ComRecGC: Global Graph Counterfactual Explainer through Common Recourse

Gregoire Fournier[1] gfourn2@uic.edu
Sourav Medya[1] medya@uic.edu

[1]University of Illinois Chicago, USA

June 23, 2025

Summary

Contribution:

We introduce a method to compute a small set of *common recourse*, graph transformations that can flip GNN decisions across many inputs.

Key Insight:

While local counterfactuals are hard to generalize, and global ones may lack actionable recourse, *Finding Common Recourse (FCR)* balances interpretability and actionability.

Motivation:

- Improve model-level understanding of GNN behavior.
- Provide compact and meaningful graph-level explanations.

Background and Problem Setup

- **Graph Neural Networks (GNNs):** Widely used for structured data, but their decisions remain hard to interpret.
- **Counterfactual Explanations:** Suggest minimal changes (recourse) to flip a prediction, helping users understand and act.
- **Limitation of Local/Global CEs:**
  - Local CEs lack generality and can be overwhelming.
  - Global CEs may not provide consistent recourse paths.
- **Problem Statement:**
  We formalize the *Finding Common Recourse (FCR)* problem and its variant *(FC)*, aiming to find a compact set of shared edits that generalize across rejected graphs.

Why Common Recourse?

Counterfactual Explanations and Recourse
Graphs are classified by a GNN as `accept` or `reject`. A counterfactual shows small edits to flip prediction; these transformations are called **recourse**.

Benefits:

- Scalable understanding of model decisions
- Actionable, minimal changes to flip outcomes
- Trust and fairness through consistent logic

Applications:

- Drug discovery: editing molecules to reduce toxicity
- Credit scoring: identifying changes to get loan approval
- Legal tech: exposing fair, generalizable decision criteria

Gregoire Fournier, Sourav Medya

Example

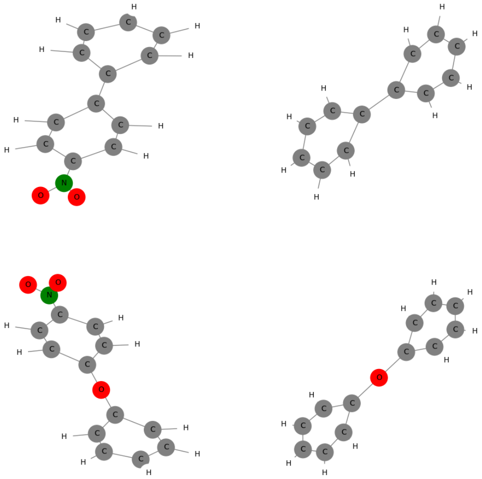

Figure: Common Recourse on Mutagenicity: Removing an NO_2 complex. On the left two mutagenetic molecules from the input, on the right two resulting non-mutagenetic molecules.

Method Overview: COMRECGC

Step 1 – Recourse Embeddings

We compute GNN embeddings for each input graph and its counterfactual to represent their transformation as a vector in latent space.

Step 2 – Generating Diverse Counterfactuals

We explore possible graph edits via a *multi-head vertex-reinforced random walk*, which favors frequently visited edits while encouraging exploration.

Step 3 – Finding Common Recourse

We cluster the recourse embeddings using a fixed radius. Each cluster captures a generalizable recourse applicable to many graphs.

Step 4 – Greedy Aggregation

We select a small number of clusters (recourse options) to maximize coverage over the input graphs.

Gregoire Fournier, Sourav Medya

Experiments: FC Problem

- Compared to local and global baselines under budgeted counterfactual generation.
- **Performance:** ComRecGC offers the best tradeoff: **highest coverage** and **lowest cost**.
- **Key insight:** Outperforms even under tighter constraints on graph edits.

	NCI1		Mutagenicity		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

Common Recourse vs Global Counterfactual Explanations

- **Performance:** ComRecGC matches the best global method (GCFEXPLAINER) on `NCI1`, `MUTAGENICITY`, and `AIDS`. ComRecGC outperforms baselines `PROTEINS`.
- **Key insight:** On sparse datasets, shared recourse better captures diverse decision boundaries.

	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%**

Conclusion

- We introduce: **common recourse for global GNN explanations**.
- Our method, ComRecGC, solves the NP-hard *FCR* and *FC* problems.
- Results show:
  - Higher-quality, shared counterfactuals
  - Competitive or better coverage than global baselines
- **Scalable, interpretable**, and effective for real-world graphs.