

OPEN
ARTICLE

Evaluating explainability for graph neural networks

Chirag Agarwal^{1,2,8}, Owen Queen^{2,3,8}, Himabindu Lakkaraju^{4,5,6} & Marinka Zitnik^{1,2,5,7}

As explanations are increasingly used to understand the behavior of graph neural networks (GNNs), evaluating the quality and reliability of GNN explanations is crucial. However, assessing the quality of GNN explanations is challenging as existing graph datasets have no or unreliable ground-truth explanations. Here, we introduce a synthetic graph data generator, SHAPEGGEN, which can generate a variety of benchmark datasets (e.g., varying graph sizes, degree distributions, homophilic vs. heterophilic graphs) accompanied by ground-truth explanations. The flexibility to generate diverse synthetic datasets and corresponding ground-truth explanations allows SHAPEGGEN to mimic the data in various real-world areas. We include SHAPEGGEN and several real-world graph datasets in a graph explainability library, GRAPHXAI. In addition to synthetic and real-world graph datasets with ground-truth explanations, GRAPHXAI provides data loaders, data processing functions, visualizers, GNN model implementations, and evaluation metrics to benchmark GNN explainability methods.

Introduction

As graph neural networks (GNNs) are being increasingly used for learning representations of graph-structured data in high-stakes applications, such as criminal justice¹, molecular chemistry^{2,3}, and biological networks^{4,5}, it becomes critical to ensure that the relevant stakeholders can understand and trust their functionality. To this end, previous work developed several methods to explain predictions made by GNNs^{6–14}. With the increase in newly proposed GNN explanation methods, it is critical to ensure their reliability. However, explainability in graph machine learning is an emerging area lacking standardized evaluation strategies and reliable data resources to evaluate, test, and compare GNN explanations¹⁵. While several works have acknowledged this difficulty, they tend to base their analysis on specific real-world² and synthetic¹⁶ datasets with limited ground-truth explanations. In addition, relying on these datasets and associated ground-truth explanations is insufficient as they are not indicative of diverse real-world applications¹⁵. To this end, developing a broader ecosystem of data resources for benchmarking state-of-the-art GNN explainers can support explainability research in GNNs.

A comprehensive data resource to correctly evaluate the quality of GNN explanations will ensure their reliability in high-stake applications. However, the evaluation of GNN explanations is a growing research area with relatively little work, where existing approaches mainly leverage ground-truth explanations associated with specific datasets² and are prone to several pitfalls (as outlined by Faber *et al.*¹⁶). Further, multiple underlying rationales can generate the correct class labels, creating redundant or non-unique explanations. A trained GNN model may only capture one or an entirely different rationale. In such cases, evaluating the explanation output by a state-of-the-art method using the ground-truth explanation is incorrect because the underlying GNN model does not rely on that ground-truth explanation. In addition, even if a unique ground-truth explanation generates the correct class label, the GNN model trained on the data could be a weak predictor using an entirely different rationale for prediction. Therefore, the ground-truth explanation cannot be used to assess post hoc explanations of such models. Lastly, the ground-truth explanations corresponding to some of the existing benchmark datasets are not good candidates for reliably evaluating explanations as they can be recovered using trivial baselines (e.g., random node or edge as explanation). The above discussion highlights a clear need for general-purpose data resources which can evaluate post hoc explanations reliably across diverse real-world applications. While various benchmark datasets (e.g., Open Graph Benchmark (OGB)¹⁷, GNNMark¹⁸, GraphGT¹⁹, MalNet²⁰, Graph

¹Media and Data Science Research Lab, Adobe, Noida, 201304, India. ²Department of Biomedical Informatics, Harvard University, Boston, MA, 02115, USA. ³Department of Electrical Engineering and Computer Science, University of Tennessee, Knoxville, TN, 37996, USA. ⁴Harvard Business School, Boston, MA, 02163, USA. ⁵Harvard Data Science Initiative, Cambridge, MA, 02138, USA. ⁶Department of Computer Science, Harvard University, Boston, MA, 02134, USA. ⁷Broad Institute of MIT and Harvard, Cambridge, MA, 02142, USA. ⁸These authors contributed equally: Chirag Agarwal, Owen Queen. ✉e-mail: marinka@hms.harvard.edu

Robustness Benchmark (GRB)²¹, Therapeutics Data Commons^{22,23}, and EFO-1-QA²⁴) and programming libraries for deep learning on graphs (e.g., Dive Into Graphs (DIG)²⁵, Pytorch Geometric (PyG)²⁶, and Deep Graph Library (DGL)²⁷) in graph machine learning literature exist, they are mainly used to only benchmark the performance of GNN predictors and are not suited to evaluate the correctness of GNN explainers because they do not capture ground-truth explanations.

Here, we address the above challenges by introducing a general-purpose data resource that is not prone to ground-truth pitfalls (e.g., redundant explanations, weak GNN predictors, trivial explanations) and can cater to diverse real-world applications. To this end, we present SHAPEGGEN (Fig. 2), a novel and flexible synthetic XAI-ready (explainable artificial intelligence ready) dataset generator which can automatically generate a variety of benchmark datasets (e.g., varying graph sizes, degree distributions, homophilic vs. heterophilic graphs) accompanied by ground-truth explanations. SHAPEGGEN ensures that the generated ground-truth explanations are not prone to the pitfalls described in Faber *et al.*¹⁶, such as redundant explanations, weak GNN predictors, and trivially correct explanations. Furthermore, SHAPEGGEN can evaluate the goodness of any explanation (e.g., node feature-based, node-based, edge-based) across diverse real-world applications by seamlessly generating synthetic datasets that can mimic the properties of real-world data in various domains.

We incorporate SHAPEGGEN and several other synthetic and real-world graphs²⁸ into GRAPHXAI, a general-purpose framework for benchmarking GNN explainers. GRAPHXAI also provides a broader ecosystem (Fig. 1) of data loaders, data processing functions, visualizers, and a set of evaluation metrics (e.g., accuracy, faithfulness, stability, fairness) to reliably benchmark the quality of any given GNN explanation (node feature-based or node/edge-based). We leverage various synthetic and real-world datasets and evaluation metrics from GRAPHXAI to empirically assess the quality of explanations output by eight state-of-the-art GNN explanation methods. Across many GNN explainers, graphs, and graph tasks, we observe that state-of-the-art GNN explainers fail on graphs with large ground-truth explanations (*i.e.*, explanation subgraphs with a higher number of nodes and edges) and cannot produce explanations that preserve fairness properties of underlying GNN predictors.

Results

To evaluate GRAPHXAI, we show how GRAPHXAI enables systematic benchmarking of eight state-of-the-art GNN explainers on both SHAPEGGEN (in the Methods section) and real-world graph datasets. We explore the utility of the SHAPEGGEN generator to benchmark GNN explainers on graphs with homophilic vs. heterophilic, small vs. large, and fair vs. unfair ground-truth explanations. Additionally, we examine the utility of GNN explanations on datasets with varying degrees of informative node features. Next, we outline the experimental setup, including details about performance metrics, GNN explainers, and underlying GNN predictors, and proceed with a discussion of benchmarking results.

Experimental setup. *GNN explainers.* The GRAPHXAI defines an Explanation class capable of storing multiple types of explanations produced by GNN explainers and provides a `graphxai.BaseExplainer` class that serves as the base for all explanation methods in GRAPHXAI. We incorporate eight GNN explainability methods, including gradient-based: Grad²⁹, GradCAM¹¹, GuidedBP⁶, Integrated Gradients³⁰; perturbation-based: GNNExplainer¹⁴, PGExplainer¹⁰, SubgraphX³¹; and surrogate-based methods: PGMExplainer¹³. Finally, following Agarwal *et al.*¹⁵, we consider random explanations as a reference: (1) Random Node Features, a node feature mask defined by an d -dimensional Gaussian distributed vector; (2) Random Nodes, a $1 \times n$ node mask is randomly sampled using a uniform distribution, where n is the number of nodes in the enclosing subgraph; and (3) Random Edges, an edge mask drawn from a uniform distribution over a node's incident edges.

Implementation details. We use a three-layer GIN model³² and a GCN model³³ as GNN predictors for our experiments. We use a model comprising three GIN convolution layers with ReLU non-linear activation function and a Softmax activation for the final layer. The hidden dimensionality of the layers is set to 16. We follow an established approach for generating explanations^{8,15} and use reference algorithm implementations. We select top- k ($k = 25\%$) important nodes, node features, or edges and use them to generate explanations for all graph explainability methods. For training GIN models, we use an Adam optimizer with a learning rate of 1×10^{-2} , weight decay of 1×10^{-5} , and the number of epochs to 1000. We use an Adam optimizer with a learning rate of 3×10^{-2} , no weight decay, and 1500 training epochs for training GNN models. We set hyperparameters of GNN explainability models following the authors' recommendations.

We use a fixed random split provided within the GRAPHXAI package to split the datasets. For each SHAPEGGEN dataset, we use a 70/5/25 split for training, validation, and testing, respectively. For MUTAG, Benzene, and Fluoride Carbonyl datasets, we use a 70/10/20 split throughout each dataset. Average performance is reported across each sample in the testing set of each dataset.

Performance metrics. In addition to the synthetic and real-world data resources, we consider four broad categories of performance metrics: (i) Graph Explanation Accuracy (GEA); (ii) Graph Explanation Faithfulness (GEF); (iii) Graph Explanation Stability (GES); and (iv) Graph Explanation Fairness (GECE, GEGF) to evaluate the explanations on the respective datasets. In particular, all evaluation metrics leverage predicted explanations, ground-truth explanations, and other user-controlled parameters, such as top- k features. Our GRAPHXAI package implements these performance metrics and additional utility functions within `graphxai.metrics` module. Figure 3 shows a code snippet for evaluating the correctness of output explanations for a given GNN prediction in GRAPHXAI.

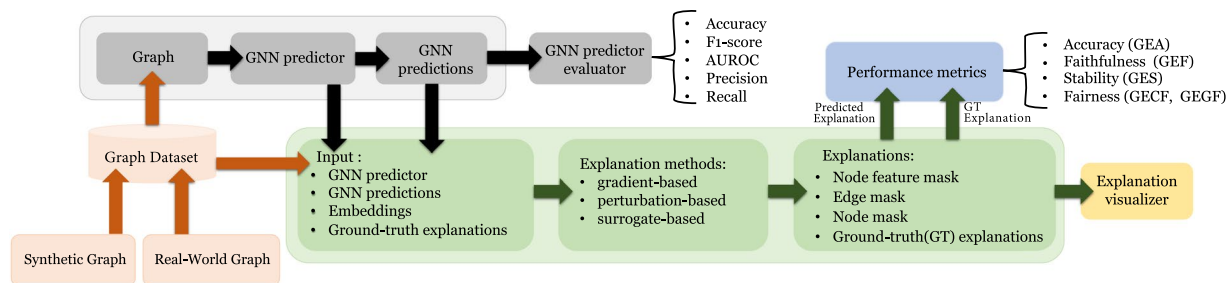


Fig. 1 Overview of GRAPHXAI. GRAPHXAI provides data loader classes for XAI-ready synthetic and real-world graph datasets with ground-truth explanations for evaluating GNN explainers, implementations of explanation methods, visualization functions for GNN explainers, utility functions to support new GNN explainers, and a diverse set of performance metrics to evaluate the reliability of explanations generated by GNN explainers.

