CF-GNNExplainer: Counterfactual Explanations for Graph Neural Networks

Ana Lucic [1] **Maartje ter Hoeve** [1] **Gabriele Tolomei** [2] **Maarten de Rijke** [1] **Fabrizio Silvestri** [2]

Abstract

Given the increasing promise of Graph Neural Networks (GNNs) in real-world applications, several methods have been developed for explaining their predictions. So far, these methods have primarily focused on generating subgraphs that are especially relevant for a particular prediction. However, such methods do not provide a clear opportunity for recourse: given a prediction, we want to understand how the prediction can be changed in order to achieve a more desirable outcome. In this work, we propose a method for generating counterfactual (CF) explanations for GNNs: the minimal perturbation to the input (graph) data such that the prediction changes. Using only edge deletions, we find that our method can generate CF explanations for the majority of instances across three widely used datasets for GNN explanations, while removing less than 3 edges on average, with at least 94% accuracy. This indicates that our method primarily removes edges that are crucial for the original predictions, resulting in minimal CF explanations.

1. Introduction

Advances in machine learning (ML) have led to breakthroughs in several areas of science and engineering, ranging from computer vision, to natural language processing, to conversational assistants. Parallel to the increased performance of ML systems, there is an increasing call for the "understandability" of ML models [13]. Understanding *why* an ML model returns a certain output in response to a given input is important for a variety of reasons such as model debugging, aiding decison-making, or fulfilling legal requirements [11]. Having certified methods for interpreting ML predictions will help enable their use across a variety of applications [29]. Explainable AI (XAI) refers to the set of techniques "*focused on exposing complex AI models to humans in a systematic and interpretable manner*" [32]. A large body of work on XAI has emerged in recent years [15; 4]. Counterfactual (CF) explanations are used to explain predictions of individual instances in the form: "If X had been different, Y would not have occurred" [36]. CF explanations are based on CF examples: modified versions of the input sample that result in an alternative output response (i.e., prediction). If the modifications recommended are also *actionable*, this is referred to as achieving recourse [40; 20].

To motivate our problem, we consider an ML application for computational biology. Drug discovery is a task that involves generating new molecules that could be used for medicinal purposes [37; 45]. Given a candidate molecule, a GNN can predict if this molecule has a certain property that would make it effective in treating a particular disease [44; 16; 30]. If the GNN predicts it does not have this desirable property, CF explanations can help identify the minimal change one should make to this molecule, such that it has this desirable property. This could help us not only generate a new molecule that has this property, but also understand what molecular structures contribute to this property.

Although GNNs have shown state-of-the-art results on tasks involving graph data [50; 6], existing methods for explaining the predictions of GNNs have primarily focused on generating subgraphs that are relevant for a particular prediction [48; 1; 9; 24; 28; 31; 33; 42; 46; 49]. However, none of these methods are able to identify the *minimal* subgraph automatically since they all require the user to specify in advance the size of the subgraph, S. We show that even if we adapt existing methods to the CF explanation problem, and try varying values for S, such methods are not able to produce valid, accurate explanations, and are therefore not well-suited to solve the CF explanation problem. To address this gap, we propose CF-GNNEXPLAINER, a method for generating CF explanations for GNNs.

Similar to other CF methods for tabular or image data proposed in the literature [41; 20], CF-GNNEXPLAINER works by perturbing input data at the instance-level. In this work, the instances are nodes in the graph since we focus on node classification. In particular, our method iteratively

[1]Informatics Institute, University of Amsterdam, Netherlands [2]Department of Computer Science, Sapienza University of Rome, Italy. Correspondence to: Ana Lucic <a.lucic@uva.nl>.

ICML (International Conference on Machine Learning) Workshop on Human in the Loop Learning. Copyright 2021 by the author(s).

removes edges from the original adjacency matrix based on matrix sparsification techniques, keeping track of the perturbations that lead to a change in prediction, and returning the perturbation with the smallest change w.r.t. the number of edges. We evaluate CF-GNNEXPLAINER on three public datasets for GNN explanations and measure its effectiveness using four metrics: fidelity, explanation size, sparsity, and accuracy. We find that CF-GNNEXPLAINER is able to generate CF examples with at least 94% accuracy, while removing fewer than 3 edges on average.

We make the following contributions: (i) we formalize the problem of generating CF explanations for GNNs (Section 3); (ii) we propose CF-GNNEXPLAINER, a novel method for explaining predictions from GNNs (Section 4); and (iii) we propose an experimental setup for holistically evaluating CF GNN explanations (Section 5).

2. Background

2.1. Graph Neural Networks

Graphs are structures that represent a set of entities (nodes) and their relations (edges). GNNs operate on graphs to produce representations that can be used in downstream tasks such as graph or node classification. The latter is the focus of this work. We refer to Battaglia et al. [3] and Chami et al. [5] for an extensive overview of existing GNN methods.

Let $f(A, X; W)$ be any GNN, where A is an $n \times n$ adjacency matrix, X is an $n \times p$ feature matrix (with p features), and W are the learned weights of f. In other words, A and X are the inputs of f, and f is parameterized by W.

A node's representation is learned by iteratively updating the node's features based on its neighbors' features. The number of layers in f determines which neighbors are included: if there are ℓ layers, then the node's final representation only includes neighbors that are at most ℓ hops away from that node in the graph $\mathcal{G}$ – the rest of the nodes in $\mathcal{G}$ are not relevant for the computation of the node's final representation. We define the *subgraph neighbourhood* of a node v as a tuple of the nodes and edges relevant for the computation of $f(v)$ (i.e., those in the ℓ-hop neighborhood of f): $\mathcal{G}_v = (A_v, X_v)$, where A_v is the subgraph adjacency matrix and X_v is the node feature matrix for nodes that are at most ℓ hops away from v. We then define a node v as a tuple of the form $v = (A_v, x)$, where x is the feature vector for v.

