COMRECGC: Global Graph Counterfactual Explainer through Common Recourse

Gregoire Fournier¹ **Sourav Medya**²

Abstract

Graph neural networks (GNNs) have been widely used in various domains such as social networks, molecular biology, or recommendation systems. Concurrently, different explanations methods of GNNs have arisen to complement its black-box nature. Explanations of the GNNs' predictions can be categorized into two types-factual and counterfactual. Given a GNN trained on binary classification into "accept" and "reject" classes, a global counterfactual explanation consists in generating a small set of "accept" graphs relevant to all of the input "reject" graphs. The transformation of a "reject" graph into an "accept" graph is called a recourse. A common recourse explanation is a small set of recourse, from which every "reject" graph can be turned into an "accept" graph. Although local counterfactual explanations have been studied extensively, the problem of finding common recourse for global counterfactual explanation remains unexplored, particularly for GNNs. In this paper, we formalize the common recourse explanation problem, and design an effective algorithm, COMRECGC, to solve it. We benchmark our algorithm against strong baselines on four different real-world graphs datasets and demonstrate the superior performance of COMRECGC against the competitors. We also compare the common recourse explanations to the graph counterfactual explanation, showing that common recourse explanations are either comparable or superior, making them worth considering for applications such as drug discovery or computational biology.

1. Introduction

Graph Neural Networks (GNNs) have been widely used for graph classification across various domains such as electrical design (Mirhoseini et al., 2020), physical simulation (Bhattoo et al., 2022), or drug discovery (Jiang et al., 2020; Cao et al., 2016; Hamilton et al., 2017). Despite their popularity and their good performance, the predictions of GNNs are not yet fully understood, and explaining GNN's behavior has become a central focus of recent research efforts (Kakkad et al., 2023) (Yuan et al., 2022). Our work focuses on the concept of Counterfactual Explanation (CE) in the context of binary classification (e.g., 'reject' vs. 'accept' classes). This involves designing the minimal changes required in each 'reject', or undesirable graph, to be able to change to an 'accept' graph, or counterfactual graph, that is often similar in structure. A global CE, in particular, aims to identify a small set of 'accept' graphs relevant to all 'reject' graphs, thereby providing insights into the critical decision regions recognized by the base GNN model. This global explanation approach highlights the GNN's model-level behavior as opposed to an instance level or local one.

Another important aspect of CE explanations is the implicit provision of recourse—graph transformations that convert a 'reject' graph into an 'accept' graph (Verma et al., 2022; Karimi et al., 2020). This is particularly valuable in applications where it is possible to take action to reverse a decision, such as in loan applications or in modifying a compound to change physical or chemical properties of molecules (e.g., mutagenicity (Riesen & Bunke, 2008)).

The instance-specific or local counterfactual explanations are insightful and capable of providing recourse for each instance (Abrate & Bonchi, 2021; Bajaj et al., 2021; Lucic et al., 2022; Tanyel et al., 2023). However, the explanations generated are generally large in number and not straightforward to interpret at the model level. Global counterfactual explainers have been proposed to address these issues (Huang et al., 2023; Ley et al., 2023). We present additional related work in Appendix A. In (Huang et al., 2023), the authors build global counterfactuals for GNNs. However, a major limitation of their approach is its generalizability when considering recourse. A single counterfactual can lead to significantly different recourse depending on the specific graph it explains, making it unsuitable as a local explanation. In (Lev et al., 2023), a list of recourse directions is determined by experts, which the algorithm then uses to create a global explanation. This approach has two main

¹Department of Mathematics and Computer Science, University of Illinois Chicago, Chicago, USA ²Department of Computer Science, University of Illinois Chicago, Chicago, USA.

weaknesses—first, it relies on experts to provide a comprehensive set of recourse directions; second, it separates the process of finding recourse directions from that of fitting them to the data, which may compromise robustness and optimality. Moreover, the latter method has not been utilized for graph data.

To answer these shortcomings, our work considers the problem of *finding common recourse* (FCR) explanation, which seeks to generate a small set of recourse that is capable of converting every 'reject' instance into an 'accept'. This essentially means that one or more graphs use the same recourse to achieve a counterfactual. To the best of our knowledge, this problem has not been studied before from a graph machine learning perspective. FCR offers both highlevel interpretability and instance level insights. Moreover, in some cases, more input graphs can be explained with fewer common recourse compared to global counterfactuals. Our main contributions are the following:

- Novel problem setting. We formalize the FCR problem. We prove that the FCR problem is NP-hard. We provide a generalized version of the FCR problem, called FC (Finding Counterfactual) and derive an approximation algorithm to a constrained version of FC.
- **COMRECGC**. We design the COMRECGC algorithm to extract high quality common recourse as an explanation, which provides a solution to both problems.
- Experiments. We experiment on real world datasets and benchmark our algorithm against popular counterfactual explainers on the Common recourse task. COMRECGC outperforms the next best solution on the FCR problem by more than 20%, 30%, 20%, and 10% in total coverage on the NCI1, MUTAGENICITY, AIDS, and PROTEINS datasets respectively. We show that COMRECGC common recourse explanation offer a comparable and sometimes higher coverage than the best baseline global counterfactual explanation.

2. Background and Problem Formulation

Counterfactual Explanation. A graph is defined as G = (V, E), where V is the set of vertices and E the set of edges. Consider a binary graph classification model—such as a GNN—where the prediction function is be denoted by Φ . A graph is predicted in the 'reject' class if $\Phi(G) = 0$ and in the 'accept' class if $\Phi(G) = 1$. For any reject graph G, we call a graph H counterfactual if $\Phi(H) = 1$ and $d(G, H) \leq \theta$, where θ is a given normalized distance value such that $0 < \theta < 1$ and d is a distance function. In other words, we say H covers G. In the rest of the paper, we consider our input to be solely constituted of reject or 'undesirable' graphs, denoted by \mathbb{G} , and we denote by $\mathcal{C}(G)$ the set of counterfactual graphs that cover G. **Common Recourse.** Given a reject graph G covered by H, we call *recourse* a transformation (a function, r) which converts G into H. Note that the cost of the recourse is the distance between G and H. While defining a common recourse, an obvious question is how to define it on different graphs. For instance, if the recourse is 'adding' a vertex connected to some specific vertices that do not exist in different graphs, it will lead to some sort of ambiguity. However, a GNN always gives a latent representation (embedding). Thus, we use graph embedding z from the space of graphs into \mathbb{R}^l for some constant l to define common recourse. Note that these embeddings also can be used to measure the distance between two graphs such as graph edit distance (Ranjan et al., 2023).

Suppose *H* covers *G*, and let *r* be the recourse that turns *G* into *H* (i.e., r(G) = H), then we define the embedding of *r* in the space of vectors of \mathbb{R}^l to be \vec{r} , the vector from z(G) to z(H), i.e., $\overline{z(G)z(H)}$. Given two recourse r_1, r_2 , we say that they form a common recourse if there exists \vec{v} , a vector in \mathbb{R}^l , such that $||\vec{v} - \vec{r_1}||_2 \leq \Delta$ and $||\vec{v} - \vec{r_2}||_2 \leq \Delta$, for a fixed $0 < \Delta < 1$. The idea is that there exists a center recourse which is close to both recourse and can be used as a summary.

Please refer to the Appendix for the proofs of the claims and Theorems.