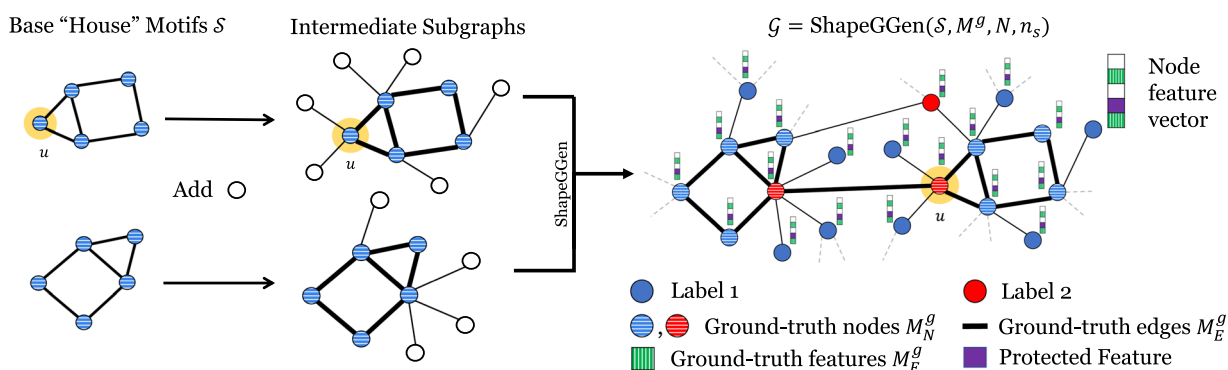


Fig. 2 Overview of SHAPEGGEN graph dataset generator. SHAPEGGEN is a novel dataset generator for graph-structured data that can be used to benchmark graph explainability methods using ground-truth explanations. Graphs are created by combining subgraphs containing any given motif and additional nodes. The number of motifs in a k -hop neighborhood determines the node label (in the figure, we use a 1-hop neighborhood for labeling, and nodes with two motifs in its 1-hop neighborhood are highlighted in red). Feature explanations are some masks over important node features (green striped), with an option to add a protected feature (shown in purple) whose correlation to node labels is controllable. Node explanations are nodes contained in the motifs (horizontal striped nodes) and edge explanations (bold lines) are edges connecting nodes within motifs.

```
import graphxai
from graphxai.explainers import GradCAM
from graphxai.metrics import graph_exp_faith

dataset = graphxai.get_dataset('BAHouses')
data = dataset.get_graph()

# Train a model ...

gcam = GradCAM(model, loss_function)
exp = gcam.get_explanation_node(data, node_idx)

feat_faith, node_faith, edge_faith = \
    graph_exp_faith(exp, data, model)

exp.visualize_node(show = True)
```

Fig. 3 Example use case of the GRAPHXAI package. An example of explaining a prediction in the GRAPHXAI package. With just a few lines of code, one can calculate an explanation for a node or graph, calculate metrics based on that explanation, and visualize the explanation.

Graph explanation accuracy (GEA). We report graph explanation accuracy as an evaluation strategy that measures an explanation's correctness using the ground-truth explanation \mathbf{M}^g . In ground-truth and predicted explanation masks, every node, node feature, or edge belongs to $\{0, 1\}$, where '0' means that an attribute is unimportant and '1' means that it is important for the model prediction. To measure accuracy, we use Jaccard index³⁴ between the ground-truth \mathbf{M}^g and predicted \mathbf{M}^p :

$$\text{JAC}(\mathbf{M}^g, \mathbf{M}^p) = \frac{\text{TP}(\mathbf{M}^g, \mathbf{M}^p)}{\text{TP}(\mathbf{M}^g, \mathbf{M}^p) + \text{FP}(\mathbf{M}^g, \mathbf{M}^p) + \text{FN}(\mathbf{M}^g, \mathbf{M}^p)}, \quad (1)$$

where TP denotes true positives, FP denotes false positives, and FN indicates false negatives. Most synthetic- and real-world graphs have multiple ground-truth explanations. For example, in the MUTAG dataset³⁵, carbon rings with both NH_2 or NO_2 chemical groups are valid explanations for the GNN model to recognize a given molecule as mutagenic. For this reason, the accuracy metric must account for the possibility of multiple equally valid explanations existing for any given prediction. Hence, we define ζ as a set of all possible ground-truth explanations, where $|\zeta|=1$ for graphs having a unique explanation. Therefore, we calculate GEA as:

$$\text{GEA}(\zeta, \mathbf{M}^p) = \max \text{JAC}(\mathbf{M}^g, \mathbf{M}^p) \quad \forall \mathbf{M}^g \in \zeta. \quad (2)$$

Here, we can calculate GEA using predicted node feature, node, or edge explanation masks. Finally, Eq. 1 quantifies the accuracy between the ground-truth and predicted explanation masks. Higher values mean a predicted explanation is more likely to match the ground-truth explanation.

Graph explanation faithfulness (GEF). We extend existing faithfulness metrics^{2,15} to quantify how faithful explanations are to an underlying GNN predictor. In particular, we obtain the prediction probability vector \hat{y}_u using the GNN, *i.e.*, $\hat{y}_u = f(\mathcal{S}_u)$, and using the explanation, *i.e.*, $\hat{y}_{u'} = f(\mathcal{S}_{u'})$, where we generate a masked subgraph $\mathcal{S}_{u'}$ by only keeping the original values of the top- k features identified by an explanation, and get their respective predictions $\hat{y}_{u'}$. Finally, we compute the graph explanation unfaithfulness metric as:

$$\text{GEF}(f, \mathcal{S}_u, \mathcal{S}_{u'}) = 1 - \exp^{-\text{KL}(f(\mathcal{S}_u) \| f(\mathcal{S}_{u'}))}, \quad (3)$$

where Kullback-Leibler (KL) divergence score quantifies the distance between two probability distributions, and the " $\|$ " operator indicates statistical divergence measure. Note that Eq. 3 is a measure of the unfaithfulness of the explanation. So, higher values indicate a higher degree of unfaithfulness.

Graph explanation stability (GES). Formally, an explanation is defined to be stable if the explanation for a given graph and its perturbed counterpart (generated by making infinitesimally small perturbations to the node feature vector and associated edges) are similar^{15,36}. We measure graph explanation stability w.r.t. the changes in the model behavior. In addition to similar output labels for \mathcal{S}_u and the perturbed $\mathcal{S}_{u'}$, we employ a second level of check where the difference between model behaviors for \mathcal{S}_u and $\mathcal{S}_{u'}$ is bounded by an infinitesimal constant δ , *i.e.*, $\|\mathcal{L}_{\mathcal{S}_u} - \mathcal{L}_{\mathcal{S}_{u'}}\|_p \leq \delta$. Here, $\mathcal{L}(\cdot)$ refers to any form of model knowledge like output logits \hat{y}_u or embeddings \mathbf{z}_u . We compute graph explanation instability as:

$$\text{GES}(\mathbf{M}_{\mathcal{S}_u}^p, \mathbf{M}_{\mathcal{S}_{u'}}^p) = \max D(\mathbf{M}_{\mathcal{S}_u}^p, \mathbf{M}_{\mathcal{S}_{u'}}^p), \quad \forall \mathcal{S}_{u'} \in \beta(\mathcal{S}_u) \quad (4)$$

where D represents the cosine distance metric, $\mathbf{M}_{\mathcal{S}_u}^p$ and $\mathbf{M}_{\mathcal{S}_{u'}}^p$ are the predicted explanation masks for \mathcal{S}_u and $\mathcal{S}_{u'}$, and β represents a δ -radius ball around \mathcal{S}_u for which the model behavior is same. Eq. 4 is a measure of instability, and higher values indicate a higher degree of instability.

Counterfactual fairness mismatch. To measure counterfactual fairness¹⁵, we verify if the explanations corresponding to \mathcal{S}_u and its counterfactual counterpart (where the protected node feature is modified) are similar (dissimilar) if the underlying model predictions are similar (dissimilar). We calculate counterfactual fairness mismatch as:

$$\text{GECF}(\mathbf{M}^p, \mathbf{M}_s^p) = D(\mathbf{M}^p, \mathbf{M}_s^p), \quad (5)$$

where, \mathbf{M}^p and \mathbf{M}_s^p are the predicted explanation mask for \mathcal{S}_u and for the counterfactual counterpart of $\mathcal{S}_{u'}$. Note that we expect a lower GECF score for graphs having weakly-unfair ground-truth explanations because explanations are similar for both original and counterfactual graphs, whereas, for graphs with strongly-unfair ground-truth explanations, we expect a higher GECF score as explanations change when we modify the protected attribute.

Group fairness mismatch. We measure group fairness mismatch¹⁵ as:

$$\text{GEGF}(\hat{y}_{\mathcal{K}}, \hat{y}_{\mathcal{K}}^{E_u}) = \left| \text{SP}(\hat{y}_{\mathcal{K}}) - \text{SP}(\hat{y}_{\mathcal{K}}^{E_u}) \right|, \quad (6)$$

where $\hat{y}_{\mathcal{K}}$ and $\hat{y}_{\mathcal{K}}^{E_u}$ are predictions for a set of \mathcal{K} graphs using the original and the essential features identified by an explanation, respectively, and SP is the statistical parity. Finally, Eq. 6 is a measure of group fairness mismatch of an explanation where higher values indicate that the explanation is not preserving group fairness.

Evaluation and analysis of GNN explainability methods. Next, we discuss experimental results that answer critical questions concerning synthetic and real-world graphs and different ground-truth explanations.

Benchmarking GNN explainers on synthetic and real-world graphs. We evaluate the performance of GNN explainers on a collection of synthetically generated graphs with various properties and molecular datasets using metrics described in the experimental setup. Results in Tables 1–5 show that, while no explanation method performs well across all properties, across different SHAPEGEN node classification datasets (Table 6), SubgraphX outperforms other methods on average. In particular, SubgraphX generates 145.95% more accurate and 64.80% less unfaithful explanations than other GNN explanation methods. Gradient-based methods, such as GradCAM and GuidedBP, perform the next best of all methods, with Grad producing the second-lowest unfaithfulness score and GradCAM achieving the second-highest explanation accuracy score. PGExplainer generates the least unstable explanations—35.35% less unstable explanations than the average instability across other GNN explainers. In summary, results of Tables 1–5 show that node explanation masks are more reliable than edge- and node feature explanation masks and state-of-the-art GNN explainers achieve better faithfulness for synthetic graph datasets as compared to real-world graphs.