2.2. Matrix Sparsification

CF-GNNEXPLAINER uses matrix sparsification to generate CF examples, inspired by Srinivas et al. [35]. They propose a method for training sparse neural networks: given a weight matrix W, a binary sparsification matrix is learned which is multiplied element-wise with W such that some of the entries in W are zeroed out. In Srinivas et al. [35], the objective is to remove entries in the weight matrix in order to reduce the number of parameters in the model. In our case, instead of learning a sparsification matrix to *zero out weights*, we want to *zero out entries in the adjacency matrix* (i.e., remove edges) in order to generate CF explanations for GNNs. That is, we want to remove the important edges – those that are crucial for the prediction.

3. Problem Formulation

In this section, we formalize the problem of generating CF explanations for GNNs. In general, a CF example $\bar{x}$ for an instance x according to a trained classifier f is found by perturbing the features of x such that $f(x) \neq f(\bar{x})$ [43]. An optimal CF example $\bar{x}^*$ is one that minimizes the distance between the original instance and the CF example, according to some distance function d, and the resulting optimal CF explanation is $\Delta_x^* = \bar{x}^* - x$ [26].

For graph data, it may not be enough to simply perturb node features, especially since they are not always available. This is why we are interested in generating CF examples by perturbing the graph structure instead. In other words, we want to change the relationships between instances, rather than change the instances themselves. Therefore, a CF example for graph data has the form $\bar{v} = (\bar{A}_v, x)$, where x is the feature vector and $\bar{A}_v$ is a perturbed version of A_v, the adjacency matrix of the subgraph neighborhood of a node v. $\bar{A}_v$ is obtained by removing some edges from A_v, such that $f(v) \neq f(\bar{v})$. Following Wachter et al. [43] and Lucic et al. [26], we generate CF examples by minimizing a loss function of the form:

$$\mathcal{L} = \mathcal{L}_{pred}(v, \bar{v} \mid f, g) + \beta \mathcal{L}_{dist}(v, \bar{v}), \tag{1}$$

where v is the original node and f is the original model; g is the CF model that generates $\bar{v}$; and $\mathcal{L}_{pred}$ is a prediction loss that encourages $f(v) \neq f(\bar{v})$. $\mathcal{L}_{dist}$ is a distance loss that encourages $\bar{v}$ to be close to v, and β controls how important $\mathcal{L}_{dist}$ is compared to $\mathcal{L}_{pred}$. We want $\bar{v}^*$ that minimizes Equation 1: this is the optimal CF example for v.

4. Method

We propose CF-GNNEXPLAINER, which generates $\bar{v} = (\bar{A}_v, x)$ given a node $v = (A_v, x)$. Our method can operate on any GNN model f. To illustrate our method and avoid cluttered notation, let f be a standard, one-layer Graph Convolutional Network [22] for node classification:

$$f(A, X; W) = \text{softmax}\left[\tilde{D}^{-1/2}\tilde{A}\tilde{D}^{-1/2}XW\right], \tag{2}$$

where $\tilde{A} = A + I$, I is the identity matrix, $\tilde{D}_{ii} = \sum_j \tilde{A}_{ij}$ are entries in the degree matrix $\tilde{D}$, X is the node feature

matrix, and W is the weight matrix [22].

4.1. Adjacency Matrix Perturbation

First, we define $\bar{A}_v = P \odot A_v$, where P is a binary perturbation matrix that sparsifies A_v. Our aim is to find P for a given node v such that $f(A_v, x) \neq f(P \odot A_v, x)$. To find P, we build upon the method by Srinivas et al. [35] for training sparse neural networks (see Section 2). Our objective is to zero out entries in the adjacency matrix (i.e., remove edges). That is, we want to find P that minimally perturbs A_v, and use it to compute $\bar{A}_v = P \odot A_v$. If an element $P_{i,j} = 0$, this results in the deletion of the edge between node i and node j. When P is a matrix of ones, this indicates that all edges in A_v are used in the forward pass.

Similar to Srinivas et al. [35], we first generate an intermediate, real-valued matrix $\hat{P}$ with entries in $[0, 1]$, apply a sigmoid transformation, then threshold the entries to arrive at a binary P: entries above 0.5 become 1, while those below 0.5 become 0. In the case of undirected graphs (i.e., those with symmetric adjacency matrices), instead of generating $\hat{P}$ directly, we first generate a perturbation vector which we then use to populate $\hat{P}$ in a symmetric manner.

4.2. Counterfactual Generating Model

We want our perturbation matrix P to only act on A_v, not $\tilde{A}_v$, in order to preserve self-loops in the message passing of f (i.e., we always want a node representation update to include the node's representation from the previous layer). To accommodate this, we first rewrite Equation 2 for our illustrative one-layer case to isolate A_v:

$$f(A_v, X_v; W) = \\ \text{softmax}\left[(D_v + I)^{-1/2}(A_v + I)(D_v + I)^{-1/2} X_v W\right] \quad (3)$$

To generate counterfactuals, we propose a new function g, which is based on f, but it is parameterized by P instead of by W. We update the degree matrix D_v based on $P \odot A_v$, add the identity matrix to account for self-loops (as in $\tilde{D}_v$ in Equation 2), and call this $\bar{D}_v$:

$$g(A_v, X_v, W; P) = \\ \text{softmax}\left[{\bar{D}_v}^{-1/2}(P \odot A_v + I){\bar{D}_v}^{-1/2} X_v W\right] \quad (4)$$

In other words, f learns the weight matrix while holding the data constant, while g is optimized to find a perturbation matrix that is then used to generate new data points (i.e., counterfactual examples) while holding the weight matrix constant. Another distinction between f and g is that the aim of f is to find the optimal set of weights that generalizes well on an unseen test set, while the objective of g is to generate an optimal counterfactual example, given a particular node (i.e., $\bar{v}$ is the output of g).