2.1. The Problem of Common Recourse

Our goal is to have a small representative set of common recourse, \mathbb{F} , and the cost of these should be small. Before formalizing our problem, we define the followings:

- COVERAGE(\mathbb{F}) := $|\{G \in \mathbb{G} | \exists H \in \mathcal{C}(G), \exists r \in \mathbb{F} \\$ such that $||\vec{r} \vec{z(G)z(H)}||_2 \leq \Delta\}|/|\mathbb{G}|$; the fraction of the input graphs for which at least one counterfactual is obtained through one of the recourse in \mathbb{F} . Note that the center recourse are being chosen from \mathbb{F} .
- $\operatorname{cost}(\mathbb{F}) := \operatorname{AGG}_{G \in \mathbb{G}} \{ \underline{\min}\{ ||\vec{r}||_2, r \in \mathbb{F} \text{ and } \exists H \in \mathcal{C}(G) \text{ such that } ||\vec{r} \overline{z(G)z(H)}||_2 \leq \Delta \} \};$ the total distance from the covered input graphs to their closest attained counterfactual through \mathbb{F} . The AGG function used in the experiments is the summation.
- size(\mathbb{F}) := $|\mathbb{F}|$.

I

Problem 1 (Finding Common Recourse (FCR)). Given input graphs \mathbb{G} and a budget R, the goal is to find a set of recourse of size R that maximizes the coverage:

$$max_{\mathbb{F}} \operatorname{coverage}(\mathbb{F})$$
 such that $\operatorname{size}(\mathbb{F}) \leq R$.

Recourse are derived from the nearby counterfactual graphs. For a set of counterfactuals \mathbb{H} , we define its associated

recourse set as $\mathbb{F}_{\mathbb{H}} := \{\overline{z(G)z(H)} | G \in \mathbb{G}, H \in \mathbb{H} \cap \mathcal{C}(G)\}$. Subsequently, another way to formalize the problem would be to look at the common recourse that can be extracted from a set of counterfactuals.

Problem 2 (Finding Counterfactual (FC)). Given input graph \mathbb{G} and two budgets R and T, the goal is to find a set of counterfactuals \mathbb{H} such that its associated common recourse set has size R:

$$max_{\mathbb{H}} \operatorname{coverage}(\mathbb{F}^*_{\mathbb{H}}) \text{ s.t. } \operatorname{size}(\mathbb{F}^*_{\mathbb{H}}) \leq R, size(\mathbb{H}) \leq T,$$

where $\mathbb{F}_{\mathbb{H}}^*$ denotes the common recourse set that achieves the maximum coverage from selecting *R* recourse of $F_{\mathbb{H}}$.

Note that the FCR problem is a specific case of the FC problem by setting T to $|\mathbb{G}|$, the number of input graphs to cover. The FC problem corresponds to a max multi-budget multi-cover problem, where we have separate budgets for counterfactuals and for recourse. To count an input graph as covered, it must be covered by both of these budgets.

Although cost is not a constraint nor an objective in both the FCR and FC problems, it remains, alongside coverage, a valuable metric for assessing the relevance of counterfactual explanations (Rawal & Lakkaraju, 2020), and by extension, the recourse explanations derived from them.

2.2. Analysis: The FCR problem

The FCR problem (Problem 1) consists in finding the best common recourse set of size R in terms of coverage from a list of common recourse. Let us define f as the function that associates to a set of common recourse the maximum coverage obtained by selecting R of them.

Theorem 1. The FCR problem is NP-hard. (Appendix B.1)

We claim that f is monotone and submodular (Appendix B.2), hence the common greedy algorithm yields a (1-1/e) approximation, which is the best poly-time approximation unless P=NP (Feige, 1998).

2.3. Analysis of the FC problem

The FC problem (see Problem 2) is an extension of the FCR problem where we are no longer given the full common recourse set, and we are required to pick at most T counterfactuals to form R recourse for best coverage. In Appendix B.3, we present a budget version of the FC problem.

Let g be the function that associates to a set of counterfactuals, its best coverage through building common recourse and selecting the R best for maximum coverage. It is not hard to see that g is monotone, and does not possess any local submodularity ratio (Santiago & Yoshida, 2020).

Theorem 2. g is not pseudo-modular. (Appendix B.4)

2.3.1. APPROXIMATION FOR THE FC PROBLEM

Since the function g is monotone but not-pseudo modular (Theorem 2), finding an approximation algorithm with some guarantees is non trivial. To make the problem tractable for approximation, we add the following constraint:

Constraint 1: (C1) To be considered as valid solution, a set of T counterfactuals, $G_1 \dots G_T$ must verify:

$$\forall p < T, \ g(\bigcup_{i \le p+1} G_i) - g(\bigcup_{i \le p} G_i) > 0.$$
 (1)

Intuitively, this means that there exists a "series" (in the union of sets sense) of counterfactual that increases the objective function at each step. Such a constraint does not seem too abstract, as most approximation algorithms look to add one element at a time. We obtain the following:

Theorem 3. There is a $1 - e^{-1/R}$ approximation algorithm in expectation for the FC problem with C1. (Appendix B.5)

Discussions. The above theoretical results have important implications: (1) FCR problem: The (1 - 1/e) approximation guarantee of the greedy algorithm is for selecting the best common recourse once the counterfactuals have been found. It is featured in our method for solving the FCR problem. (2) FC problem: The $1 - e^{-1/R}$ approximation for the FC problem might be less useful in practice, as typically we want to allow numerous ($R \approx 100$) recourse to explain a dataset. Though the corresponding randomized greedy algorithm provides the guarantee, instead of using this algorithm, we will assess the importance of each counterfactual with its number of visits through a "common recourse random walk".

3. Our Method: COMRECGC

To find good counterfactual for common recourse, COM-RECGC operates in different stages.

- A graph embedding algorithm: First, COMRECGC maps each graph from the set of input graphs into ℝ^l, an embedding space. From the representation of an input graph and one of its counterfactual, we derive the recourse embedding.
- COMRECGC finds counterfactuals for common recourse through a **multi-head vertex reinforced random walk** in the graph edit space. This is a variation of a vertex reinforced random walk (Pemantle, 2004).
- A clustering algorithm for common recourse: we form clusters over the embeddings of the generated recourse. Each cluster of a certain radius (Δ) corresponds to a *common recourse*. Finally, we aggregate the common recourse greedily to obtain the maximum coverage.

3.1. Graph embedding algorithm

One of the essential notions to define counterfactual is the distance function as described earlier. Our method COM-RECGC begins with the efficient computation of distance (or similarity) between two graphs. To assess distance between graphs, we use the graph edit distance (GED) (Sanfeliu & Fu, 1983) that accounts for the minimal number of transformations—such as edge/vertex deletion/addition and label change—to make the graphs isomorphic. For a graph G, we denote $\mathcal{N}(G)$ as the neighbor set of graphs that are only one edit distance $\widehat{\text{GED}}(G_1, G_2) = \text{GED}(G_1, G_2)/(|V_1| + |V_2| + |E_1| + |E_2|)$ for our framework, which has the advantage to compare two graphs of different sizes.

Since the GED is NP-hard to compute, we use a graph embedding algorithm GREED from (Ranjan et al., 2023) as a proxy, which aims to learn a GED metric through a projection of the graph dataset in \mathbb{R}^l . Note that it is possible to employ any other embedding algorithm that can estimate the GED between two graphs, which allows us to validate the closeness of counterfactuals, and to define common recourse on the space of graph through vectors in \mathbb{R}^l . We denote as z(G) the embedding of a graph G in the rest of the paper.