Analyzing homophilic vs. heterophilic ground-truth explanations. We compare GNN explainers by generating explanations on GNN models trained on homophilic and heterophilic graphs generated using the SG-HETEROPHILIC generator. Then, we compute the graph explanation unfaithfulness scores of output explanations generated using state-of-the-art GNN explainers. We find that GNN explainers produce 55.98% more faithful explanations when ground-truth explanations are homophilic than when ground-truth explanations are heterophilic (*i.e.*, low unfaithfulness scores for light green bars in Fig. 4). These results reveal an important gap in existing GNN explainers. Namely, existing GNN explainers fail to perform well on diverse graph types, such as homophilic, heterophilic and attributed graphs. This observation, enabled by the flexibility of SHAPEGEN generator, highlights an opportunity for future algorithmic innovation in GNN explainability.

Analyzing the reliability of graph explainers to smaller vs. larger ground-truth explanations. Next, we examine the reliability of GNN explainers when used to predict explanations for models trained on graphs generated using the SG-SMALLEX graph generator. Results in Fig. 5 show that explanations from existing GNN explainers are faithful (*i.e.*, lower GEF scores) to the underlying GNN models when ground-truth explanations are smaller, *i.e.*, $S = \text{'triangle'}$. On average, across all eight GNN explainers, we find that existing GNN explainers are highly unfaithful to graphs with large ground-truth explanations with an average GEF score of 0.7476. Further, we observe that explanations generated on ‘triangular’ (smaller) ground-truth explanations are 59.98% less unfaithful than explanations for ‘house’ (larger) ground-truth explanations (*i.e.*, low unfaithfulness scores for light purple bars in Fig. 5). However, the Grad explainer, on average, achieves 9.33% lower unfaithfulness on large ground-truth explanations compared to other explanation methods. This limited behavior of existing GNN explainers has not been previously known and highlights an urgent need for additional analysis of GNN explainers.

Examining fair vs. unfair ground-truth explanations. To measure the fairness of predicted explanations, we train GNN models on SG-Unfair, which generates graphs with controllable fairness properties. Next, we compute the average GECF and GEGF values for predicted explanations from eight GNN explainers. The fairness results in Fig. 6 show that GNN explainers do not preserve counterfactual fairness and are highly prone to producing unfair explanations. We note that for weakly-unfair ground-truth explanations (light red in Fig. 6), explanations M^p should not change as the label-generating process is independent of the protected attribute. Still, we observe high GECF scores for most explanation methods. For strongly-unfair ground-truth explanations, we find that explanations from most GNN explainers fail to capture (*i.e.*, low GECF scores for dark red bars in Fig. 6) the unfairness enforced using the protected attribute and generate similar explanations even when we flip/modify the respective protected attribute. We see that GradCAM and PGEx explanations outperform other GNN explainers in preserving counterfactual explanations for weakly-unfair ground-truth explanations. In contrast, the PGMEx explainer preserves counterfactual fairness better than other explanation methods on strongly-unfair ground truth explanations. Our results highlight the importance of studying fairness in XAI as they can enhance a user’s confidence in the model and assist in detecting and correcting unwanted bias.

Faithfulness shift with varying degrees of node feature information. Using SHAPEGEN’s support for node features and ground-truth explanations on those features, we evaluate explainers that generate explanations for node features. Results for node feature explanations on SG-BASE are given in Table 4. In addition, we explore the performance of explainers under varying proportions of informative node features. Informative node features, defined in the SHAPEGEN construction (Algorithm 1), are node features correlated with the label of a given node, as opposed to redundant features, which are sampled randomly from a Gaussian distribution. Figure 7 shows the results of experiments on three datasets, SG-MOREINFORM, SG-BASE, and SG-LESSINFORM. All datasets have similar graph topology, but SG-MOREINFORM has a higher proportion of informative features while SG-LESSINFORM has a lower proportion of these informative features. SG-BASE is used as a baseline with a proportion of informative features greater than SG-LESSINFORM but less than SG-MOREINFORM. There are minimal differences between explainers’ faithfulness across datasets, however, unfaithfulness tends to increase with fewer informative node features. As seen in Table 4, the Gradient explainer shows the best faithfulness score across all datasets for node feature explanation. Still, this faithfulness is relatively weak, only 0.001 better than

Method	GEA (\uparrow)	GEF (\downarrow)	GES (\downarrow)	GECF (\downarrow)	GEGF (\downarrow)
Random	0.148 \pm 0.002	0.579 \pm 0.007	0.920 \pm 0.002	0.763 \pm 0.003	0.023 \pm 0.002
Grad	0.193 \pm 0.002	0.392 \pm 0.006	0.806 \pm 0.004	0.159 \pm 0.004	0.039 \pm 0.003
GradCAM	0.222 \pm 0.002	0.452 \pm 0.006	0.263 \pm 0.004	0.010 \pm 0.001	0.020 \pm 0.002
GuidedBP	0.194 \pm 0.001	0.557 \pm 0.007	0.432 \pm 0.004	0.067 \pm 0.002	0.021 \pm 0.002
IG	0.142 \pm 0.002	0.545 \pm 0.007	0.727 \pm 0.005	0.110 \pm 0.003	0.021 \pm 0.002
GNNExplainer	0.102 \pm 0.003	0.534 \pm 0.007	0.431 \pm 0.008	0.233 \pm 0.006	0.027 \pm 0.002
PGMExplainer	0.133 \pm 0.002	0.541 \pm 0.007	0.984 \pm 0.001	0.791 \pm 0.003	0.096 \pm 0.004
PGExplainer	0.194 \pm 0.002	0.557 \pm 0.007	0.217 \pm 0.004	0.009 \pm 0.000	0.029 \pm 0.002
SubgraphX	0.324 \pm 0.004	0.254 \pm 0.006	0.745 \pm 0.005	0.241 \pm 0.006	0.035 \pm 0.003

Table 1. Evaluation of GNN explainers on SG-Base graph dataset based on node explanation masks M_N^p . Arrows (\uparrow/\downarrow) indicate the direction of better performance. SubgraphX far outperforms other methods in accuracy and faithfulness while PGExplainer is best for stability and counterfactual fairness. In general, gradient methods produce the most fair explanations across both counterfactual and group fairness metrics. See Tables 3–4 for results on edge and feature explanation masks.

Method	GEA (\uparrow)	GEF (\downarrow)	GES (\downarrow)	GECF (\downarrow)	GEGF (\downarrow)
Random	0.075 \pm 0.002	0.638 \pm 0.007	1.55 \pm 0.004	1.01 \pm 0.010	0.027 \pm 0.002
Grad	0.194 \pm 0.002	0.498 \pm 0.007	0.745 \pm 0.005	0.157 \pm 0.004	0.068 \pm 0.003
GradCAM	0.188 \pm 0.001	0.620 \pm 0.006	0.295 \pm 0.005	0.029 \pm 0.003	0.027 \pm 0.002
GuidedBP	0.190 \pm 0.001	0.653 \pm 0.007	0.430 \pm 0.004	0.074 \pm 0.003	0.020 \pm 0.002
IG	0.140 \pm 0.002	0.672 \pm 0.007	0.639 \pm 0.004	0.114 \pm 0.004	0.011 \pm 0.001
GNNExplainer	0.103 \pm 0.003	0.632 \pm 0.007	0.431 \pm 0.008	0.249 \pm 0.007	0.028 \pm 0.002
PGMExplainer	0.133 \pm 0.002	0.622 \pm 0.007	0.974 \pm 0.001	0.798 \pm 0.003	0.083 \pm 0.003
PGExplainer	0.165 \pm 0.002	0.635 \pm 0.007	0.224 \pm 0.004	0.005 \pm 0.000	0.030 \pm 0.002
SubgraphX	0.383 \pm 0.004	0.344 \pm 0.006	0.585 \pm 0.004	0.225 \pm 0.004	0.114 \pm 0.004

Table 2. Evaluation of GNN explainers on SG-Base graph dataset based on node explanation masks M_N^p . Base GNN is a GCN³³ as opposed to Table 1 which is based on explaining a GIN model³². Overall, explainer performance is very similar to that of the GIN with SubgraphX performing the best on faithfulness and accuracy metrics while gradient-based methods and PGExplainer typically perform best for fairness and stability.

Method	GEA (\uparrow)	GEF (\downarrow)	GES (\downarrow)	GECF (\downarrow)	GEGF (\downarrow)
Random	0.135 \pm 0.001	0.419 \pm 0.007	1.167 \pm 0.001	0.997 \pm 0.002	0.064 \pm 0.003
GNNExplainer	0.152 \pm 0.002	0.302 \pm 0.006	0.995 \pm 0.001	0.957 \pm 0.002	0.047 \pm 0.003
PGExplainer	0.117 \pm 0.002	0.171 \pm 0.004	1.000 \pm 0.000	1.000 \pm 0.000	0.037 \pm 0.003
SubgraphX	0.271 \pm 0.003	0.548 \pm 0.007	0.815 \pm 0.005	0.290 \pm 0.007	0.030 \pm 0.002

Table 3. Evaluation of GNN explainers on SG-Base graph dataset based on edge explanation masks M_E^p . Arrows (\uparrow/\downarrow) indicate the direction of better performance. SubgraphX method, on average, produces the most reliable edge explanations when evaluated across all five performance metrics. Note that of the explainers tested in this study, only the above four methods produce edge explanations.

Method	GEA (\uparrow)	GEF (\downarrow)	GES (\downarrow)	GECF (\downarrow)	GEGF (\downarrow)
Random	0.281 \pm 0.003	0.016 \pm 0.001	0.997 \pm 0.001	0.810 \pm 0.005	0.023 \pm 0.002
Grad	0.306 \pm 0.002	0.015 \pm 0.001	0.925 \pm 0.003	0.259 \pm 0.006	0.027 \pm 0.003
GuidedBP	0.240 \pm 0.003	0.016 \pm 0.001	0.899 \pm 0.004	0.275 \pm 0.006	0.025 \pm 0.002
IG	0.278 \pm 0.003	0.016 \pm 0.001	0.917 \pm 0.004	0.119 \pm 0.004	0.022 \pm 0.002
GNNExplainer	0.281 \pm 0.003	0.017 \pm 0.001	0.999 \pm 0.001	0.826 \pm 0.005	0.023 \pm 0.003

Table 4. Evaluation of GNN explainers on SG-Base graph dataset based on node feature explanation masks M_F^p . Arrows (\uparrow/\downarrow) indicate the direction of better performance. All GNN explainers produce highly faithful node feature explanations. However, the stability of these methods on node features is more similar to random explanations than is observed for node explanations in Table 1 and Table 2. Note that of the explainers tested in this study, only the above five methods produce node feature explanations.