Algorithm 1 CF-GNNEXPLAINER: given a node $v = (A_v, x)$ where $f(v) = y$, generate the minimal perturbation, $\bar{v} = (\bar{A}_v, x)$, such that $f(\bar{v}) \neq y$.

Input: node $v = (x, A_v)$, trained GNN model f, CF model g, loss function $\mathcal{L}$, learning rate α, trade-off parameter β, number of iterations K, distance function d.

$f(v) = y$ *# Get GNN prediction*
$\hat{P} \leftarrow J_n$ *# Initialization*

for $k \in range(K)$ **do**
  $v^{(k)}$ = GET_CF_EXAMPLE()
  $\mathcal{L} \leftarrow \mathcal{L}(v, \bar{v}^{(k)})$ *# Eq 1 & Eq 5*
  $\hat{P} \leftarrow P^{(k)} + \alpha \nabla_{\hat{P}} \mathcal{L}$ *# Update* $\hat{P}$
end for

Function GET_CF_EXAMPLE()
  $P \leftarrow \text{threshold}(\sigma(\hat{P}^{(k)}))$
  $\bar{A}_v \leftarrow P \odot A_v$
  $\bar{v}^{(k)}_{cand} \leftarrow (\bar{A}_v, x)$
  if $f(v) \neq f(\bar{v}^{(k)}_{cand})$ **then**
    $\bar{v}^{(k)} \leftarrow \bar{v}^{(k)}_{cand}$
    if $\mathcal{L}_{dist}(v, \bar{v}) \leq \mathcal{L}_{dist}(v, \bar{v}^{(k)})$ **then**
      $\bar{v}^{*} \leftarrow \bar{v}^{(k)}$ *# Keep track of best CF*
    end if
  end if
  return $\bar{v}^{*}$

4.3. Loss Function Optimization

We generate P by minimizing Equation 1, adopting the negative log-likelihood (NLL) loss for $\mathcal{L}_{pred}$:

$$\mathcal{L}_{pred}(v, \bar{v} | f, g) = -\mathbb{1}\left[f(v) = f(\bar{v})\right] \cdot \mathcal{L}_{NLL}(f(v), g(\bar{v})). \quad (5)$$

Since we do not want $f(\bar{v})$ to match $f(v)$, we put a negative sign in front of $\mathcal{L}_{pred}$, and include an indicator function to ensure the loss is active as long as $f(\bar{v}) = f(v)$. Note that f and g have the same weight matrix W – the main difference is that g also includes the perturbation matrix P.

For $\mathcal{L}_{dist}$, we take d to be the element-wise difference between v and $\bar{v}$, corresponding to the difference between A_v and $\bar{A}_v$ – the number of edges removed. For undirected graphs, we divide this value by 2 to account for the symmetry in the adjacency matrices. When updating P, we take the gradient of Equation 1 with respect to the intermediate $\hat{P}$, *not* the binary P.

4.4. CF-GNNEXPLAINER

We call our method CF-GNNEXPLAINER and summarize its details in Algorithm 1: given an instance in the test

set v, we first obtain its original prediction from f and initialize $\hat{P}$ as a matrix of ones, J_n, to initially retain all edges. Next, we run CF-GNNEXPLAINER for K iterations. To find a CF example, we use Equation 4. First, we compute P by thresholding $\hat{P}$ (see Section 4.1). Then we use P to obtain the sparsified adjacency matrix that gives us a candidate CF example. This example is then fed to the original GNN, f, and if f predicts a different output than for the original node, we have found a valid CF example, $\bar{v}$. We keep track of the "best" CF example (i.e., the most minimal according to d), and return this as the optimal CF example $\bar{v}^*$ after K iterations. Between iterations, we compute the loss following Equations 1 and 5, and update $\hat{P}$ based on the gradient of the loss. In the end, we retrieve the optimal CF explanation $\Delta_v^* = v - \bar{v}^*$.

5. Experimental Setup

5.1. Datasets and Models

Given how challenging it is to define and evaluate accuracy of XAI techniques [8], we first focus on synthetic tasks where we know the ground-truth explanations. Although there exist real graph classification datasets with ground-truth explanations [7], we are not aware of any real node classification datasets with ground-truth explanations, which is the task we focus on in this paper. We believe that building such a dataset would be an excellent contribution, but is outside the scope of this paper.

In our experiments, we use the TREE-CYCLES, TREE-GRIDS, BA-SHAPES datasets from Ying et al. [46]. These are synthetic datasets that were created specifically for the task of explaining node classification predictions from GNNs. Each dataset consists of (i) a base graph, (ii) motifs that are attached to random nodes of the base graph, and (iii) additional edges that are randomly added to the overall graph. They are all undirected graphs. The classification task is to determine whether or not the nodes are part of the motif. The purpose of these datasets is to have a ground-truth for the "correctness" of an explanation: for nodes in the motifs, the explanation is the motif itself [27]. The dataset statistics are available in Table 1.

TREE-CYCLES consists of a binary tree base graph with cycle-shaped motifs, TREE-GRIDS also has a binary tree as its base graph, but with 3×3 grids as the motifs. For BA-SHAPES, the base graph is a Barabasi-Albert (BA) with house-shaped motifs, where each motif consists of 5 nodes (one for the top of the house, two in the middle, and two on the bottom). Here, there are four possible classes (not in motif, in motif: top, middle, bottom). We note that compared to the other two datasets, the BA-SHAPES dataset is much more densely connected – the node degree is more than twice as high as that of the TREE-CYCLES or TREE-GRID datasets, and the average number of nodes and edges in each node's computation graph is order(s) of magnitude larger. We use the same experimental setup (i.e., dataset splits, model architecture) as Ying et al. [46] to train a 3-layer GCN (hidden size = 20) for each task. Our GCNs have at least 87% accuracy on the test set.

Table 1: Dataset statistics.