3.2. Multi-head vertex reinforced random walk

To identify counterfactuals, we explore the graph edit map through a random walk. The random walk occurs in the space of graphs, where each state is a distinct (nonisomorphic) graph. Two states are connected if and only if their corresponding graphs can be transformed into one another by a single edit.

COMRECGC uses a variation of a vertex reinforced random walk (VRRW). VRRW (Pemantle, 2004) performs random walks on a finite state space where the transition probability depends on the number of visits. This family of random walk has the advantage of experimentally converging, as well as providing an interpretation of the diversity and exploration performances (Mei et al., 2010). We now describe how COMRECGC uses random walks:

Initially, the random walk begins with k-heads each placed on different input graphs randomly, $(G_i)_{i \le k} \subset \mathbb{G}$. At each step, we either continue the walk or, with a small probability, all heads teleport back to starting graphs. It is crucial for our walk to keep track of the graph from which each head started or was teleported to. We represent the state of the k-heads random walk as $(u_i)_{i \le k}$.

In each step, we randomly select one of the heads as the **lead**, denoted by the index ℓ , then proceed as follows:

• First, we move the lead head. The goal for the walker

is to go towards a potential counterfactual graph with the following transition rule, for $v \in \mathcal{N}(u_{\ell})$:

$$p(u_{\ell}, v) \propto p_{\phi}(v)C(v) \tag{2}$$

Where C(v) is one plus the number of visits to graph v, and $p_{\phi}(v)$ is the probability, assigned by the GNN, of v being a counterfactual.

 Each non-lead head moves to the next state based on the recourse available among its neighbors that is closest to the lead's recourse. More formally, if the lead head is in state u_l after the previous step, then the next state for the *i*-th head is:

$$argmin_{v \in \mathcal{N}(u_i) \cup \{u_i\}} || \overrightarrow{z(G_\ell)z(u_\ell)} - \overrightarrow{z(G_i)z(v)} ||_2$$

Teleportation. The graph edit space is exponentially large. Thus, to explore the search space around small neighborhoods of the input graphs, we restrict the random walk by adding a probability of returning to input graphs of $0 < \tau < 1$ at each step. We call this operation teleportation, and each head state is reset to one of the input graphs optimizing for **coverage** as follows: define t(G) be the number of walks started from the input graph G. Then the probability of a head to teleport to $G \in \mathbb{G}$ is:

$$p_{\tau}(G) = \frac{\exp(-t(G))}{\sum_{G' \in \mathbb{G}} \exp(-t(G'))}$$
(3)

The pseudo code of the random walk procedure and some elements of analysis are presented in the Appendix (C.1 and C.2).

Counterfactual Candidates. After M steps, the random walk is terminated. For the next phase, we select the counterfactuals that have been visited at least a certain number of times, referring to these as the counterfactual candidates that represent the outcomes of the walk.

3.3. Aggregation with clustering

Once counterfactuals have been found, COMRECGC forms clusters among the resulting recourse to find common recourse. For this task, we use a spatial clustering, DBSCAN (Ester et al., 1996) to find high density areas to aggregate recourse with close embedding representation. We associate a common recourse to each cluster of radius Δ . Finally we pick *R* common recourse to be our explanation using a greedy approach. The algorithm on the set \mathbb{G} of input graphs, with the set of counterfactuals \mathbb{S} is given below. Here **gain** $(r; \mathbb{F})$ is the gain in coverage from adding the recourse *r* to \mathbb{F} . Algorithm 1 CR CLUSTERING(\mathbb{G} , \mathbb{S} , R)

1: $\mathbb{C} \leftarrow \Delta$ -clusterize $\{\overline{z(G)z(v)} : v \in \mathbb{S}, G \in \mathbb{G}\}$ 2: $\mathbb{F} \leftarrow \emptyset$ 3: for $i \in 1 : R$ do 4: $r \leftarrow \arg \max_{r \in \mathbb{C}} \operatorname{gain}(r; \mathbb{F})$ 5: $\mathbb{F} \leftarrow \mathbb{F} + \{r\}$ 6: end for 7: return \mathbb{F}

3.4. COMRECGC: algorithm and complexity

Finally we present COMRECGC's solution to the FCR problem, which consists in a Multi-head VRRW, a selection of counterfactual candidates for recourse, and a clustering process for identifying common recourse. The final COM-RECGC algorithm is given below. Line 2 is the filtering based on the number of visits, where n is a threshold set as a hyper-parameter for the dataset.

Algorithm 2 COMRECGC(ϕ , \mathbb{G} , k , M , τ , R , n)	
1: $\mathbb{S} \leftarrow \text{Multi-head VRRW}(\phi, \mathbb{G}, k, M, \tau)$	
2: $\mathbb{S} \leftarrow \text{top } n \text{ frequently visited counterfactuals in } \mathbb{S}$	
3: return CR CLUSTERING($\mathbb{G}, \mathbb{S}, R$)	

Complexity Analysis. The complexity of the random walk is O(Mhk), where M is the number of steps in the VRRW, k is the number of heads in the VRRW, and h is the maximum node degree in the graph-edit space map. The complexity of using DBSCAN clustering on $n|\mathbb{G}|$ recourse is $O(n^2|\mathbb{G}|^2)$, where n is the number of top-visited counterfactuals, and $|\mathbb{G}|$ is the number of input graphs. In practice, we have a small constant due to being in \mathbb{R}^l for l = 64. Finally, the complexity of the greedy summary of $n|\mathbb{G}|$ recourse over $|\mathbb{G}|$ features R times is $O(nR|\mathbb{G}|^2)$, where Ris the size of the recourse set. The overall complexity is $O(Mhk + n^2|\mathbb{G}|^2)$.

3.5. Variation of COMRECGC for the FC problem

Our goal is to solve a generic framework to solve both FCR and FC problem. Note that we can also use COMRECGC to generate solution to the FC problem. In the counterfactual candidates step, we need to significantly reduce the number of counterfactuals that we use to generate common recourse to match the constraints in FC. COMRECGC selects the counterfactuals that are closest to the input graphs, intuitively this is because the common recourse threshold Δ is geometrically tight, so we want recourse (cost) to be as small as possible to increase the chance of getting common ones in the clustering step. Algorithm 3 represents COM-RECGC's solution to the FC problem. Line 3 is the specific to the FC framework. In particular, we select $T = |\mathbb{G}|$ counterfactuals. This allows our method to be compared to the existing local explainers where one counterfactual is given per input graph (Ying et al., 2019). The complexity of this extra step is $O(n|\mathbb{G}|)$, hence COMRECGC's solution to the FC problem has a complexity of $O(Mhk + n^2|\mathbb{G}|^2)$.

Algorithm 3 COMRECGC for FC (ϕ , \mathbb{G} , k , M , R)	
1: $\mathbb{S} \leftarrow \text{Multi-head VRRW}(\phi, \mathbb{G}, k, M)$	
2: $\mathbb{S} \leftarrow \text{top } n$ frequently visited counterfactuals in \mathbb{S}	
3: $\mathbb{S} \leftarrow \bigcup_{G \in \mathbb{G}} \operatorname{argmin}_{v \in \mathbb{S}} \overline{z(G)z(v)} _2$	
4: return CR CLUSTERING($\mathbb{G}, \mathbb{S}, R$)	

4. Experiments

We evaluate COMRECGC on four real-world datasets against the recourse from benchmark explainers, and show:

- COMRECGC produces global common recourse that are of higher quality than those of the baselines on both the FCR and FC problems.
- The explanations produced by COMRECGC are significantly less costly in terms of recourse than the ones generated by the baseline counterfactual explainers.
- The common recourse from COMRECGC can explain a similar, and in some cases higher, number of input graphs compared to the counterfactual graphs generated by a global counterfactual explainer.