Dataset	Method	GEA (\uparrow)	GEF (\downarrow)
MUTAG	Random	0.044 \pm 0.007	0.590 \pm 0.031
	Grad	0.022 \pm 0.006	0.598 \pm 0.030
	GradCAM	0.085 \pm 0.012	0.672 \pm 0.029
	GuidedBP	0.036 \pm 0.007	0.649 \pm 0.030
	Integrated Grad (IG)	0.049 \pm 0.010	0.443 \pm 0.031
	GNNExplainer	0.031 \pm 0.005	0.618 \pm 0.030
	PGMExplainer	0.042 \pm 0.007	0.503 \pm 0.031
	PGExplainer	0.046 \pm 0.007	0.504 \pm 0.031
BENZENE	SubgraphX	0.039 \pm 0.007	0.611 \pm 0.030
	Random	0.108 \pm 0.003	0.513 \pm 0.012
	Grad	0.122 \pm 0.007	0.262 \pm 0.011
	GradCAM	0.291 \pm 0.007	0.551 \pm 0.012
	GuidedBP	0.205 \pm 0.007	0.438 \pm 0.012
	Integrated Grad (IG)	0.044 \pm 0.003	0.182 \pm 0.010
	GNNExplainer	0.129 \pm 0.005	0.444 \pm 0.012
	PGMExplainer	0.154 \pm 0.006	0.433 \pm 0.012
Fl-CARBONYL	PGExplainer	0.169 \pm 0.007	0.375 \pm 0.012
	SubgraphX	0.371 \pm 0.009	0.513 \pm 0.012
	Random	0.087 \pm 0.007	0.440 \pm 0.26
	Grad	0.132 \pm 0.010	0.210 \pm 0.021
	GradCAM	0.005 \pm 0.007	0.500 \pm 0.026
	GuidedBP	0.089 \pm 0.010	0.315 \pm 0.024
	Integrated Grad (IG)	0.091 \pm 0.007	0.174 \pm 0.019
	GNNExplainer	0.094 \pm 0.009	0.423 \pm 0.026
	PGMExplainer	0.078 \pm 0.008	0.426 \pm 0.026
	PGExplainer	0.079 \pm 0.009	0.372 \pm 0.025
	SubgraphX	0.008 \pm 0.002	0.466 \pm 0.026

Table 5. Evaluation of GNN explainers for real-world molecular datasets with ground-truth explanations. Arrows (\uparrow/\downarrow) indicate the direction of better performance. Integrated Gradient explanations obtain the lowest unfaithfulness score across all three datasets. Note that stability and fairness do not apply to these datasets because generating plausible perturbations for molecules is non-trivial, and they do not contain protected features.

Dataset	SG-BASE	SG-HETEROPHILIC	SG-SMALLEX	SG-UNFAIR
Nodes	13150	13150	15505	13150
Edges	46472	46472	51782	46472
Node features	11	11	11	11
Average node degree	3.53 \pm 0.02	3.53 \pm 0.02	3.34 \pm 0.02	3.53 \pm 0.02
Class 0 Nodes	4382	4382	7777	4382
Class 1 Nodes	8768	8768	7728	8768
Shape of the planted motif (S)	'house'	'house'	'triangle'	'house'
Number of initial subgraphs (N_s)	1200	1200	1300	1200
Probability of subgraph connection (p)	0.006	0.006	0.006	0.006
Subgraph size (n_s)	11	11	12	11
Number of classes (K)	2	2	2	2
Number of node features (n_f)	11	11	11	11
Number of informative features (n_i)	4	4	4	4
Class separation factor (s_f)	0.6	0.6	0.5	0.6
Number of clusters per class (c_f)	2	2	2	2
Protected feature noise factor (ϕ)	0.5	0.5	0.5	0.75
Homophily coefficient (η)	1	-1	1	1
Number of GNN layers (L)	3	3	3	3

Table 6. Statistics of graphs generated using SHAPEGGEN data generator for evaluating different properties of GNN explanations.

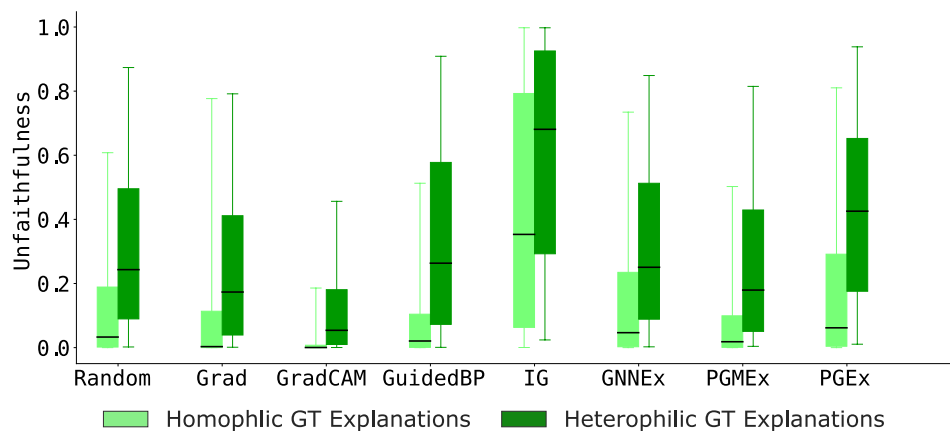


Fig. 4 Unfaithfulness scores across eight GNN explainers on SG-HETEROPHILIC graph dataset consisting of either homophilic or heterophilic ground-truth (GT) explanations. GNN explainers produce more faithful explanations (lower GEF scores) on homophilic graphs than heterophilic graphs, revealing an important limitation of existing GNN explainers.

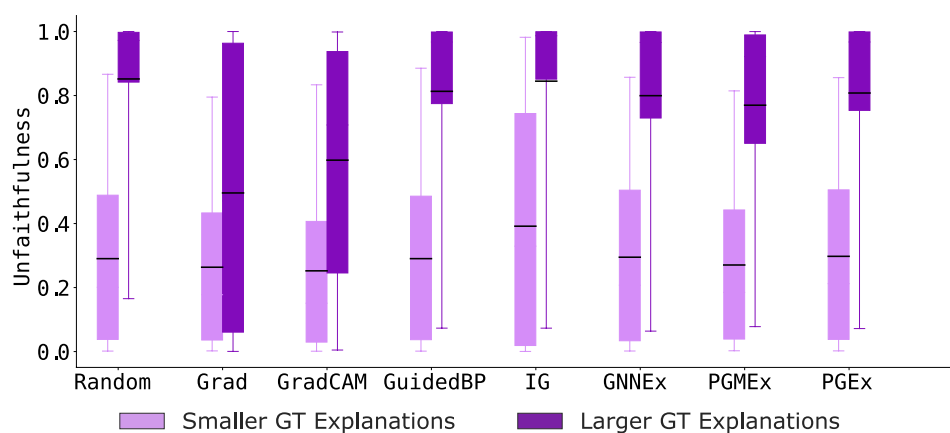


Fig. 5 Unfaithfulness scores across eight GNN explainers on SG-SMALLEX graph dataset with smaller (triangle shapes) or (house shapes) ground-truth (GT) explanations. Results show that GNN explainers produce more faithful explanations (lower GEF scores) on graphs with smaller GT explanations than on graphs with larger GT explanations.




Fig. 6 Counterfactual fairness mismatch scores across eight GNN explainers on SG-Unfair graph dataset with weakly-unfair or strongly-unfair ground-truth (GT) explanations. Results show that explanations produced on graphs with strongly-unfair ground-truth explanations do not preserve fairness and are sensitive to flipping/modifying the protected node feature.

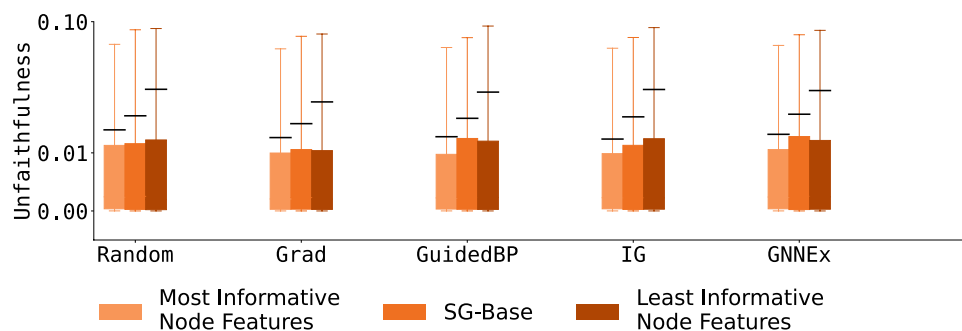


Fig. 7 Unfaithfulness scores across five GNN explainers that produce node feature explanations. Every GNN explainer is evaluated on three datasets whose network topology is equivalent to SG-Base and by varying the ratio between informative and redundant node features: most informative node features, control node features, and least informative node features. Results show that across all explainers, unfaithfulness decreases as the proportion of informative to redundant features increases, with explainers trained on the graph with the most informative node features having consistently lower unfaithfulness scores than explainers trained on graphs with the least informative node features.




Fig. 8 Visualization of four explainers from the G-XAI Bench library on the BA-Shapes dataset. The visualization is for explaining the prediction of node u . We show the $L + 1$ -hop around node u , where L is the number of layers of the GNN model predicting on the dataset. Two color bars indicate the intensity of attribution scores for the node and edge explanations. Note that edge importance is not defined for every method, so edges are set to black to indicate that the method does not provide edge scores. Visualization tools are a native part of the GRAPHXAI package, including user-friendly functions `graphxai.Explanation.visualize_node` and `graphxai.Explanation.visualize_graph` to visualize GNN explanations. The visualization tools in GRAPHXAI allow users to compare the explanations of different GNN explainers, such as gradient-based methods (Gradient and Grad-CAM) and perturbation-based methods (GNNExplainer and SubgraphX).

random explanation. These results show that the faithfulness of node feature explanations worsens under sparse node feature signals.

Visualization results. GRAPHXAI provides functions that visualize explanations produced by GNN explainability methods. Users can compare both node- and graph-level explanations. In addition, function implementations for visualization are parameterized, allowing users to change colors and weight interpretation. Functions are compatible with `matplotlib`³⁷ and `networkx`³⁸. Visualizations are generated by `graphxai.Explanation.visualize_node` for node-level explanations and `graphxai.Explanation.visualize_graph` functions for graph-level explanations. In Fig. 8, we show the output explanation from four different GNN explainers as produced by our visualization function. Figure 9 shows example outputs from multiple explainers in the GRAPHXAI package on a SHAPEGGEN-generated dataset.