	TREE CYCLES	TREE GRID	BA SHAPES
# classes	2	2	4
# nodes	871	1231	700
# edges	1950	3410	4100
Avg node degree	2.27	2.77	5.87
Avg # nodes in A_v	19.12	30.69	304.40
Avg # edges in A_v	18.99	33.94	1106.24

5.2. Baselines

To evaluate CF-GNNEXPLAINER, we compare against 4 baselines: RANDOM, ONLY-1HOP, RM-1HOP, and GNNEXPLAINER. The random perturbation is meant as a sanity check. We randomly initialize the entries of $\hat{P} \in [-1, 1]$ and apply the same sigmoid transformation and thresholding as described in Section 4.1. We repeat this K times and keep track of the most minimal perturbation resulting in a CF example. Next, we compare against baselines that are based on the 1-hop neighbourhood of v (i.e., its ego graph): ONLY-1HOP keeps all edges in the ego graph of v, while RM-1HOP removes all edges in the ego graph of v.

Our fourth baseline is based on GNNEXPLAINER by Ying et al. [46], which identifies the S most relevant edges for the prediction (i.e., the most relevant subgraph of size S). To generate CF explanations, we remove the subgraph generated by GNNEXPLAINER. We include this method in our experiments in order to have a baseline based on a prominent GNN XAI method, but we note that subgraph-retrieving methods like GNNEXPLAINER are not meant for generating CF explanations. Unlike our method, GNNEXPLAINER cannot automatically find a "minimal" subgraph and therefore requires the user to determine the number of edges to keep in advance (i.e., the value of S). As a result, we cannot evaluate how "minimal" its CF explanations are, but we can compare it against our method in terms of its ability to generate valid CF examples (*Fidelity*) and how accurate those CF examples are (*Accuracy*) (see Section 5.3). We perform a hyperparameter search over S and choose the setting that produces the most CF examples.

5.3. Metrics

We generate separate CF examples for each node in the graph, and evaluate in terms of four metrics: (i) *Fidelity*,

(ii) *Explanation Size*, (iii) *Sparsity*, and (iv) *Accuracy*.

Fidelity: is defined as the proportion of nodes where the original predictions match the prediction for the explanations [48]. Since we generate CF examples, we do not want the original prediction to match the prediction for the explanation, so we want a low value for *Fidelity*.

Explanation Size: is the number of removed edges. It corresponds to the $\mathcal{L}_{dist}$ term in Equation 1: the difference between the original A_v and the CF $\bar{A}_v$. Since we want to have *minimal* explanations, we want a small value for this metric.

Sparsity: measures the proportion of edges in A_v that are removed [48]. A value of 0 indicates all edges in A_v were removed. Since we want *minimal* explanations, we want a value close to 1.

Accuracy: is the proportion of explanations that are "correct". Following Ying et al. [46]; Luo et al. [27], we only compute accuracy for nodes that are originally predicted as being part of the motifs, since accuracy can only be computed on instances for which we know the ground truth explanations. Since we want *minimal* explanations, we consider an explanation to be correct if it exclusively involves edges that are inside the motifs (i.e., only removes edges that are within the motifs).

5.4. Hyperparameters, Complexity, and Resources

We experiment with different optimizers and hyperparameter values for the number of iterations K, the trade-off parameter β, the learning rate α, and the Nesterov momentum m (when applicable) and choose the setting that produces the most CF examples. We test the number of iterations $K \in \{100, 300, 500\}$, the trade-off parameter $\beta \in \{0.1, 0.5\}$, learning rate $\alpha \in \{0.005, 0.01, 0.1, 1\}$, and Nesterov momentum $m \in \{0, 0.5, 0.7, 0.9\}$. We test Adam, SGD and AdaDelta as optimizers. We find that for all three datasets, the SGD optimizer gives the best results, with $k = 500$, $\beta = 0.5$, and $\alpha = 0.1$. For the TREE-CYCLES and TREE-GRID datasets, we set $m = 0$, while for the BA-SHAPES dataset, we use $m = 0.9$.

CF-GNNEXPLAINER has time complexity $O(KN^2)$, where N is the number of nodes in the subgraph neighbourhood and K is the number of iterations. We note that high complexity is common for local XAI methods (i.e., SHAP, GNNExplainer, etc), but in practice, one typically only generates explanations for a subset of the dataset. We run approximately 375 hours of experiments on one Nvidia TitanX Pascal GPU with access to 12GB RAM.

6. Results

We evaluate CF-GNNEXPLAINER in terms of the metrics outlined in Section 5.3. The results are shown in Table 2. In almost all settings, we find that CF-GNNEXPLAINER outperforms the baselines in terms of *Explanation Size*, *Subgraph Impact*, and *Accuracy*, which shows that CF-GNNEXPLAINER satisfies our objective of finding accurate, minimal CF examples. In cases where other methods outperform CF-GNNEXPLAINER on a particular metric, these methods perform poorly on the rest of the metrics, or on other datasets.

Fidelity: CF-GNNEXPLAINER outperforms ONLY-1HOP and GNNEXPLAINER across all three datasets, and outperforms RM-1HOP for TREE-CYCLES and TREE-GRID in terms of *Fidelity*. We find that RANDOM has the lowest *Fidelity* in all cases – it is able to find CF examples for every single node. In the following subsections, we will see that this corresponds to poor performance on the other metrics.

Explanation Size: Figures 1 to 5 show histograms of the *Explanation Size* for the five methods we compare.[1] We see that across all three datasets, CF-GNNEXPLAINER has the smallest (i.e., most minimal) *Explanation Sizes*. This is especially true when comparing to RANDOM and ONLY-1HOP for the BA-SHAPES dataset, where we had to use a different scale for the x-axis due to how different the *Explanation Sizes* were. We postulate that this difference could be because BA-SHAPES is a much more densely connected graph; it has fewer nodes but more edges compared to the other two datasets, and the average number of nodes and edges in the subgraph neighbourhood is order(s) of magnitude larger (see Table 1). Therefore, when performing random perturbations, there is substantial opportunity to remove edges that do not necessarily need to be removed, leading to much larger *Explanation Sizes*. When there are many edges in the subgraph neighbourhood, removing everything except the 1-hop neighbourhood, as is done in ONLY-1HOP, also results in large *Explanation Sizes*. In contrast, the loss function used by CF-GNNEXPLAINER ensures that only a few edges are removed, which is the desirable behavior since we want minimal explanations.