Reproducibility. We have made our code available at this anonymous link: https://anonymous.4open.science/r/COMRECGC-3E4E/.

4.1. Datasets

We consider the datasets MUTAGENICITY (Riesen & Bunke, 2008; Kazius et al., 2005), NCI1 (Wale & Karypis, 2006), AIDS (Riesen & Bunke, 2008), and PROTEINS (Borgwardt et al., 2005; Dobson & Doig, 2003). In the first three, each graph accounts for a molecule, where nodes represent atoms and edges chemical bonds between them. The molecules are classified by whether they are mutagenic, anticancer, and active against HIV, respectively. The PROTEINS dataset is composed of different proteins classified into enzymes and non-enzymes, with nodes representing secondary structure elements and edges representing structural proximity. For each dataset, we remove graphs containing rare nodes (with a label count less than 50). The statistics of the datasets are detailed in Table 1.

4.2. Experimental set up

We train a base GNN model (GCN) (Kipf & Welling, 2017) for a binary classification task, consisting of three convolutional layers, a max pooling layer, and a fully connected layer, following best practices from the literature (Vu &

Table 1. Datasets overview.								
	NCI1	MUTAGENICITY	AIDS	PROTEINS				
#Graphs	3978	4308	1837	1113				
#Nodes	118714	130719	28905	43471				
#Edges	128663	132707	29985	81044				
#Node Labels	10	10	9	3				

Thai, 2020). The model is trained with the Adam optimizer (Kingma & Ba, 2014) and a learning rate of 0.001 for 1000 epochs. The training/validation/testing split is 80%/10%/10%, and the corresponding accuracy measures are presented in Table 2.

Table 2. Accuracy of GCN on the graph binary classification task on NCI1, MUTAGENICITY, AIDS and PROTEINS.

	NCI1	MUTAGENICITY	AIDS	PROTEINS
Training	0.844	0.882	0.998	0.780
Validation	0.816	0.830	0.973	0.820
Testing	0.781	0.800	0.978	0.730

Choice of the parameters θ and Δ . The parameter θ accounts for the distance to a graph $G \in \mathbb{G}$ within which a graph with an accepted prediction can be used as a counterfactual explanation for G. There is a trade-off: a small θ provides instance-wise explanations, if they exist, and a larger θ allow for more effective summaries, as counterfactuals tend to cover a broader range of input graphs. Following (Huang et al., 2023), we set $\theta = 0.1$ in the experiments on the NCI1, MUTAGENICITY and AIDS datasets. However, we set $\theta = 0.15$ for the PROTEINS dataset, as we have found experimentally that counterfactuals tend to be more distant from their input graphs.

On the other hand, the parameter Δ accounts for the maximum difference between two recourse to be considered common. A larger Δ yields better summaries but the insights become less precise. We set $\Delta = 0.02$, although this may seem large compared to the value of θ , the embedding space dimension used for the datasets is l = 64, making the margin for common recourse sensibly smaller than the one for finding counterfactuals, as expected. An example illustrating a common recourse explanation with those parameters is provided in Figure 1.

COMRECGC parameters. Across all of our experiments, COMRECGC uses k = 5 heads, has probability of teleportation $\tau = 0.05$, performs the random walk for M = 50000steps and selects R = 100 common recourse.

4.3. Results for the FCR problem

Baselines. Since the FCR problem does not constrain the number of counterfactuals used to generate common recourse, explainers used for benchmarking this problem must

be able to generate a large number of counterfactuals. In the literature, we find that only GCFEXPLAINER(Huang et al., 2023) aims at generating global counterfactuals on a large scale. This is also the only global method to generate counterfactual. To generate common recourse with GCFEX-PLAINER counterfactuals, we use the same recourse clustering algorithm (Algorithm 1) that is part of COMRECGC. For COMRECGC, its number of heads variants, and GCFEX-PLAINER, we limit the number of counterfactuals entering the clustering stage to 100,000 graphs.

Results. The results are presented in Table 3. We observe that the common recourse explanation from COMRECGC yield better coverage by a significant margin on all datasets, thus providing a better solution for the FCR problem than the baselines. The cost of the common recourse is also lower on the NCI1, MUTAGENICITY datasets, comparable on the AIDS dataset, and slightly above on the PROTEINS dataset. This difference in coverage is explained by the multi-head random walk used to generate counterfactuals in COMRECGC, that tends to facilitate forming common recourse afterwards. The lower cost is also explained by the fact that finding common recourse with similarity threshold $\Delta = 0.02$ is difficult, and is easier on small recourse than on larger ones.

4.4. Results for the FC problem

Baselines. The FC problem constraints the number of counterfactual used to form common recourse, which allows use to compare our method to common local counterfactual explainers. To evaluate the performance of our method compared to baselines, we construct a common recourse explanation for a counterfactual explainer as follows: a counterfactuals explainer takes the set of input graphs G and returns a set S of counterfactuals of the same cardinality. We form clusters on S using Algorithm 1 to generate a common recourse explanation. The baseline counterfactual explainers are GCFEXPLAINER and a local random walk explainer. The second one performs a random walk around each input graphs to find close counterfactuals, we select the closest counterfactual to each input graph to be part of the set of counterfactuals candidates used to generate recourse. We have not used explainers such as CFF EXPLAINER (Bajaj et al., 2021), RCEXPLAINER (Tan et al., 2022), as they mainly focus on local explanation and GCFEXPLAINER already out performs these methods (Huang et al., 2023). They also only generate smaller subgraphs counterfactual explanation, whereas a random walk can also adds elements to graph to find counterfactuals. Finally we also add a baseline that corresponds to Algorithm 1 processing the counterfactual graphs given in the original dataset.

Results. Table 4 presents the outcome. We show more detailed result for different number of recourse in Figure 2.

Table 3. Results on the FCR problem. Higher coverage corresponds to generating common recourse that is shared by more input graphs. The cost (lower is better) refers to the total length of the recourse. COMRECGC and its variants outperforms the baseline almost in all settings.

	NCI1		MUTAGEN	ICITY	Y AIDS		Prote	
	Coverage	Cost	Coverage	Cost	Coverage	Cost	Coverage	Cost
GCFEXPLAINER	21.4%	5.75	20.6%	6.91	14.2%	6.97	32.8%	10.65
COMRECGC 2-HEAD	40.5%	5.12	45.9%	5.74	32.8%	6.71	45.9%	11.44
ComRecGC 3-head	44.5%	5.07	52.6 %	5.61	33.6%	6.62	45.9%	11.51
COMRECGC 4-HEAD	44.6%	4.70	52.0%	5.81	34.8%	6.71	46.2%	11.47
COMRECGC 5-HEAD	42.9%	4.95	51.8%	5.63	34.7%	6.74	46.4%	11.59
COMRECGC 6-HEAD	44.9 %	4.51	52.0%	5.68	35.2%	6.66	47.3%	11.59

Table 4. Results on the FC problem. Higher coverage and lower cost are desirable. COMRECGC performs better than the baselines on all datasets and metrics except for the cost on the AIDS and PROTEINS datasets.