Discussion

GRAPHXAI provides a general-purpose framework to evaluate GNN explanations produced by state-of-the-art GNN explanation methods. GRAPHXAI provides data loaders, data processing functions, visualizers, real-world graph datasets with ground-truth explanations, and evaluation metrics to benchmark the quality of GNN explanations. GRAPHXAI introduces a novel and flexible synthetic dataset generator called SHAPEGGEN to automatically generate benchmark datasets and corresponding ground truth explanations robust against known pitfalls of GNN explainability methods. Our experimental results show that existing GNN explainers perform well on graphs with homophilic ground-truth explanations but perform considerably worse on heterophilic and

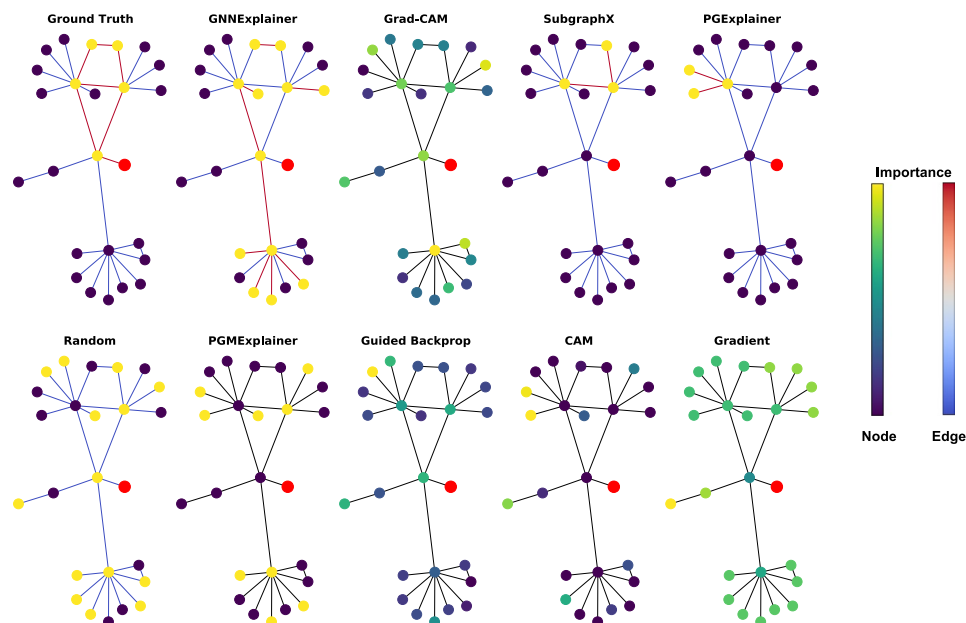


Fig. 9 Example of a particularly challenging example in a SHAPEGGEN dataset. All explanation methods that output node-wise importance scores are shown, including the ground-truth explanation at the top of the figure. Importance and edge scores are highlighted by relative value across each explanation method, as shown by the scales at right in the figure. The central node, *i.e.*, the node being classified in this example, is shown in red on each subgraph. Visualizations are generated by `graphxai.Explanation.visualize_node`, a function native to the `graphxai` package. Some explainers can capture portions of the ground-truth explanation, such as SubgraphX and GNNExplainer, but others attribute no importance to the ground-truth shape, such as CAM and Gradient.

attributed graphs. Across multiple graph datasets and types of downstream prediction tasks, we show that existing GNN explanation methods fail on graphs with larger ground-truth explanations and cannot generate explanations that preserve the fairness properties of the underlying GNN model. In addition, GNN explainers tend to underperform on sparse node feature signals compared to more densely informative node features. These findings indicate the need for methodological innovation and a thorough analysis of future GNN explainability methods across performance dimensions.

GRAPHXAI provides a flexible framework for evaluating GNN explanation methods and promotes reproducible and transparent research. We maintain GRAPHXAI as a centralized library for evaluating GNN explanation methods and plan to add newer datasets, explanation methods, diverse evaluation metrics, and visualization features to our existing framework. In the current version of GRAPHXAI, we mostly employ real-world molecular chemistry datasets as the need for model understanding is motivated by experimental evaluation of model predictions in the laboratory, and it includes a wide variety of graph sizes (ranging from 1,768 to 12,000 instances in the dataset), node feature dimensions (ranging from 13 to 27 dimensions), and class imbalance ratios. In addition to the scale-free SHAPEGGEN dataset generator, we will include other random graph model generators in the next GRAPHXAI release to support benchmarking of GNN explainability methods on other graph types. Our evaluation metrics can be also extended to explanations from self-explaining GNNs, *e.g.*, GraphMASK¹² identifies edges at each layer of the GNN during training that can be ignored without affecting the output model predictions. In general, self-explaining GNNs also return a set of edge masks as an output explanation for a GNN prediction that can be converted to edge importance scores for computing GraphXAI metrics. We anticipate that GRAPHXAI can help algorithm developers and practitioners in graph representation learning develop and evaluate principled explainability techniques.

Methods

SHAPEGGEN is a key component of GRAPHXAI and serves as a synthetic dataset generator of XAI-ready graph datasets. It is founded in graph theory and designed to address the pitfalls (see Introduction) of existing graph datasets in the broad area of explainable AI. As such, SHAPEGGEN can facilitate the development, analysis, and evaluation of GNN explainability methods (see Results). We proceed with the description of SHAPEGGEN data generator.

Notation. *Graphs.* Let $\mathcal{G} = (\mathcal{V}_{\mathcal{G}}, \mathcal{E}_{\mathcal{G}}, \mathbf{X}_{\mathcal{G}})$ denote an undirected graph comprising of a set of nodes $\mathcal{V}_{\mathcal{G}}$ and a set of edges $\mathcal{E}_{\mathcal{G}}$. Let $\mathbf{X}_{\mathcal{G}} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$ denote the set of node feature vectors for all nodes in $\mathcal{V}_{\mathcal{G}}$, where $\mathbf{x}_v \in \mathbf{X}_{\mathcal{G}}$ is an d -dimensional vector which captures the attribute values of a node v and $N = |\mathcal{V}_{\mathcal{G}}|$ denotes the number of

nodes in the graph. Let $\mathbf{A} \in \mathbb{R}^{N \times N}$ be the graph adjacency matrix where element $\mathbf{A}_{uv} = 1$ if there exists an edge $e \in \mathcal{E}_{\mathcal{G}}$ between nodes u and v and $\mathbf{A}_{uv} = 0$ otherwise. We use \mathcal{N}_u to denote the set of immediate neighbors of node u , $\mathcal{N}_u = \{v \in \mathcal{V}_{\mathcal{G}} | \mathbf{A}_{uv} = 1\}$. Finally, the function $\text{deg}: \mathcal{V}_{\mathcal{G}} \mapsto \mathbb{Z}_{\geq 0}$ is defined as $\text{deg}(u) = |\mathcal{N}_u|$ and outputs the degree of a node $u \in \mathcal{V}_{\mathcal{G}}$.

Graph neural networks. Most GNNs can be formulated as message passing networks³⁹ using three operators: Msg, Agg, and Upd. In a L -layer GNN, these operators are recursively applied on \mathcal{G} , specifying how neural messages (*i.e.* embeddings) are exchanged between nodes, aggregated, and transformed to arrive at node representations in the last layer of transformations. Formally, a message between a pair of nodes (u, v) in layer l is defined as a function of hidden representations of nodes \mathbf{h}_u^{l-1} and \mathbf{h}_v^{l-1} from the previous layer: $\mathbf{m}_{uv}^l = \text{Msg}(\mathbf{h}_u^{l-1}, \mathbf{h}_v^{l-1})$. In Agg, messages from all nodes $v \in \mathcal{N}_u$ are aggregated as: $\mathbf{m}_u^l = \text{Agg}(\mathbf{m}_{uv}^l | v \in \mathcal{N}_u)$. In Upd, the aggregated message \mathbf{m}_u^l is combined with \mathbf{h}_u^{l-1} to produce u 's representation for layer l as $\mathbf{h}_u^l = \text{Upd}(\mathbf{m}_u^l, \mathbf{h}_u^{l-1})$. Final node representation $\mathbf{z}_u = \mathbf{h}_u^L$ is the output of the last layer. Lastly, let f denote a downstream GNN classification model that maps the node representation \mathbf{z}_u to a softmax prediction vector $\hat{y}_u \in \mathbb{R}^C$, where C is the total number of labels.

GNN explainability methods. Given the prediction $f(u)$ for node u made by a GNN model, a GNN explainer \mathcal{O} outputs an explanation mask \mathbf{M}^p that provides an explanation of $f(u)$. These explanations can be given with respect to node attributes $\mathbf{M}_a \in \mathbb{R}^d$, nodes $\mathbf{M}_n \in \mathbb{R}^N$, or edges $\mathbf{M}_e \in \mathbb{R}^{N \times N}$, depending on specific GNN explainer, such as GNNExplainer¹⁴, PGExplainer¹⁰, and SubgraphX³¹. For all explanation methods, we use a graph masking function MASK that outputs a new graph $\mathcal{G}' = (\mathcal{V}_{\mathcal{G}'}, \mathcal{E}'_{\mathcal{G}'}, \mathbf{X}'_{\mathcal{G}'})$ by performing element-wise multiplication operation between the masks $(\mathbf{M}_a, \mathbf{M}_n, \mathbf{M}_e)$ and their respective attributes in the original graph \mathcal{G} , *e.g.* $\mathbf{A}' = \mathbf{A} \odot \mathbf{M}_e$. Finally, we denote the ground-truth explanation mask as \mathbf{M}^g that is used to evaluate the performance of GNN explainers.

SHAPEGGEN dataset generator. We propose a novel and flexible synthetic dataset generator called SHAPEGGEN that can automatically generate a variety of benchmark datasets (*e.g.*, varying graph sizes, degree distributions, homophilic vs. heterophilic graphs) accompanied by ground-truth explanations. Furthermore, the flexibility to generate diverse synthetic datasets and corresponding ground-truth explanations allows us to mimic the data generated by various real-world applications. SHAPEGGEN is a generator of XAI-ready graph datasets supported by graph theory and is particularly suitable for benchmarking GNN explainers and studying their limitations.