Sparsity: CF-GNNEXPLAINER outperforms all four baselines for all three datasets in terms of *Sparsity*. We note CF-GNNEXPLAINER and RM-1HOP perform much better on this metric in comparison to the other methods, which aligns with the results from *Explanation Size*.

[1] We remind the reader that the *Explanation Size* for GNNEXPLAINER is equal to the user-input parameter S (see Section 5.2). It is possible for the *Explanation Size* to be less than S if and only if there are less than S edges in the subgraph neighbourhood of a node v. This is possible only for nodes outside of the motifs. However, based on Figure 4, we see that this is rarely the case, and that in almost 100% of cases, *Explanation Size* $= S$.

Table 2: Results comparing our method (denoted CF-GNN) to the baselines. Below each metric, ▼ indicates a low value is desirable, while ▲ indicates a high value is desirable.

Metric	TREE-CYCLES				TREE-GRID				BA-SHAPES			
	Fid. ▼	*Size* ▼	*Spars.* ▲	*Acc.* ▲	*Fid.* ▼	*Size* ▼	*Spars.* ▲	*Acc.* ▲	*Fid.* ▼	*Size* ▼	*Spars.* ▲	*Acc.* ▲
RANDOM	**0.00**	4.70	0.79	0.63	**0.00**	9.06	0.75	0.77	**0.00**	503.31	0.58	0.17
ONLY-1HOP	0.32	15.64	0.13	0.45	0.32	29.30	0.09	0.72	0.60	504.18	0.05	0.18
RM-1HOP	0.46	2.11	0.89	—	0.61	2.27	0.92	—	0.21	10.56	0.97	**0.99**
GNNEXP	0.55	6.00	0.57	0.46	0.34	8.00	0.68	0.74	0.81	6.00	0.81	0.27
CF-GNN	0.21	**2.09**	**0.90**	**0.94**	0.07	**1.47**	**0.94**	**0.96**	0.39	**2.39**	**0.99**	0.96

Accuracy: We observe that CF-GNNEXPLAINER has the highest *Accuracy* for the TREE-CYCLES and TREE-GRID datasets, whereas RM-1HOP has the highest *Accuracy* for BA-SHAPES. However, we are unable to calculate the accuracy of RM-1HOP for the other two datasets since it is unable to generate *any* CF examples for motif nodes, contributing to the low *Coverage* on those datasets. We observe *Accuracy* levels upwards of 94% for CF-GNNEXPLAINER across *all* datasets, indicating that it is consistent in correctly removes edges that are crucial for the initial predictions in the vast majority of cases (see Table 2).

7. Related Work

Several GNN XAI approaches have been proposed – a recent survey of the most relevant work is presented by Yuan et al. [48]. However, unlike our work, *none* of the methods in this survey generate CF explanations.

Kang et al. [19] generate CF examples for GNNs, but their work focuses on a different task: link prediction. Other GNN XAI methods identify important node features [17] or similar examples [12]. Yuan et al. [47] and Schnake et al. [34] generate model-level explanations for GNNs, which differs from our work since we produce instance-level explanations.

Adversarial attacks [38] are also related to CF examples: they both represent instances obtained from minimal perturbations to the input, which induce changes in the prediction made by the learned model. One difference between the two is in the intent: adversarial examples are meant to fool the model, while CF examples are meant to explain the prediction [26]. In the context of graph data, adversarial attack methods make minimal perturbations to the *overall graph* with the intention of degrading overall model performance. In contrast, we are interested in generating CF examples for individual nodes, as opposed to identifying perturbations to the overall graph.

There exists a substantial body of work on CF explanations for tabular, image, and text data [41; 20; 36]. Some methods treat the underlying classification model as a black-box [23; 14; 25], whereas others make use of the model's inner workings [39; 43; 40; 18; 26]. All of these methods are based on perturbing feature values to generate CF examples – they are not equipped to handle graph data with relationships (i.e., edges) between data points. In contrast, CF-GNNEXPLAINER provides CF examples for graph data.

8. Broader Impact

Concerns have been raised about the hidden assumptions behind CF examples [2] and potentials for misuse [21]. When explaining ML systems through CF examples, it is important to account for the context in which the system is deployed. We believe it is crucial for the ML community to invest in developing more rigorous evaluation protocols for XAI methods. We suggest that researchers in XAI collaborate with researchers in human-computer interaction to design human-centered user studies about evaluating the utility of XAI methods in practice. We are glad to see initiatives for such collaborations already taking place [10].

9. Conclusion

We propose CF-GNNEXPLAINER, which generates CF explanations for any GNN. Our simple and effective method is able to generate CF explanations that are (i) minimal, both in terms of the absolute number of edges removed, as well as the proportion of the subgraph neighbourhood that is perturbed, and (ii) accurate, in terms of removing edges that we know to be crucial for the initial predictions. We evaluate our method on three commonly used datasets for GNN explanation tasks and find that these results hold across all three datasets. For future work, we plan to incorporate node feature perturbations in our framework and extend CF-GNNEXPLAINER to accommodate graph classification tasks. We also plan to investigate the potential of adapting graph attack methods for generating CF explanations and conduct a user study to determine if humans find CF-GNNEXPLAINER useful in practice.