	NCI1		MUTAGEN	ICITY	AIDS		PROTEINS	
	Coverage	Cost	Coverage	Cost	Coverage	Cost	Coverage	Cost
DATASET COUNTERFACTUALS	8.52%	9.02	10.4%	8.34	0.41%	0.97	29.0%	12.95
LOCALRWEXPLAINER	19.0%	5.89	18.2%	7.19	12.9%	7.31	22.1%	11.33
GCFEXPLAINER	14.7%	7.12	11.9%	7.80	14.2%	7.07	29.8%	11.13
ComRecGC	33.4%	5.60	46.7%	6.56	24.3%	6.59	39.6%	12.04



Figure 1. Common Recourse on MUTAGENICITY: Removing an NO₂ complex. On the left two mutagenetic molecules from the input, on the right two resulting non-mutagenetic molecules.

We observe that COMRECGC generates the best coverage compared to all the baselines on the FC problem across all datasets. The cost of the common recourse is also lower than the recourse from other method, except on the Proteins dataset where they are comparable. The low cost of the dataset counterfactuals common recourse for the AIDS dataset is explained by the low number of common recourse formed, as few θ -close counterfactuals are in the dataset. This difference in coverage is even more striking in the FC problem than in the FCR problem. When the number of counterfactual is reduced from 100,000 to 2000, mostly close counterfactuals are grouped together, this is why the local RW explainer performances are better than the GCF-EXPLAINER global counterfactual for forming common recourse on the NCI1 and MUTAGENICITY datasets.

4.5. Common recourse explanations vs global counterfactual explanations

Lastly, we compare common recourse explanations and global counterfactual explanations in terms of coverage. Since COMRECGC is the only method designed to explain through common recourse, we will compare the common recourse explanation generated by COMRECGC to the global graph counterfactuals generated by GCFEX-PLAINER(Huang et al., 2023), CFF EXPLAINER (Bajaj et al., 2021) and RCEXPLAINER (Tan et al., 2022). We also add the global graph counterfactual explanation generated by the counterfactual given in the dataset.

For the definition of coverage for a set of counterfactual graph (S), we refer to (Huang et al., 2023), and define it as the ratio of the input graphs covered by at least one counterfactual in S. More formally COVERAGE(S) := $|\{G \in \mathbb{G} | min_{S \in \mathbb{S}} | | z(G) - z(S) | |_2 \le \theta \}| / |\mathbb{G}|.$

The results, shown in table 5 for 10 explanations, indicate that the common recourse explanations from COM-RECGC are comparable to the best graph counterfactual

	NCI1 Coverage	MUTAGENICITY Coverage	AIDS Coverage	PROTEINS Coverage
DATASET COUNTERFACTUALS	16.5%	29.0%	0.4%	8.5%
RCEXPLAINER	15.2%	32.0%	9.0%	8.7%
CFF	17.6%	30.4%	3.4%	3.8%
GCFEXPLAINER	27.9%	37.1%	14.7%	10.9%
ComRecGC	26.1%	39.4%	15.2%	18.0%

Table 5. Common recourse explanation vs global counterfactual explanation for 10 explanations. COMRECGC's common recourse explanations outperform the baselines' graph counterfactual explanations on most of the datasets.



Figure 2. Common Recourse coverage and cost comparison between COMRECGC and baselines for the FC problem where $\Delta = 0.02, T = |\mathbb{G}|$ and R = 1 to 100 common recourse.

method, GCFEXPLAINER, on the NCI1, MUTAGENICITY, and AIDS datasets, while being significantly more comprehensive on the PROTEINS dataset. This difference may be due to the sparsity of the Proteins dataset, which makes it challenging to identify a central counterfactual graph to explain the sparse neighborhood. Common recourse seems to be a more suitable explanation for this dataset, at least under the conditions of our experiments.

4.6. Discussions

We make a few observations including a few limitations of COMRECGC. *First*, since the FCR and FC problems are novel, our method is the only one specifically designed to generate common recourse. It is therefore unsurprising that we outperform the existing counterfactual explainers. *Second*, determining the appropriate parameters θ and Δ for the FCR problem is challenging and application dependent. Since these parameters serve as thresholds for a normalized distance, larger graphs generally work well with smaller values of θ and Δ , while smaller graphs require larger values. *Third*, we believe that relaxing the values of θ and Δ could lead to richer and more diverse explanations. However, as more common recourse would be generated in that process, we need an appropriate filtering system to select the most meaningful ones. *Fourth*, we explore the graph edit map without considering the physical feasibility of the changes, especially if the domain has specific constraints (e.g., chemical bonds in molecules). As a result, the generated counterfactual graphs may not correspond to real-world entities.

We discuss the choice of the random walk parameters in Appendix D.1 and bring additional results for different values of θ and Δ in Appendix D.2. We present the running time of COMRECGC in Appendix D.3, an ablation study in Appendix D.4 and other examples of common recourse in Appendix D.5.

5. Conclusion

In this work, we have formalized the problem of generating global counterfactual explanations for GNNs with common recourse. This novel problem setting includes the FCR and FC problems. These problems are NP-hard and thus, we have designed COMRECGC, a method specifically tailored to extract high-quality common recourse explanations. Through extensive experiments on real-world datasets, we have benchmarked COMRECGC against popular counterfactual explainers in the common recourse task. Our results demonstrate that COMRECGC produces global common recourse of significantly higher quality than the baselines across both the FCR and FC problems. Finally, COM-RECGC's common recourse explanations can account for a similar number of input graphs as those generated by global counterfactual explainers, providing a robust and scalable alternative to global graph counterfactual explanation.

6. Impact Statement

This paper presents work to advance the field of explainable Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

- Abrate, C. and Bonchi, F. Counterfactual graphs for explainable classification of brain networks. In *SIGKDD*, 2021.
- Abrate, C., Preti, G., and Bonchi, F. Counterfactual explanations for graph classification through the lenses of density, 2023.
- Bajaj, M., Chu, L., Xue, Z. Y., Pei, J., Wang, L., Lam, P. C.-H., and Zhang, Y. Robust counterfactual explanations on graph neural networks. In *NeurIPS*, 2021.
- Barman, S., Fawzi, O., Ghoshal, S., and Gürpinar, E. Tight approximation bounds for maximum multi-coverage. *CoRR*, 2019.
- Benaïm, M. and Tarrès, P. Dynamics of vertex-reinforced random walks. *The Annals of Probability*, 39(6), November 2011.
- Bhattoo, R., Ranu, S., and Krishnan, N. Learning articulated rigid body dynamics with lagrangian graph neural network. In *NeurIPS*, 2022.
- Borgwardt, K. M., Ong, C. S., Schönauer, S., Vishwanathan, S., Smola, A. J., and Kriegel, H.-P. Protein function prediction via graph kernels. *Bioinformatics*, 21(suppl_1): i47–i56, 2005.
- Buchbinder, N., Feldman, M., Naor, J., and Schwartz, R. Submodular maximization with cardinality constraints. In Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014, pp. 1433–1452. SIAM, 2014.
- Cao, S., Lu, W., and Xu, Q. Deep neural networks for learning graph representations. In *Proceedings of the Thirtieth* AAAI Conference on Artificial Intelligence, AAAI'16. AAAI Press, 2016.
- Cellinese, F., D'Angelo, G., Monaco, G., and Velaj, Y. The multi-budget maximum weighted coverage problem. In *Algorithms and Complexity*, 2021.
- Dobson, P. D. and Doig, A. J. Distinguishing enzyme structures from non-enzymes without alignments. *Journal of molecular biology*, 330(4):771–783, 2003.