Flexible parameterization of SHAPEGGEN. SHAPEGGEN has tunable parameters for data generation. By varying these parameters, SHAPEGGEN can generate diverse types of graphs, including graphs with varying degrees of class imbalance and graph sizes. Formally, a graph is generated as: $\mathcal{G} = \text{SHAPEGGEN}(\mathcal{S}, N_s, p, n_s, K, n_f, n_i, s_f, c_f, \varphi, \eta, L)$, where:

- \mathcal{S} is the motif, defined as a subgraph, to be planted within the graph.
- N_s denotes the number of subgraphs used in the initial graph generation process.
- p represents the connection probability between two subgraphs and controls the average shortest path length for all possible pairs of nodes. Ideally, a smaller p value for larger N_s is preferred to avoid low average path length and, therefore, the poor performance of GNNs.
- n_s is the expected size of each subgraph in the SHAPEGGEN generation procedure. For a fixed \mathcal{S} shape, large n_s values produce graphs with long-tailed degree distributions. Note that the expected total number of nodes in the generated graph is $N \times n_s$.
- K is the number of distinct classes defined in the graph downstream task.
- n_f represents the number of dimensions for node features in the generated graph.
- n_i is the number of informative node features. These are features correlated to the node labels instead of randomly generated non-informative features. The indices for the informative features define the ground-truth explanation mask for node features in the final SHAPEGGEN instance. In general, larger n_i results in an easier classification and explanation task, as it increases the node feature-level ground-truth explanation size.
- s_f is defined as the class separation factor that represents the strength of discrimination of class labels between features for each node. Higher s_f corresponds to a stronger signal in the node features, *i.e.*, if a classifier is trained only on the node features, a higher s_f value would result in an easier classification task.
- c_f is the number of clusters per class. A larger c_f value increases the difficulty of the classification task with respect to node features.
- $\varphi \in [0, 1]$ is the protected feature noise factor that controls the strength of correlation r between the protected features and the node labels. This value corresponds to the probability of “flipping” the protected feature with respect to the node’s label. For instance, $\varphi = 0.5$ results in zero correlation ($r = 0$) between the protected feature and the label (*i.e.* complete fairness), $\varphi = 0$ results in a positive correlation ($r = 1$), and $\varphi = 1$ results in a negative correlation ($r = -1$) between the label and the protected feature.
- η is the homophily coefficient that controls the strength of homophily or heterophily in the graph. Positive values ($\eta > 0$) produce a homophilic graph while negative values ($\eta < 0$) produce a heterophilic graph.
- L is the number of layers in the GNN predictor corresponding to the GNN’s receptive field. For the purposes of ShapeGGen, L is used to define the size of the GNN’s receptive field and thus the size of the ground-truth explanation generated for each node.

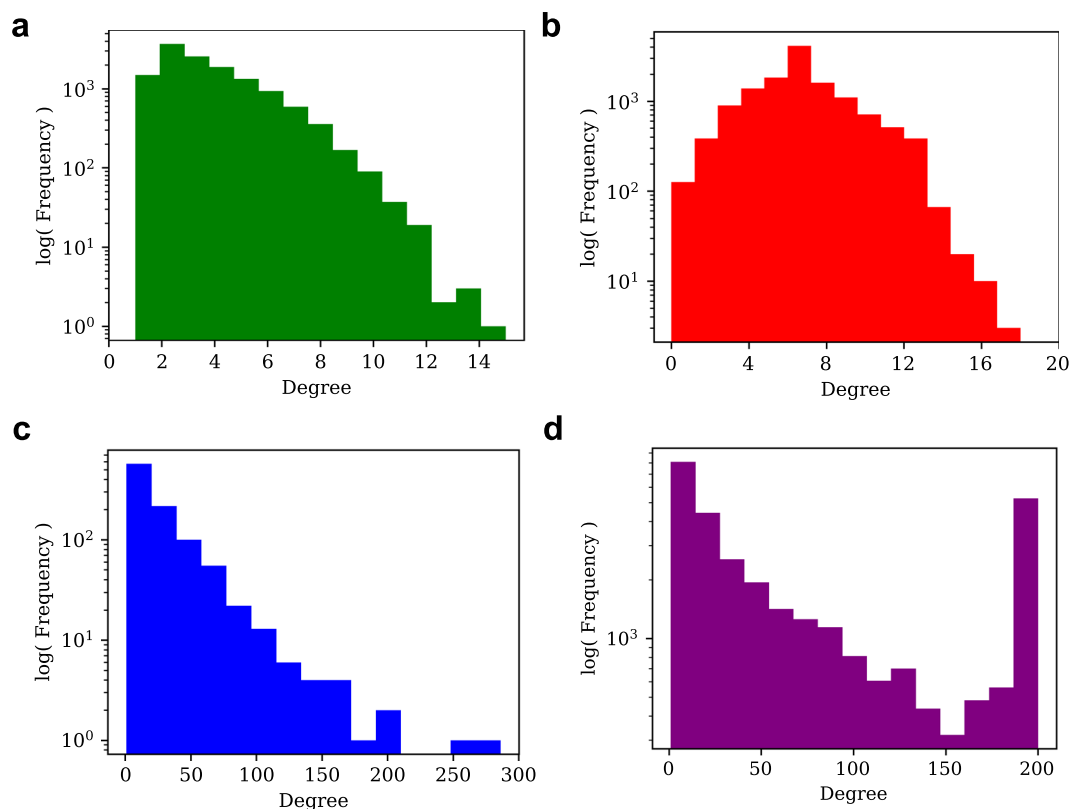


Fig. 10 Comparison of degree distribution for (a) SHAPEGGEN dataset (SG-Base), (b) random Erdős-Rényi graph ($p = 5 \times 10^{-4}$) graph, (c) German Credit dataset, and (d) Credit Defaulter dataset. All plots are shown with a log scale of frequency for the y-axis. SG-Base and both real-world graphs show a power-law degree distribution commonly observed in real-world datasets exhibiting preferential attachment properties. Datasets generated by SHAPEGGEN are designed to present power-law degree distributions to match real-world dataset topologies, such as those observed in German Credit and Credit Defaulter. The degree distribution of SG-Base is much different than the binomial distribution exhibited in Erdős-Rényi graph (b), an unstructured random graph model.

This wide array of parameters for SHAPEGGEN allows for the generation of graph instances with vastly differing properties Fig. 10.

Generating graph structure. Figure 2 summarizes the process to generate a graph $\mathcal{G} = (\mathcal{V}_{\mathcal{G}}, \mathcal{E}_{\mathcal{G}}, \mathbf{X}_{\mathcal{G}})$. SHAPEGGEN generates N_s subgraphs that exhibit the preferential attachment property⁴⁰, which occurs in many real-world graphs. Each subgraph is first given a motif, *i.e.*, a subgraph $\mathcal{S} = (\mathcal{V}_{\mathcal{S}}, \mathcal{E}_{\mathcal{S}}, \mathbf{X}_{\mathcal{S}})$. A preferential attachment algorithm is then performed on base structure \mathcal{S} , adding n' ($n' \sim \text{Poisson}(\lambda = n_s - |\mathcal{V}_{\mathcal{S}}|)$) nodes by creating edges to nodes in $\mathcal{V}_{\mathcal{S}}$. The Poisson distribution is used for determining the sizes of each subgraph used in the generation process, with $\lambda = n_s - |\mathcal{V}_{\mathcal{S}}|$, the difference between the number of nodes in the motif and the expected subgraph size n_s . After creating a list of randomly-generated subgraphs $\mathbf{S} = \{\mathcal{S}_1, \dots, \mathcal{S}_{N_s}\}$, edges are created to connect subgraphs, creating the structure of an instance of SHAPEGGEN. Subgraph connections and local subgraph construction is performed in such a way that each node in the final graph \mathcal{G} has between 1 and K motifs within its neighborhood, *i.e.*, $|\bigcup_{i=1}^N \mathcal{V}_{\mathcal{S}_i} \cap \mathcal{N}_v| \in \{1, 2, \dots, K\}$ for any v and \mathcal{S}_i . This naturally defines a classification task in the domain of f to $\{0, 1, \dots, K-1\}$. More details on the SHAPEGGEN structure generation can be found in Algorithm 2.

Generating labels for prediction. A motif defined as a subgraph $\mathcal{S} = (\mathcal{V}_{\mathcal{S}}, \mathcal{E}_{\mathcal{S}}, \mathbf{X}_{\mathcal{S}})$ occurs randomly throughout \mathcal{G} , with the set $\mathbf{S} = \{\mathcal{S}_1, \dots, \mathcal{S}_{N_s}\}$. The task on this graph is a motif detection problem, where each node has exactly 1, 2, or K motifs in its 1-hop neighborhood. A motif \mathcal{S}_i is considered to be within the neighborhood of a node $v \in \mathcal{V}_{\mathcal{G}}$ if any node $s \in \mathcal{V}_{\mathcal{S}_i}$ is also in the neighborhood of v , *i.e.*, if $|\mathcal{V}_{\mathcal{S}_i} \cap \mathcal{N}_v| > 0$. Therefore, the task that a GNN predictor needs to solve is defined by:

$$f(v \in \mathcal{V}_{\mathcal{G}}) = \sum_{\mathcal{S}_i \in \mathbf{S}} \mathbf{1}_{\mathcal{V}_{\mathcal{S}_i}(\mathcal{N}_v)} - 1, \quad (7)$$

where $\mathbb{1}_{\mathcal{V}_i}(\mathcal{N}_v) = 0$ if $|\mathcal{V}_i \cap \mathcal{N}_v| = 0$ and 1 otherwise.

Generating node feature vectors. SHAPEGGEN uses a latent variable model to create node feature vectors and associate them with network structure. This latent variable model is based on that developed by Guyon⁴¹ for the MADEON dataset and implemented in Scikit-Learn's `make_classification` function⁴². The latent variable model creates n_i informative features for each node based on the node's generated label and also creates non-informative features as noise. Having non-informative/redundant features allows us to evaluate GNN explainers, such as GNNExplainer¹⁴, that formulate explanations as node feature masks. More detail on node feature generation is given in Algorithm 1.

SHAPEGGEN can generate graphs with both homophilic and heterophilic ground-truth explanations. We optimize between homophily vs. heterophily by taking advantage of redundant node features, *i.e.*, features that do not associate with the generated labels, and manipulate them appropriately based on a user-specified homophily parameter η . The magnitude of the η parameter determines the degree of homophily/heterophily in the generated graph. The algorithm for node features is given in Algorithm 3. While not every node feature in the feature vector is optimized for homophily/heterophily indication, we experimentally verified the cosine similarity between node feature vectors produced by Algorithm 3 reveals a strong homophily/heterophily pattern. Finally, SHAPEGGEN can generate protected features to enable the study of fairness¹. By controlling the value assignment for a selected discrete node feature, SHAPEGGEN introduces bias between the protected feature and node labels. The biased feature is a proxy for a protected feature, such as gender or race. The procedure for generating node features is outlined in **NODEFEATUREVECTORS** within Algorithm 1.