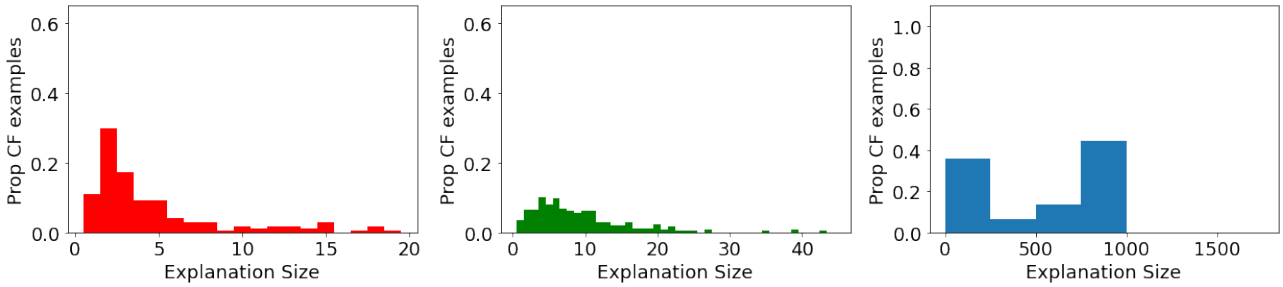

Figure 1: Histograms showing *Explanation Size* from RANDOM. Note the x-axis for BA-SHAPES goes up to 1500. Left: TREE-CYCLES, Middle: TREE-GRID, Right: BA-SHAPES.

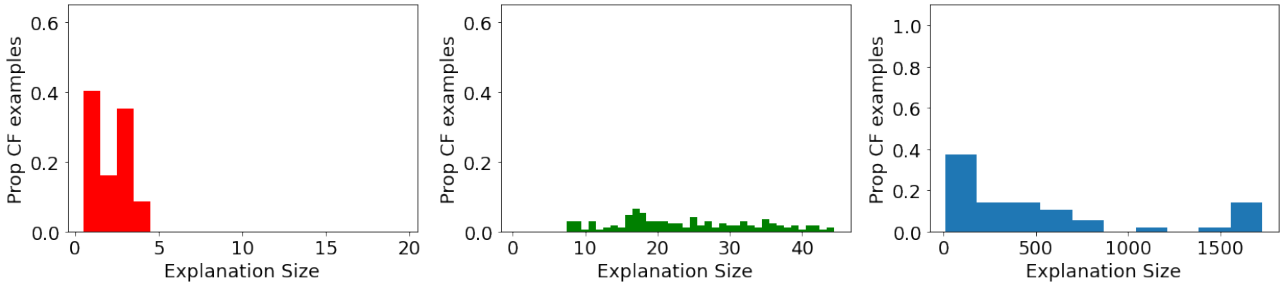

Figure 2: Histograms showing *Explanation Size* from ONLY-1HOP. Note the x-axis for BA-SHAPES goes up to 1500. Left: TREE-CYCLES, Middle: TREE-GRID, Right: BA-SHAPES.

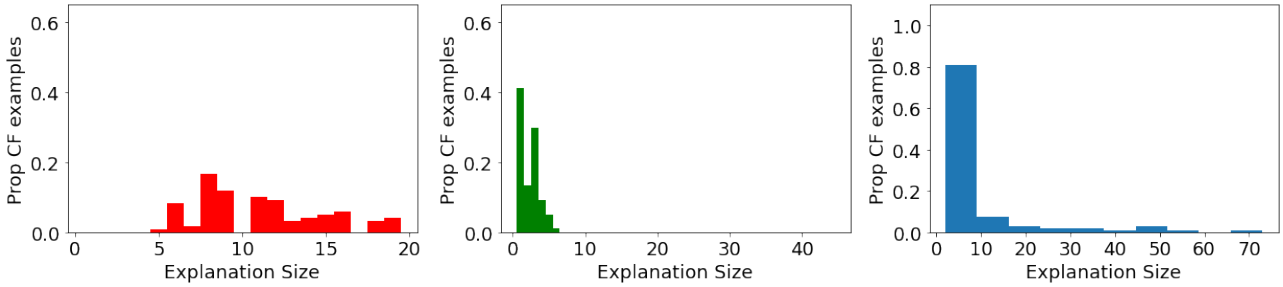

Figure 3: Histograms showing *Explanation Size* from RM-1HOP. Note the x-axis for BA-SHAPES goes up to 70. Left: TREE-CYCLES, Middle: TREE-GRID, Right: BA-SHAPES.

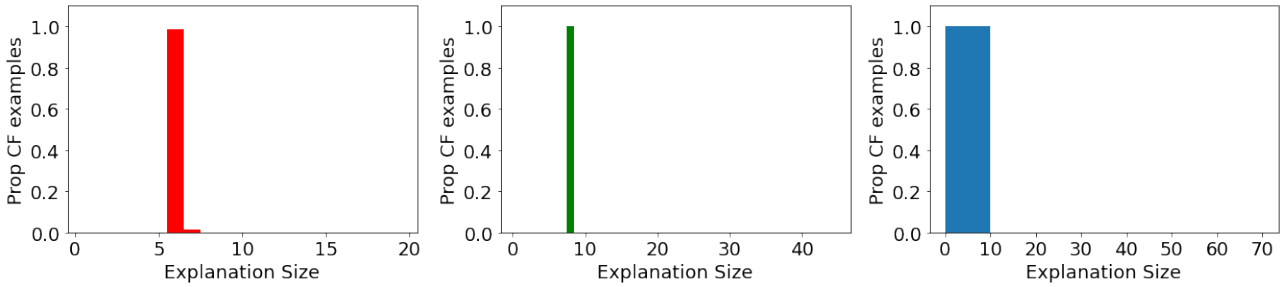

Figure 4: Histograms showing *Explanation Size* from GNNEXPLAINER. Note that the y-axis goes up to 1. Left: TREE-CYCLES, Middle: TREE-GRID, Right: BA-SHAPES.