- Ester, M., Kriegel, H.-P., Sander, J., and Xu, X. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Knowledge Discovery and Data Mining*, 1996. URL https://api.semanticscholar.org/CorpusID:355163.
- Feige, U. A threshold of ln *n* for approximating set cover. *J. ACM*, 45(4):634–652, 1998.
- Gorgi, A., Ouali, M. E., Srivastav, A., and Hachimi, M. Approximation algorithm for the multicovering problem. *J. Comb. Optim.*, 2021.
- Hall, N. G. and Hochbaum, D. S. A fast approximation algorithm for the multicovering problem. *Discret. Appl. Math.*, 1986.
- Hamilton, W. L., Ying, Z., and Leskovec, J. Inductive representation learning on large graphs. In *NeurIPS*, 2017.
- Hua, Q., Yu, D., Lau, F. C. M., and Wang, Y. Exact algorithms for set multicover and multiset multicover problems. In *ISAAC*, 2009.
- Huang, Z., Kosan, M., Medya, S., Ranu, S., and Singh, A. K. Global counterfactual explainer for graph neural networks. In WSDM, 2023.
- Jiang, M., Li, Z., Zhang, S., Wang, S., Wang, X., Yuan, Q., and Wei, Z. Drug–target affinity prediction using graph neural network and contact maps. *RSC advances*, 10(35): 20701–20712, 2020.
- Kakkad, J., Jannu, J., Sharma, K., Aggarwal, C., and Medya, S. A survey on explainability of graph neural networks, 2023.
- Karimi, A.-H., Schölkopf, B., and Valera, I. Algorithmic recourse: from counterfactual explanations to interventions, 2020.
- Kazius, J., McGuire, R., and Bursi, R. Derivation and validation of toxicophores for mutagenicity prediction. *Journal of medicinal chemistry*, 48(1):312–320, 2005.
- Khuller, S., Moss, A., and Naor, J. S. The budgeted maximum coverage problem. *Information Processing Letters*, 1999.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv:1412.6980*, 2014.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017.
- Ley, D., Mishra, S., and Magazzeni, D. Globe-ce: A translation based approach for global counterfactual explanations. In *ICML*, 2023.

- Lovász, L. Random walks on graphs. *Combinatorics, Paul erdos is eighty*, 2(1-46):4, 1993.
- Lucic, A., Ter Hoeve, M. A., Tolomei, G., De Rijke, M., and Silvestri, F. Cf-gnnexplainer: Counterfactual explanations for graph neural networks. In *AISTATS*, 2022.
- Magister, L. C., Kazhdan, D., Singh, V., and Liò, P. Gcexplainer: Human-in-the-loop concept-based explanations for graph neural networks, 2021.
- Mei, Q., Guo, J., and Radev, D. R. Divrank: the interplay of prestige and diversity in information networks. In *SIGKDD*, 2010.
- Mirhoseini, A., Goldie, A., Yazgan, M., Jiang, J. W. J., Songhori, E. M., Wang, S., Lee, Y., Johnson, E., Pathak, O., Bae, S., Nazi, A., Pak, J., Tong, A., Srinivasa, K., Hang, W., Tuncer, E., Babu, A., Le, Q. V., Laudon, J., Ho, R., Carpenter, R., and Dean, J. Chip placement with deep reinforcement learning. *CoRR*, abs/2004.10746, 2020.

Pemantle, R. Vertex-reinforced random walk, 2004.

- Pemantle, R. A survey of random processes with reinforcement. *Probability Surveys*, 4, January 2007.
- Ranjan, R., Grover, S., Medya, S., Chakaravarthy, V., Sabharwal, Y., and Ranu, S. Greed: A neural framework for learning graph distance functions, 2023.
- Rawal, K. and Lakkaraju, H. Beyond individualized recourse: Interpretable and interactive summaries of actionable recourses. In *NeurIPS*, 2020.
- Riesen, K. and Bunke, H. Iam graph database repository for graph based pattern recognition and machine learning. In *Joint IAPR International Workshops on Statistical Techniques in Pattern Recognition (SPR) and Structural and Syntactic Pattern Recognition (SSPR)*, pp. 287–297. Springer, 2008.
- Sanfeliu, A. and Fu, K.-S. A distance measure between attributed relational graphs for pattern recognition. *IEEE* transactions on systems, man, and cybernetics, (3):353– 362, 1983.
- Santiago, R. and Yoshida, Y. Weakly submodular function maximization using local submodularity ratio. In *ISAAC*, volume 181 of *LIPIcs*, 2020.
- Tan, J., Geng, S., Fu, Z., Ge, Y., Xu, S., Li, Y., and Zhang, Y. Learning and evaluating graph neural network explanations based on counterfactual and factual reasoning. In *WebConf*, 2022.
- Tanyel, T., Ayvaz, S., and Keserci, B. Beyond known reality: Exploiting counterfactual explanations for medical research. *CoRR*, abs/2307.02131, 2023.

- Verma, S., Boonsanong, V., Hoang, M., Hines, K. E., Dickerson, J. P., and Shah, C. Counterfactual explanations and algorithmic recourses for machine learning: A review, 2022.
- Volkov, S. Vertex-reinforced random walk on arbitrary graphs, 1999.
- Vu, M. and Thai, M. T. Pgm-explainer: Probabilistic graphical model explanations for graph neural networks. In *NeurIPS*, 2020.
- Wale, N. and Karypis, G. Comparison of descriptor spaces for chemical compound retrieval and classification. In *ICDM*, 2006.
- Ying, R., Bourgeois, D., You, J., Zitnik, M., and Leskovec, J. Gnnexplainer: Generating explanations for graph neural networks, 2019.
- Yuan, H., Yu, H., Gui, S., and Ji, S. Explainability in graph neural networks: A taxonomic survey, 2022.

A. Additional Related Work

We summarize recent approaches for counterfactual explanations. (Abrate et al., 2023) propose counterfactual explanations by sparsifying and densifying graphs, to find counterfactuals through pattern deletion or generation. ARES (Rawal & Lakkaraju, 2020) produce actionable two level recourse summaries on non-graphical datasets through an optimization problem. Lastly, in (Magister et al., 2021) the "Human-in-the-Loop" framework integrates human feedback to enhance the relevance and interpretability of counterfactual explanations.

B. Additional Proofs and Analyses for Section 2

B.1. Proof of Theorem 1

Proof. We prove that the FCR problem is NP-Hard through a reduction from the maximum coverage problem. Consider the problem of covering U using the sets $(S_i)_i$, we build the following instance of the FCR problem:

Let X be a graph binary classifier that accepts a graph if and only if the graph has two nodes of the same color. To each $u_j \in U$, we associate G_j , a star graph with an unlabelled central node and m peripheral vertices. If u_j is in k elements of S, we one-color k peripheral vertices with the colors $\{i : u_j \in S_i\}$. We define the common recourse r_i , for $i \leq m$, as follows: r_i colors the middle vertex of a star graph using color i, and we take the corresponding counterfactuals as the inputs to this FCR problem.

Thus $X(G_j) = 0$ for all j and $X(r_i(G_j)) = 1$ if and only if $u_j \in S_i$. Hence from an optimal solution to this FCR problem, we derive an optimal solution for the maximum coverage problem $(U, (S_i)_i)$. Therefore FCR is NP-hard.