Generating ground-truth explanations. In addition to generating ground-truth labels, SHAPEGGEN has a unique capability to generate unique ground-truth explanations. To accommodate diverse types of GNN explainers, every ground-truth explanation in SHAPEGGEN contains information on a) the importance of nodes, b) the importance of node features, and c) the importance of edges. This information is represented by three distinct masks defined over enclosing subgraphs of nodes $v \in \mathcal{V}_S$, *i.e.*, the L -hop neighborhood around the node. We denote the enclosing subgraph of node $v \in \mathcal{V}_S$ for a given GNN predictor with L layers as: $\text{Sub}(v; L) = (\mathcal{V}_{\text{Sub}(v)}, \mathcal{E}_{\text{Sub}(v)}, \mathbf{X}_{\text{Sub}(v)}) \subseteq \mathcal{G}$. Let motifs within this enclosing subgraph be: $\mathbf{S}_v = (\mathcal{V}_{\mathbf{S}_v}, \mathcal{E}_{\mathbf{S}_v}, \mathbf{X}_{\mathbf{S}_v}) = \mathbf{S} \cap \text{Sub}(v)$. Using this notation, we define ground-truth explanation masks:

- Node explanation mask.** Nodes in $\text{Sub}(v)$ are assigned a value of 0 or 1 based on whether they are located within a motif or not. For any node $v_i \in \mathcal{V}_{\text{Sub}(v)}$, we set $\mathbb{1}_{\mathcal{V}_v}(v_i) = 1$ if $v_i \in \mathcal{V}_{\mathbf{S}_v}$ and 0 otherwise. This function $\mathbb{1}_{\mathcal{V}_v}$ is applied to all nodes in the enclosing subgraph of v to produce an importance score for each node, yielding the final mask as: $\mathbf{M}_n = \{\mathbb{1}_{\mathcal{V}_v}(u) | u \in \mathcal{V}_{\text{Sub}(v)}\}$.
- Node feature explanation mask.** Each feature in v 's feature vector is labeled as 1 if it represents an informative feature and 0 if it is a random feature. This procedure produces a binary node feature mask for node v as: $\mathbf{M}_f \in \{0, 1\}^d$.
- Edge explanation mask.** To each $e = (v_i, v_j) \in \mathcal{E}_{\text{Sub}(v)}$ we assign a value of either 0 or 1 based on whether e connects any two nodes in $\text{Sub}(v)$. The masking function is defined as follows:

$$\mathbb{1}_{\mathcal{E}_{\mathbf{S}_v}}(e) = \begin{cases} 0 & \text{if } v_i \notin \mathcal{V}_{\mathbf{S}_v} \cup \{v\} \vee v_j \notin \mathcal{V}_{\mathbf{S}_v} \cup \{v\} \\ 1 & \text{if } v_i \in \mathcal{V}_{\mathbf{S}_v} \cup \{v\} \wedge v_j \in \mathcal{V}_{\mathbf{S}_v} \cup \{v\} \end{cases} \quad (8)$$

This function $\mathbb{1}_{\mathcal{E}_{\mathbf{S}_v}}$ is applied to all edges $e \in \mathcal{E}_{\text{Sub}(v)}$ to produce ground-truth edge explanation as: $\mathbf{M}_e = \{\mathbb{1}_{\mathcal{E}_{\mathbf{S}_v}}(e) | e \in \mathcal{E}_{\text{Sub}(v)}\}$. The procedure to generate these ground-truth explanations is thoroughly described in Algorithm 1.

Datasets in GRAPHXAI. We proceed with a detailed description of synthetic and real-world graph data resources included in GRAPHXAI.

Synthetic graphs. The SHAPEGGEN generator outlined in the Methods section is a dataset generator that can be used to generate any number of user-specified graphs. In GRAPHXAI, we provide a collection of pre-generated XAI-ready graphs with diverse properties that are readily available for analysis and experimentation.

Base SHAPEGGEN graphs (SG-Base). We provide a base version of SHAPEGGEN graphs. This instance of SHAPEGGEN is homophilic, large, and contains house-shaped motifs for ground-truth explanations, formally described as $\mathcal{G} = \text{SHAPEGGEN}$ ($S = \text{'house'}$, $N_s = 1200$, $p = 0.006$, $n_s = 11$, $K = 2$, $n_f = 11$, $n_i = 4$, $s_f = 0.6$, $c_f = 2$, $\phi = 0.5$, $\eta = 1$, $L = 3$). The node features in this graph exhibit homophily, a property commonly found in social networks. With over 10,000 nodes, this graph also provides enough examples of ground-truth explanations for rigorous statistical evaluation of explainer performance. The house-shaped motifs follow one of the earliest synthetic graphs used to evaluate GNN explainers¹⁴.

Homophilic and heterophilic SHAPEGEN graphs. GNN explainers are evaluated on homophilic graphs^{1,43–45} as low homophily levels in graphs can degrade the performance of GNN predictors^{46,47}. To this end, there are no heterophilic graphs with ground-truth explanations in current GNN XAI literature despite such graphs being abundant in real-world applications⁴⁶. To demonstrate the flexibility of the SHAPEGEN data generator, we use it to generate graphs with: i) homophilic ground-truth explanations (SG-BASE) and ii) heterophilic ground-truth explanations (SG-HETEROPHILIC), *i.e.*, $\mathcal{G} = \text{SHAPEGEN}$ ($\mathcal{S} = \text{'house'}$, $N_s = 1200$, $p = 0.006$, $n_s = 11$, $K = 2$, $n_f = 11$, $n_i = 4$, $s_f = 0.6$, $c_f = 2$, $\phi = 0.5$, $\eta = -1$, $L = 3$).

Weakly and strongly unfair SHAPEGEN graphs. We utilize the SHAPEGEN data generator to generate graphs with controllable fairness properties, *i.e.*, leverage SHAPEGEN to generate synthetic graphs with real-world fairness properties where we can enforce unfairness w.r.t. a given protected attribute. We use the SHAPEGEN to generate graphs with: i) weakly-unfair ground-truth explanations (SG-BASE) and ii) strongly-unfair ground-truth explanations (SG-UNFAIR) $\mathcal{G} = \text{SHAPEGEN}$ ($\mathcal{S} = \text{'house'}$, $N_s = 1200$, $p = 0.006$, $n_s = 11$, $K = 2$, $n_f = 11$, $n_i = 4$, $s_f = 0.6$, $c_f = 2$, $\phi = 0.75$, $\eta = 1$, $L = 3$). Here, for the first time, we generated unfair synthetic graphs that can serve as pseudo-ground-truth for quantifying whether current GNN explainers preserve counterfactual fairness.

Small and large SHAPEGEN explanations. We explore the faithfulness of explanations w.r.t. different ground-truth explanation sizes. This is important because a reliable explanation should identify important features independent of the explanation size. However, current data resources only provide graphs with smaller-size ground-truth explanations. Here, we use the SHAPEGEN data generator to generate graphs having (i) smaller ground-truth explanations size (SG-SMALLEX) $\mathcal{G} = \text{SHAPEGEN}$ ($\mathcal{S} = \text{'triangle'}$, $N_s = 1200$, $p = 0.006$, $n_s = 12$, $K = 2$, $n_f = 11$, $n_i = 4$, $s_f = 0.5$, $c_f = 2$, $\phi = 0.5$, $\eta = 1$, $L = 3$) and (ii) larger ground-truth explanations size using house motifs.

Low and high proportions of salient features. We examine the faithfulness of node feature masks produced by explainers under different levels of sparsity for class-associated signal in the node features. The feature generation procedure in SHAPEGEN specifies node feature parameters n_i , the number of informative features that are generated to correlate with node labels, and n_f number of total node features. The remaining $n_f - n_i$ features are redundant features that are randomly distributed and have no correlation to the node label. Using an equivalent graph topology to SG-Base, we change the relative proportion of node features which are attributed to the label by adjusting n_i and n_f . We create SG-MOREINFORM, a SHAPEGEN instance with a high proportion of ground-truth features to total features (8:11). Likewise, we create SG-LESSINFORM, a SHAPEGEN instance with a low proportion of ground-truth features to total features (4:21). This proportion in SG-Base falls in the middle of SG-MOREINFORM and SG-LESSINFORM instances with a proportion of ground-truth to total features of 4:11. Formally, we define SG-MoreInform as $\mathcal{G} = \text{SHAPEGEN}$ ($\mathcal{S} = \text{'house'}$, $N_s = 1200$, $p = 0.006$, $n_s = 11$, $K = 2$, $n_f = 11$, $n_i = 8$, $s_f = 0.6$, $c_f = 2$, $\phi = 0.5$, $\eta = 1$, $L = 3$) and SG-LessInform as $\mathcal{G} = \text{SHAPEGEN}$ ($\mathcal{S} = \text{'house'}$, $N_s = 1200$, $p = 0.006$, $n_s = 11$, $K = 2$, $n_f = 21$, $n_i = 4$, $s_f = 0.6$, $c_f = 2$, $\phi = 0.5$, $\eta = 1$, $L = 3$).

BA-Shapes. In addition to SHAPEGEN, we provide a version of BA-Shapes¹⁴, a synthetic graph data generator for node classification tasks. We start with a base Barabasi-Albert (BA)⁴⁸ graph using N nodes and a set of five-node “house”-structured motifs K randomly attached to nodes of the base graph. The final graph is perturbed by adding random edges. The nodes in the output graph are categorized into two classes corresponding to whether the node is in a house (1) or not in a house (0).

Real-world graphs. We describe the real-world graph datasets with and without ground-truth explanations provided in GRAPHXAI. To this end, we provide data resources from crime forecasting, financial lending, and molecular chemistry and biology^{1,2,35}. We consider these datasets as they are used to train GNNs for generating predictions in high-stakes downstream applications. In particular, we include chemical and biological datasets because they are used to identify whether an input graph (*i.e.*, a molecular graph) contains a specific pattern (*i.e.*, a chemical group with a specific property in the molecule). Knowledge about such patterns, which determine molecular properties, represents ground-truth explanations². We provide a statistical description of real-world graphs in Tables 6–8. Below, we discuss the details of each of the real-world datasets that we employ:

MUTAG. The MUTAG³⁵ dataset contains 1,768 graph molecules labeled into two different classes according to their mutagenic properties, *i.e.*, effect on the Gram-negative bacterium *S. Typhimurium*. Kazius *et al.*³⁵ identifies several toxicophores - motifs in the molecular graph - that correlate with mutagenicity. The dataset is trimmed from its original 4,337 graphs to 1,768, based on those whose labels directly correspond to the presence or absence of our chosen toxicophores: NH_2 , NO_2 , aliphatic halide, nitroso, and azo-type (terminology, as referred to in Kazius *et al.*³⁵).

Alkane carbonyl. The Alkane Carbonyl² dataset contains 1,125 molecular graphs labeled into two classes where a positive sample indicates a molecule that contains an unbranched alkane and a carbonyl ($\text{C}=\text{O}$) functional group. The ground-truth explanations include any combinations of alkane and carbonyl functional groups within a given molecule.