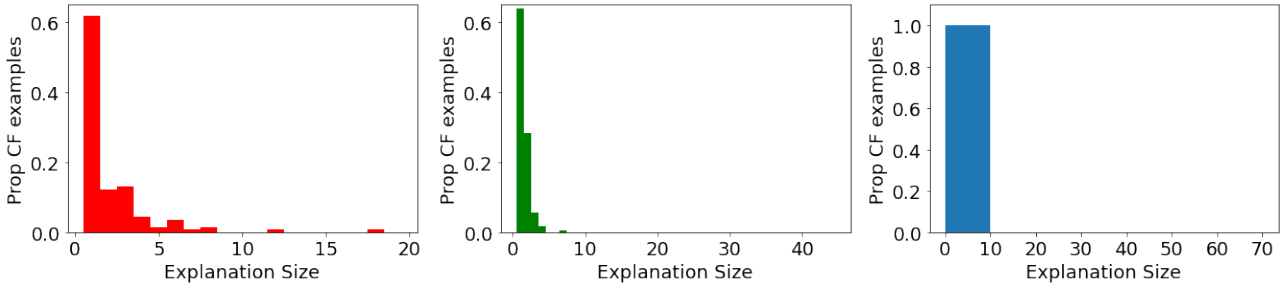

Figure 5: Histograms showing *Explanation Size* from CF-GNNEXPLAINER. Note the x-axis for BA-SHAPES goes up to 70. Left: TREE-CYCLES, Middle: TREE-GRID, Right: BA-SHAPES.

Acknowledgments

This research was supported by the Netherlands Organisation for Scientific Research (NWO) under project nr. 652.001.003, the Dutch Police, the Italian Ministry of Education, University and Research (MIUR) under the grant "Dipartimenti di eccellenza 2018-2022" of the Department of Computer Science and the Department of Computer Engineering at Sapienza University of Rome. All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

References

[1] Baldassarre, F. and Azizpour, H. Explainability techniques for graph convolutional networks. *arXiv preprint arXiv:1905.13686*, May 2019.

[2] Barocas, S., Selbst, A. D., and Raghavan, M. The hidden assumptions behind counterfactual explanations and principal reasons. In *ACM Conference on Fairness, Accountability, and Transparency*, December 2020.

[3] Battaglia, P. W., Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R., Gulcehre, C., Song, F., Ballard, A., Gilmer, J., Dahl, G., Vaswani, A., Allen, K., Nash, C., Langston, V., Dyer, C., Heess, N., Wierstra, D., Kohli, P., Botvinick, M., Vinyals, O., Li, Y., and Pascanu, R. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, October 2018.

[4] Bodria, F., Giannotti, F., Guidotti, R., Naretto, F., Pedreschi, D., and Rinzivillo, S. Benchmarking and survey of explanation methods for black box models, 2021.

[5] Chami, I., Abu-El-Haija, S., Perozzi, B., Re, C., and Murphy, K. Machine learning on graphs: A model and comprehensive taxonomy. *arXiv preprint arXiv:2005.03675*, 2021.

[6] Deac, A., Huang, Y.-H., Velickovic, P., Lio, P., and Tang, J. Drug-drug adverse effect prediction with graph co-attention. *arXiv preprint arXiv:1905.00534*, May 2019.

[7] Debnath, A., Lopez de Compadre, R., Debnath, G., Shusterman, A., and Hansch, C. Structure-activity relationship of mutagenic aromatic and heteroaromatic nitro compounds. correlation with molecular orbital energies and hydrophobicity. *Journal of Medicinal Chemistry*, 34:786–797, 1991.

[8] Doshi-Velez, F. and Kim, B. Towards a Rigorous Science of Interpretable Machine Learning. *arXiv preprint arXiv:1702.08608v2*, 2017.

[9] Duval, A. and Malliaros, F. D. Graphsvx: Shapley value explanations for graph neural networks. 2021.

[10] Ehsan, U., Liao, Q. V., and et al. Operationalizing Human-Centered Perspectives in Explainable AI, 2021.

[11] EU. Regulation (EU) 2016/679 of the European Parliament (GDPR). *Official Journal of the European Union*, L119:1–88, 2016.

[12] Faber, L., Moghaddam, A. K., and Wattenhofer, R. Contrastive graph neural network explanation. *ICML 2020 Workshop on Graph Representation Learning and Beyond*, pp. 6, 2020.

[13] Goebel, R., Chander, A., Holzinger, K., Lecue, F., Akata, Z., Stumpf, S., Kieseberg, P., and Holzinger, A. Explainable AI: The new 42? In *CD-Make 2018*, pp. 295–303, 2018.

[14] Guidotti, R., Monreale, A., Ruggieri, S., Pedreschi, D., Turini, F., and Giannotti, F. Local rule-based explanations of black box decision systems. *arXiv preprint arXiv:1805.10820*, May 2018.

[15] Guidotti, R., Monreale, A., Turini, F., Pedreschi, D., and Giannotti, F. A survey of methods for explaining black box models. *arXiv preprint arXiv:1802.01933*, 2018.

[16] Guo, Z., Zhang, C., Yu, W., Herr, J., Wiest, O., Jiang, M., and Chawla, N. V. Few-shot graph learning for molecular property prediction. In *Proceedings of The Web Conference*, 2021.

[17] Huang, Q., Yamada, M., Tian, Y., Singh, D., Yin, D., and Chang, Y. GraphLIME: Local interpretable model explanations for graph neural networks. *arXiv preprint arXiv:2001.06216*, January 2020.

[18] Kanamori, K., Takagi, T., Kobayashi, K., and Arimura, H. DACE: Distribution-aware counterfactual explanation by mixed-integer linear optimization. *IJCAI*, pp. 2855–2862.

[19] Kang, B., Lijffijt, J., and De Bie, T. ExplaiNE: An approach for explaining network embedding-based link predictions. *arXiv preprint arXiv:1904.12694*, 2019.

[20] Karimi, A.-H., Barthe, G., Scholkopf, B., and Valera, I. A survey of algorithmic recourse: Definitions, formulations, solutions, and prospects. *arXiv preprint arXiv:2010.04050*, 2020.

[21] Kasirzadeh, A. and Smart, A. The use and misuse of counterfactuals in ethical machine learning. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pp. 228–236, 2021.