B.2. Analysis for The FCR problem

Let us define f as the function that associate to a set of common recourse the maximum coverage obtained by selecting R of them. It is not hard to see that f is monotone, i.e $f(\mathbb{F} \cup \{r\}) \ge f(\mathbb{F})$ for every set of recourse \mathbb{F} , and recourse r. We recall the definition of submodularity (Santiago & Yoshida, 2020), a function h is submodular if for any two sets $A \subseteq B$ and for any element e:

$$h(A \cup \{e\}) - h(A) \ge h(B \cup \{e\}) - h(B)$$
(4)

It is easy to see that f is submodular. In the rest of the paper, we denote by $h_A(B)$ the marginal gain of adding the set B to A, i.e $h(A \cup B) - h(A)$.

Proof. We give an instance of the FC problem, and two sets A, B such that equation 5 is only true for $\gamma_{A,B} = 0$, which means that there is no local pseudo-submodularity bound. Consider the instance of the FC problem where $R = 2, T \ge 4$ and $(G_i)_{1 \le i \le 4}$ some counterfactual graphs such that they each cover a different input graphs, G_1 is part of R_1 (recourse 1), G_2 part of R_2 (recourse 2), G_3, G_4 parts of R_3 (recourse 3). We set A to $\{G_1, G_2\}$ and B to $\{G_3, G_4\}$. Then $g_A(G_3) = g_A(G_4) = 0$ and $g_A(B) = 1$.

B.3. Budget Version of the FC problem

When we assume that the counterfactuals and recourse are all known, the FC problem becomes a budget problem: **Problem** 3 (max 2-budget 2-cover problem). Given two budgets k_1 and k_2 , and two families of sets S_1 , S_2 , the goal is to find $S_1 \subset S_1$, $S_2 \subset S_2$ verifying :

$$max_{\mathcal{S}_1,\mathcal{S}_2} S_1 \cap S_2$$
 st. $(\mathcal{S}_1) \leq k_1$ and $size(\mathcal{S}_2) \leq k_2$,

where $S_i = \bigcup_{S \in S_i} S$.

To the best of our knowledge, this problem has not been studied. This is a variation of the maximum coverage problem. The budgeted maximum coverage problem (Khuller et al., 1999) is another budget variation, where only one budget is considered, there is no constraint about the intersection, but the sets costs and the rewards for covering elements can be set to be different from 1. The multi-budget maximum coverage problem is another variation between knapsack and maximum coverage (Cellinese et al., 2021), but there is only one family of set to pick from, and we do not require double coverage. Another variation of the maximum coverage problem is called multi-coverage (Barman et al., 2019), where covering multiple times one element raises the objective value.

Another variation is called multi-set multi-cover, uses one budget, and each element must be covered a certain number of times (Hall & Hochbaum, 1986; Hua et al., 2009; Gorgi et al., 2021).

B.4. Proof of Theorem 2

A function h is called pseudo-modular if there exists $\gamma \in (0, 1]$ such that for any pair of disjoint sets A, B:

$$\sum_{e \in B} h_A(e) \ge \gamma h_A(B) \tag{5}$$

Note that this γ is the minimum over all the $\gamma_{A,B}$, for all sets A, B, where $\gamma_{A,B}$ is a specific value of γ for given A, B in equation 5. We now proceed to the proof of Theorem 2:

Proof. We give an instance of the FC problem, and two sets A, B such that equation 5 is only true for $\gamma_{A,B} = 0$, which means that there is no local pseudo-submodularity bound. Consider the instance of the FC problem where $R = 2, T \ge 4$ and $(G_i)_{1 \le i \le 4}$ some counterfactual graphs such that they each cover a different input graphs, G_1 is part of R_1 (recourse 1), G_2 part of R_2 (recourse 2), G_3, G_4 parts of R_3 (recourse 3). We set A to $\{G_1, G_2\}$ and B to $\{G_3, G_4\}$. Then $g_A(G_3) = g_A(G_4) = 0$ and $g_A(B) = 1$.

B.5. Proof of Theorem 3

We first prove that g, the optimization function for the FC problem with the added constraint C1, is pseudo-modular.

Proof. For every set A, and for every graph e in the problem space we have that $R \ge g_A(e)$ by problem definition, we also have that $g_A(e) > 0$ by C1. So for two disjoint sets A, B part of the solution :

$$\sum_{e \in B} g_A(e) \ge \frac{1}{R|B|} g_A(B).$$

Hence g is pseudo-modular.

The pseudo code of the Greedy algorithm for the FC problem with C1 is as follows, where g denotes the function that associates to a set of counterfactual its best coverage through building common recourse and selecting the R best for coverage, and T is the maximum number of counterfactuals to use for the explanations.

Algorithm 4 GREEDY FC(q, T)

1: Initialize the set of counterfactuals: $S_0 \leftarrow \emptyset$ 2: for $i \in 1 : T$ do 3: Let $M_i \subseteq E \setminus S_{i-1}$ be a subset of size T maximizing $\sum_{e \in M_i} g_{S_{i-1}}(e)$ 4: Let e_i be an element uniformly chosen at random of M_i 5: $S_i \leftarrow S_{i-1} + e_i$ 6: end for 7: return S_T

This algorithm was first introduced in the work of (Buchbinder et al., 2014). Applying Theorem 1.10 of (Santiago & Yoshida, 2020) to the monotone and pseudo-modular function g (Theorem 3), shows that Algorithm 4 yields a $1 - e^{-1/R}$ approximation in expectation to the FC problem, where R is the maximum number of common recourse of the FC problem.

 \square

C. Detail of our method, COMRECGC for Section 3

C.1. The Pseudocode of MULTI-HEAD VRRW

The pseudo code for our multi-head VRRW is as follows, where G_i denotes the input graph where the random walk started or was last teleported for the *i*-th head, and H_i denotes the current graph head *i* is occupying:

Algorithm 5 MULTI-HEAD VRRW(ϕ , \mathbb{G} , k, M, τ)

```
1: (G_1, \ldots, G_k) \leftarrow random input graph from \mathbb{G}, \mathbb{S} = \emptyset
 2: (H_1,\ldots,H_k) \leftarrow (G_1,\ldots,G_k)
 3: for t \in 1 : M do
         Let \epsilon, \ell \sim Bernoulli(\tau), \mathcal{U}\{1, \ldots k\}
 4:
 5:
         if \epsilon = 0 then
 6:
             for v \in \mathcal{N}(H_{\ell}) do
 7:
                Compute P(H_{\ell}, v) based on equation 2.
 8:
             end for
             v_{\ell} \leftarrow random neighbor of H_{\ell} based on p(H_{\ell}, v)
 9:
             for i \in 1: k, i \neq \ell do
10:
                 v_j \leftarrow argmin_{v \in \mathcal{N}(u_i)} || \overline{z(G_\ell)z(u_\ell)} - \overline{z(G_i)z(v)} ||_2
11:
12:
             end for
13:
         else
             (G_1, \ldots, G_k) \leftarrow random input graphs from \mathbb{G} based on equation 3.
14:
15:
             (v_1,\ldots,v_k) \leftarrow (G_1,\ldots,G_k).
16:
         end if
17:
         for i \in 1: k do
18:
             C(v_i) \leftarrow C(v) + 1
19:
             if \phi(v_i) > 0.5 then
20:
                 \mathbb{S} \leftarrow \mathbb{S} + \{v_i\}
21:
             end if
22:
         end for
23:
         (H_1,\ldots,H_k) \leftarrow (v_1,\ldots,v_k).
         N(H_1) = N(H_1) + 1, \dots, N(H_k) = N(H_k) + 1.
24:
25: end for
26: return S
```

Lines 1-2 describe the initialization of the random walk, where S represents the set of counterfactuals we aim to find. Lines 5-12 outline the regular behavior of the random walk as it moves towards counterfactuals. Lines 13-15 cover the teleportation process. Finally, lines 17-21 detail the update of the counterfactual set. In line 23, we advance one step in the random walk, and in line 24, we update the visit count of the graphs reached by the *k* heads during this step, following the principles of VRRW. This visit count influences the transition probabilities in Equation 2.

