.1 ± 6.0	80.9 ± 1.0	58.4 ± 1.0	53.9 ± 2.0	82.8 ± 1.0
	PGExplainer	99.9 ± .01	99.9 ± .01	94.1 ± 9.0	92.3 ± 5.0	79.4 ± 2.0	99.9 ± .01
	PGMExplainer	56.2 ± 0.3	53.0 ± 0.4	68.9 ± 0.2	62.1 ± 0.2	66.3 ± 0.4	53.4 ± 0.5
	DnX	95.7 ± .09	87.5 ± .09	97.3 ± .03	80.5 ± .02	72.6 ± 0.1	94.5 ± 0.1
	FastDnX	99.4 ± NA	93.0 ± NA	96.7 ± NA	89.1 ± NA	77.4 ± NA	99.5 ± NA
ARMA	GNNExplainer	80.8 ± 1.0	75.1 ± 1.0	69.4 ± 0.1	59.2 ± 2.0	59.9 ± 1.0	85.8 ± .01
	PGExplainer	93.3 ± 1.0	82.7 ± .01	93.1 ± .01	92.4 ± 2.0	84.0 ± .03	98.0 ± .02
	PGMExplainer	82.2 ± 0.2	49.6 ± 0.7	40.4 ± 0.2	64.7 ± 0.2	69.0 ± 0.3	63.1 ± 0.3
	DnX	95.7 ± .06	87.9 ± .07	97.3 ± .04	80.4 ± 0.3	72.8 ± 0.1	94.4 ± .08
	FastDnX	99.7 ± NA	96.8 ± NA	97.9 ± NA	87.1 ± NA	78.4 ± NA	99.7 ± NA
GATED	GNNExplainer	79.1 ± 2.0	59.2 ± 3.0	69.4 ± 2.0	62.6 ± 2.0	53.6 ± 1.0	73.9 ± 2.0
	PGExplainer	99.9 ± 2.0	99.1 ± 2.0	99.8 ± 1.0	76.5 ± 9.0	97.0 ± 3.0	99.9 ± 0.1
	PGMExplainer	49.3 ± 0.4	47.2 ± 0.5	46.2 ± 0.3	45.8 ± 0.6	53.1 ± 0.8	49.1 ± 0.7
	DnX	95.8 ± .07	84.4 ± .09	97.3 ± .02	81.1 ± 0.1	72.8 ± 0.1	94.4 ± .09
	FastDnX	99.5 ± NA	93.7 ± NA	96.6 ± NA	89.2 ± NA	78.6 ± NA	99.6 ± NA
GIN	PGMExplainer	52.4 ± 0.3	52.0 ± 0.4	52.2 ± 1.0	67.4 ± 0.2	63.9 ± 0.2	53.3 ± 0.4
	DnX	95.7 ± 0.1	87.2 ± 0.1	97.3 ± 0.1	81.8 ± 0.2	72.0 ± 0.1	94.2 ± 0.1
	FastDnX	97.9 ± NA	91.0 ± NA	97.1 ± NA	86.3 ± NA	72.3 ± NA	97.0 ± NA

Table 10: Performance (AUC) of edge-level explanations for real-world datasets, measured over 500 test nodes.

Model	Explainer	Bitcoin-Alpha	Bitcoin-OTC
GCN (3-hop)	GNNex	94.0	97.3
	PGEx	59.7	54.2
	PGMEx	75.8	52.7
	DnX	97.1	98.1
	FastDnX	97.3	99.1
ARMA (3-hop)	GNNex	79.6	92.3
	PGEx	47.8	45.4
	PGMEx	76.2	73.8
	DnX	96.4	98.7
	FastDnX	97.4	99.1
GATED (3-hop)	GNNex	92.5	93.6
	PGEx	34.7	33.7
	PGMEx	86.3	83.1
	DnX	96.3	98.6
	FastDnX	97.2	99.1
GIN (3-hop)	PGMEx	52.8	76.1
	DnX	96.4	98.7
	FastDnX	97.5	98.9

More results for node-level explanations. Table 11 shows results for ARMA, GATED, and GIN. As observed in Table 2 for GCNs, DnX FastDnX are the best methods.

Table 11: Performance (average precision) of node-level explanations for real-world datasets, measured over 500 test nodes.

Model	Explainer	Bitcoin-Alpha			Bitcoin-OTC		
		top 3	top 4	top 5	top 3	top 4	top 5
ARMA (3-hop)	GNNEx	80.9	78.9	74.5	73.2	69.5	64.0
	PGEx	72.5	67.8	65.0	69.7	68.7	61.45
	PGMEx	73.8	66.4	58.3	69.1	62.8	54.9
	DnX	95.0	90.1	84.9	93.6	88.1	83.2
GATED (3-hop)	FastDnX	90.2	85.8	81.0	87.7	83.3	79.2
	GNNEx	80.1	75.1	769.6	75.9	70.9	66.0
	PGEx	75.4	74.8	67.8	72.5	70.5	65.1
	PGMEx	80.2	76.5	72.4	77.5	72.1	67.3
GIN (3-hop)	DnX	94.4	89.6	84.4	93.4	88.9	83.5
	FastDnX	89.6	85.3	80.0	88.0	83.4	79.0
GIN (3-hop)	PGM-Ex	58.7	49.2	40.8	57.6	47.8	40.2
	DnX	94.3	88.9	83.2	94.0	89.2	83.6
	FastDnX	85.0	80.4	74.6	82.6	77.1	71.3