[22] Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, February 2017.

[23] Laugel, T., Lesot, M.-J., Marsala, C., Renard, X., and Detyniecki, M. Inverse classification for comparison-based interpretability in machine learning. *arXiv preprint arXiv:1712.08443*, December 2017.

[24] Lin, W., Lan, H., and Li, B. Generative causal explanations for graph neural networks. *arXiv preprint arXiv:2104.06643*, April 2021.

[25] Lucic, A., Haned, H., and de Rijke, M. Why does my model fail? Contrastive local explanations for retail forecasting. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pp. 90–98, 2020.

[26] Lucic, A., Oosterhuis, H., Haned, H., and de Rijke, M. Focus: Flexible optimizable counterfactual explanations for tree ensembles. *arXiv preprint arXiv:1911.12199*, 2020.

[27] Luo, D., Cheng, W., Xu, D., Yu, W., Zong, B., Chen, H., and Zhang, X. Parameterized explainer for graph neural network. *NeurIPS*, 2020.

[28] Luo, D., Cheng, W., Xu, D., Yu, W., Zong, B., Chen, H., and Zhang, X. Parameterized explainer for graph neural network. *arXiv preprint arXiv:2011.04573*, November 2020.

[29] Miller, T. Explanation in artificial intelligence: Insights from the social sciences. *arXiv preprint arXiv:1706.07269*, 2017.

[30] Nguyen, C. Q., Kreatsoulas, C., and Branson, K. M. Meta-learning gnn initializations for low-resource molecular property prediction. In *ICML 2020 Workshop on Graph Representation Learning and Beyond (GRL+)*, 2020.

[31] Pope, P. E., Kolouri, S., Rostami, M., Martin, C. E., and Hoffmann, H. Explainability methods for graph convolutional neural networks. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10764–10773, Long Beach, CA, USA, June 2019. IEEE.

[32] Samek, W., Montavon, G., Vedaldi, A., Hansen, L. K., and Müller, K.-R. *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*. Springer, 2019.

[33] Schlichtkrull, M. S., De Cao, N., and Titov, I. Interpreting graph neural networks for NLP with differentiable edge masking. *arXiv preprint arXiv:2010.00577*, October 2020.

[34] Schnake, T., Eberle, O., Lederer, J., Nakajima, S., Schutt, K. T., Muller, K.-R., and Montavon, G. XAI for graphs: Explaining graph neural network predictions by identifying relevant walks. *arXiv preprint arXiv:2006.03589*, June 2020.

[35] Srinivas, S., Subramanya, A., and Babu, R. V. Training sparse neural networks. *arXiv preprint arXiv:1611.06694*, November 2016.

[36] Stepin, I., Alonso, J. M., Catala, A., and Pereira-Fariña, M. A survey of contrastive and counterfactual explanation generation methods for explainable artificial intelligence. *IEEE Access*, 9:11974–12001, 2021.

[37] Stokes, J. M., Yang, K., Swanson, K., Jin, W., Cubillos-Ruiz, A., Donghia, N. M., MacNair, C. R., French, S., Carfrae, L. A., Bloom-Ackermann, Z., Tran, V. M., Chiappino-Pepe, A., Badran, A. H., Andrews, I. W., Chory, E. J., Church, G. M., Brown, E. D., Jaakkola, T. S., Barzilay, R., and Collins, J. J. A deep learning approach to antibiotic discovery. *Cell*, 180:688–702, 2020.

[38] Sun, L., Dou, Y., Yang, C., Wang, J., Yu, P. S., He, L., and Li, B. Adversarial attack and defense on graph data: A survey. *arXiv preprint arXiv:1812.10528*, 2018.

[39] Tolomei, G., Silvestri, F., Haines, A., and Lalmas, M. Interpretable predictions of tree-based ensembles via actionable feature tweaking. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 465–474, 2017.

[40] Ustun, B., Spangher, A., and Liu, Y. Actionable recourse in linear classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pp. 10–19, 2019.

[41] Verma, S., Dickerson, J., and Hines, K. Counterfactual explanations for machine learning: A review. *arXiv preprint arXiv:2010.10596*, 2020.

[42] Vu, M. N. and Thai, M. T. PGM-Explainer: Probabilistic graphical model explanations for graph neural networks. 2020.

[43] Wachter, S., Mittelstadt, B., and Russell, C. Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harvard Journal of Law & Technology*, 31(2):841–888, 2018.

[44] Wieder, O., Kohlbacher, S., Kuenemann, M., Garon, A., Ducrot, P., Seidel, T., and Langer, T. A compact review of molecular property prediction with graph neural networks. *Drug Discovery Today: Technologies*, December 2020.

[45] Xie, Y., Shi, C., Zhou, H., Yang, Y., Zhang, W., Yu, Y., and Li, L. Mars: Markov molecular sampling for multi-objective drug discovery. In *Proceedings of the International Conference on Learning Representations*, 2021.

[46] Ying, R., Bourgeois, D., You, J., Zitnik, M., and Leskovec, J. GNNExplainer: Generating explanations for graph neural networks. *arXiv preprint arXiv:1903.03894*, November 2019.

[47] Yuan, H., Tang, J., Hu, X., and Ji, S. XGNN: Towards model-level explanations of graph neural networks. *arXiv preprint arXiv:2006.02587*, June 2020.

[48] Yuan, H., Yu, H., Gui, S., and Ji, S. Explainability in graph neural networks: A taxonomic survey. *arXiv preprint arXiv:2012.15445*, 2020.

[49] Yuan, H., Yu, H., Wang, J., Li, K., and Ji, S. On explainability of graph neural networks via subgraph explorations. *arXiv preprint arXiv:2102.05152*, February 2021.

[50] Zitnik, M., Agrawal, M., and Leskovec, J. Modeling polypharmacy side effects with graph convolutional networks. *Bioinformatics*, 34(13):i457–i466, July 2018.