C.2. Analysis of the Random Walks

VRRW. A vertex reinforced random walk is a random process with reinforcement that focuses on the number of visits of vertices (Pemantle, 2004) (Pemantle, 2007). To the best of our knowledge, theorical guarantees have only been derived for the special case where the initial transition matrix is symmetric (Volkov, 1999) (Benaïm & Tarrès, 2011), which is not relevant for our application, as we want to identify counterfactuals with the random walk, as described in Equation 2.

RW. Let us consider a classic, non-reinforced, random walk with the following transition rule:

$$p(u_{\ell}, v) \propto p_{\phi}(v) \tag{6}$$

Where $p_{\phi}(v)$ is the probability, assigned by the GNN, of vertex v being a counterfactual, and uniform teleportation:

$$p_{\tau}(G) = \frac{1}{|\mathbb{G}|} \tag{7}$$

We refer to (Lovász, 1993) for the theoretical results. The main purpose of the analysis for our application is to determine the mixing time and mixing rate of our walk, to tune the parameter M, the number of steps. We add to the random walk the following constraint: the maximum distance of the head from its starting point is at most $3\theta/2$, to make the space of the

random walk finite. By Corollary 5.2 in (Lovász, 1993), the mixing rate of the walk is $max(|\lambda_2|, |\lambda_n|)$, where λ_i is the *i*-th largest eigenvalue of the matrix $D^{-1/2}M_{rw}D^{1/2}$, where M_{rw} is the transition matrix associated to our random walk and D is the diagonal matrix with value 1/degree(i).

Unfortunately, there are two main obstacles that make this approach untractable: (1) The search space, the number of nodes of the graph on which we do the random walk, is of size $O(|\mathbb{G}|e^r\theta)$, hence making it hard to extract the eigenvalues of the matrix M_{rw} (2) Each node's transition rule is determined by the GNN's prediction, making the construction of the matrix M_{rw} difficult.

D. Additional Details on Experiments & Results

D.1. Choice of the random walk parameters k, M, τ

The value of parameters k and τ have been chosen through experiments on the Mutagenicity datasets. We chose τ , the probability of teleport, to have a compromise between being able to look far enough to find counterfactuals, and still be able to explore a good amount of paths with the random walk. Although higher values of k seem to yield slightly better coverage, as seen on table 3, k increases the running time, the number of steps for the random walk to converge, and the memory usage. Hence we decide to limit k to 5 heads and to set M to a relatively low 50,000 steps.

D.2. Experiments on different values of θ, Δ

Setting. We experiment on different values of θ , Δ . Those parameters are only used in the clustering algorithm part of COMRECGC. The results are presented in table 6, for COMRECGC with parameters $M = 50,000, k = 5, \tau = 0.05$.

Results. We do not run into issues with increasing the value of Δ , the common recourse threshold. As expected, the coverage goes up, and so does the cost as COMRECGC is able to cover close to the whole of MUTAGENICITY and NCI1 datasets, and follows the same trend for the AIDS dataset. However, it is quite surprising that the coverage on the PROTEINS dataset went down by 1%. An explanation for this is a possible issue with the clustering algorithm DBScan, which may be caused by a too big overlap on some common recourse clusters.

Increasing the counterfactual threshold θ is proven to be challenging, as raising it from 0.1 to 0.15 across datasets dramatically increased the number of recourse entering the clustering stage. For example, on the MUTAGENICITY and NCI1 datasets, the number of recourse increased by a factor of 100, making DBScan intractable in its current form. As expected, increasing θ results in greater coverage in the common recourse explanations produced by COMRECGC, but it also raises the cost, as longer recourse are allowed. In particular, on the AIDS dataset, the increase in coverage was minimal, suggesting a potential bottleneck in the effectiveness of common recourse explanations for that dataset with common recourse threshold set at $\Delta = 0.02$.

Table 6. Results on the FCR problem of COMRECGC for different values of θ and Δ . We find that increasing the common recourse and counterfactual threshold tend improve coverage, particularly or more dense datasets, such as MUTAGENICITY and NCI1. The effect on more sparse datasets such as PROTEINS and AIDS is more contrasted, with moderate to no added coverage. The cost rises as we allow more distant counterfactuals.

	NCI1		MUTAGENICITY		AIDS		PROTEINS	
	Coverage	Cost	Coverage	Cost	Coverage	Cost	Coverage	Cost
$\theta = 0.1, \Delta = 0.02$	42.9%	4.95	51.8%	5.63	34.7%	6.74	42.8%	7.21
$\theta = 0.1, \Delta = 0.04$	86.5%	8.05	90.0%	8.11	47.0%	8.4	41.5%	6.36
$\theta=0.15, \Delta=0.02$	х	Х	Х	Х	34.8%	7.65	46.4%	11.59

D.3. Running time

We show the running time of COMRECGC, its VRRW (Algorithm 5) and Clustering (Algorithm 1) components, for the FCR problem in the datasets of our study.

	NCI1	MUTAGENICITY	AIDS	PROTEINS
Algorithm 5	101	172	61	152
Algorithm 1	23	27	14	14
COMRECGC	124	199	75	166

Table 7. The running time in minutes of the different parts of COMRECGC for the experiments of Table 3.

D.4. Ablation Study

We give the performance of the variation of COMRECGC without some of its features on the FCR problem for the datasets in our study. We find the clustering to be very important on large datasets such as MUTAGENICITY, but less so on the smaller PROTEINS. This is possibly explained by the ability of the VRRW to spend a longer time pairing the same elements through different recourse, hence somewhat harmonizing some recourse.

Table 8. Results on the FCR	problem of different variations	of COMRECGC, in the same	parameters setting as Section 4.
ruore of results on the r ort	proofern of aniferent (anialions)		purumeters setting us seetion m

	NCI1		MUTAGEN	ENICITY A		5	PROTEINS	
	Coverage	Cost	Coverage	Cost	Coverage	Cost	Coverage	Cost
COMRECGC	42.9%	4.95	51.8 %	5.63	34.7%	6.74	46.4%	11.59
COMRECGC w/o clustering	10.1%	Х	8.2%	х	13.6%	Х	33.9%	Х
COMRECGC w/o multi-head	21.4%	5.75	20.6%	6.91	14.2%	6.97	32.8%	10.65

D.5. Additional examples of common recourse

We provide additional example of common recourse identified through COMRECGC, in the settings $\theta = 0.1$ and $\Delta = 0.02$. in Figures 3 and 4.

In Figure 3, we observe larger molecules, which allow for a 'larger recourse' consisting of three transformations. This is because transformations on larger molecules correspond to smaller variations in the embedding for the "same transformation", which is capped by Δ .

In Figure 4, on smaller molecules, we observe two types of transformations in the recourse, likely due to the increased value of the θ parameter compared to Figure 1. Indeed, a larger θ allows for more distant counterfactuals.



Figure 3. Common Recourse on the MUTAGENICITY dataset: removing two Hydrogen and one Carbon, on the left two mutagenetic input graphs, on the right two non-mutagenetic graph.



Figure 4. Common Recourse on the MUTAGENICITY dataset: removing one Nitrogen, one Hydrogen, adding one Oxygen and one Carbon, on the left two mutagenetic input graphs, on the right two non-mutagenetic graph.