Dataset	MUTAG	Benzene	Fluoride-Carbonyl	Alkane-Carbonyl
Graphs	1,768	12,000	8,671	4,326
Average nodes	29.15 ± 0.35	20.58 ± 0.04	21.36 ± 0.04	21.13 ± 0.05
Average edges	60.83 ± 0.75	43.65 ± 0.08	45.37 ± 0.09	44.95 ± 0.12
Node features	14	14	14	14
GT Explanation	NH ₂ , NO ₂ chemical group	Benzene Ring	F ⁻ and C=O chemical group	Alkane and C=O chemical group

Table 7. Statistics of real-world graph classification datasets in GRAPHXAI with ground-truth (GT) explanations.

Dataset	German credit graph	Recidivism graph	Credit defaulter graph
Nodes	1,000	18,876	30,000
Edges	22,242	321,308	1,436,858
Node features	27	18	13
Average node degree	44.48 ± 26.51	34.04 ± 46.65	95.79 ± 85.88

Table 8. Statistics of real-world node classification datasets in GRAPHXAI without ground-truth (GT) explanations.

Benzene. The Benzene² dataset contains 12,000 molecular graphs extracted from the ZINC15⁴⁹ database and labeled into two classes where the task is to identify whether a given molecule has a benzene ring or not. Naturally, the ground truth explanations are the nodes (atoms) comprising the benzene rings, and in the case of multiple benzenes, each of these benzene rings forms an explanation.

Fluoride carbonyl. The Fluoride Carbonyl² dataset contains 8,671 molecular graphs labeled into two classes where a positive sample indicates a molecule that contains a fluoride (F⁻) and a carbonyl (C=O) functional group. The ground-truth explanations consist of combinations of fluoride atoms and carbonyl functional groups within a given molecule.

German credit. The German Credit¹ graph dataset contains 1,000 nodes representing clients in a German bank connected based on the similarity of their credit accounts. The dataset includes demographic and financial features like gender, residence, age, marital status, loan amount, credit history, and loan duration. The goal is to associate clients with credit risk.

Recidivism. The Recidivism¹ dataset includes samples of bail outcomes collected from multiple state courts in the USA between 1990–2009. It contains past criminal records, demographic attributes, and other demographic details of 18,876 defendants (nodes) who got released on bail at the U.S. state courts. Defendants are connected based on the similarity of past criminal records and demographics, and the goal is to classify defendants into bail vs. no bail.

Credit defaulter. The Credit defaulter¹ graph has 30,000 nodes representing individuals that we connected based on the similarity of their spending and payment patterns. The dataset contains applicant features like education, credit history, age, and features derived from their spending and payment patterns. The task is to predict whether an applicant will default on an upcoming credit card payment.

Algorithm 1: Overview of SHAPEGGEN Algorithm

Input: Shape $\mathcal{S} = (\mathcal{V}_{\mathcal{S}}, \mathcal{E}_{\mathcal{S}})$; number of subgraphs N_s ; probability of connection p ; subgraph size n_s ; number of classes K ; number of features n_f ; number of informative features n_i ; class separation factor s_f ; number of clusters per class c_f ; protected feature noise factor ϕ ; homophily coefficient η ; model layers L
 $\mathcal{G} \leftarrow \text{SHAPEGGENSTRUCTURE}(\mathcal{S}, N_s, p, n_s, K)$ (See Algorithm 2)
 $\mathbf{Y} \leftarrow$ Labels for each node $v \in \mathcal{G}$ by Equation (7)
 $\mathbf{X}, \mathbf{M}_F \leftarrow \text{NODEFEATUREVECTORS}(\mathcal{V}_{\mathcal{G}}, \mathbf{Y}, K, n_f, n_i, s_f, c_f, \phi, \eta)$
 $\mathbf{M}_N, \mathbf{M}_E \leftarrow \text{CREATEGROUNDTRUTHEXPLANATIONS}(\mathcal{G}, L)$

NODEFEATUREVECTORS

Input: Nodes \mathcal{V} ; labels of all nodes \mathbf{Y} ; number of classes/labels K ; number of features n_f ; number of informative features n_i ; separation factor s_f ; number of clusters per class c_f ; protected feature noise factor ϕ ; homophily coefficient η
Output: Node features \mathbf{X} , Ground-truth feature explanation \mathbf{M}_F
Partition \mathcal{V} into K sets, separated by labels, $\{\mathcal{V}_1, \dots, \mathcal{V}_K\}$
Generate hypercube \mathbf{H} in n_i dimensions scaled by size s_f
Sample $K \times c_f$ vertices $\{h_1, \dots, h_K\}$ from \mathbf{H} such that $h_i \in \mathbb{R}^{n_i}$

```

for  $i \leftarrow 1$  to  $K$  do
  Assign features  $x \sim \mathcal{N}(h_i, 1)$  to each node  $v \in \mathcal{V}_i$ 
  Noise vectors  $r \sim \mathcal{N}(0, 1)$ ,  $r \in \mathbb{R}^{n_i - m_i}$ 
  Set protected feature  $q$  equal to label  $i$ 
  For each node, flip  $q$  with probability of  $\phi$  to another random label
   $\mathbf{X}[v_j] \leftarrow x[v_j] \oplus r[v_j] \oplus q[v_j] \forall v_j \in \mathcal{V}_i$ 
end
 $\mathbf{M}_F \leftarrow$  Mask over all informative dimensions of features
 $\mathbf{X} \leftarrow$  OPTIMIZEHOMOPHILYHETEROPHILY( $\mathcal{G}, \mathbf{X}, \mathbf{Y}, \eta$ ) in Algorithm 3
Return  $\mathbf{X}, \mathbf{M}_F$ 

```

CREATEGROUNDTRUTHEXPLANATIONS

```

Input: SHAPEGGEN  $\mathcal{G} = (\mathcal{V}_{\mathcal{G}}, \mathcal{E}_{\mathcal{G}})$ ; number of layers  $L$ 
Output: Ground-truth node explanation  $\mathbf{M}_N$  and edge explanation  $\mathbf{M}_E$ 
Extract all planted motifs in  $\mathcal{G}$ ,  $\mathbf{S} = \{\mathcal{S}_1, \dots, \mathcal{S}_N\}$ 
Define SUB :  $\mathcal{V}_{\mathcal{G}} \mapsto (\mathcal{V}_{\text{SUB}}, \mathcal{E}_{\text{SUB}})$  as  $L$ -hop subgraph around  $v_i \in \mathcal{V}_{\mathcal{G}}$ 
for  $i \leftarrow 1$  to  $|\mathcal{V}_{\mathcal{G}}|$  do
   $\mathcal{V}_{\text{SUB}} \leftarrow$  Nodes in SUB( $v_i$ );  $\mathcal{E}_{\text{SUB}} \leftarrow$  Edges in SUB( $v_i$ )
   $\mathbf{S}_{\text{SUB}(v_i)} \leftarrow$  SUB( $v_i$ )  $\cap$  ( $\bigcup_{j=1}^{|\mathbf{S}|} \mathcal{S}_j$ )
  for  $j \leftarrow 1$  to  $|\text{SUB}(v_i)|$  do
     $\mathbf{M}_N^{v_i}[j] \leftarrow 1$  if  $v_j \in \mathbf{S}_{\text{SUB}(v_i)}$ , else 0
  end
   $\mathbf{M}_E^{v_i} \leftarrow 1$  for any edge  $(v_a, v_b)$  such that  $\mathbf{M}_N^{v_i}[a] = \mathbf{M}_N^{v_i}[b] = 1$ , else 0
end
Return  $\mathbf{M}_N, \mathbf{M}_E$ 

```

Algorithm 2: SHAPEGGENSTRUCTURE

```

Input: Shape  $\mathcal{S} = (\mathcal{V}_{\mathcal{S}}, \mathcal{E}_{\mathcal{S}})$ ; number of subgraphs  $N_s$ ; probability of connection  $p$ ; subgraph size  $n_s$ ;
number of classes  $K$ 
Output: SHAPEGGEN  $\mathcal{G}$ 
 $\mathbf{S} \leftarrow \{\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_{N_s}\}$ 
for  $i \leftarrow 1$  to  $N_s$  do
   $\mathcal{V}_{\mathcal{S}_i} \leftarrow$  nodes in  $\mathcal{S}_i$ 
   $\Delta s_i \leftarrow n_s - |\mathcal{V}_{\mathcal{S}_i}|$ 
   $n_i \sim$  Poisson( $\lambda = \Delta s_i$ )
  for  $j \leftarrow 1$  to  $n_i$  do
     $g_i \propto \frac{\text{deg}(v)}{\sum_{v_j \in \mathcal{V}_{\mathcal{S}_i}} \text{deg}(v_j)} \forall v \in \mathcal{V}_{\mathcal{S}_i}$ 
     $v^* \leftarrow$  Node sampled with  $g_i$  from  $\mathcal{V}_{\mathcal{S}_i}$ 
    Introduce new node  $v'$  to subgraph  $\mathcal{S}_i$  with edge  $(v', v^*)$ 
  end
end
 $\mathcal{G} \leftarrow \bigcup_{i=1}^{N_s} \mathcal{S}_i$ 
for  $i \leftarrow 1$  to  $N_s$  do
  for  $j \leftarrow i$  to  $N_s$  do
    if  $p$  then
      continue
    else
       $\mathcal{E}_{ij} \leftarrow$  All possible edges that could connect  $\mathcal{S}_i$  and  $\mathcal{S}_j$ 
      while  $\mathcal{E}_{ij} \neq \emptyset$  do
         $g_{ij} \propto \frac{\text{deg}(v_1) + \text{deg}(v_2)}{\sum_{(v_k, v_m) \in \mathcal{E}_{ij}} \text{deg}(v_k) + \text{deg}(v_m)} \forall (v_1, v_2) \in \mathcal{E}_{ij}$ 
         $e^* \leftarrow$  Edge sampled with  $g_{ij}$  from  $\mathcal{E}_{ij}$ 
        x Remove  $e^*$  from  $\mathcal{E}_{ij}$ 
        Add edge  $e^*$  to  $\mathcal{G}$ 
        if  $\exists v \in \mathcal{V}_{\mathcal{G}}$  such that  $f(v) > K - 1$  (Eqn. 7) then
          Remove  $e^*$  from  $\mathcal{G}$ 
        else
          Break
        end
      end
    end
  end
end
 $\mathcal{G}^* \leftarrow$  Largest connected component of  $\mathcal{G}$ 
Return  $\mathcal{G}^*$ 

```

Algorithm 3: OPTIMIZEHOMOPHILYHETEROPHILY

Input: Graph \mathcal{G} , node features \mathbf{X} , node labels \mathbf{Y} , homophily coefficient η
Output: Optimized node features \mathbf{X}_η
 $\mathbf{C} \leftarrow$ Sample connected set of nodes with same labels in \mathbf{Y}
 $\mathbf{D}_{\text{dis}} \leftarrow$ Sampled set of *disconnected* pairs of nodes with different labels
 $\mathbf{D}_{\text{con}} \leftarrow$ Sampled set of *connected* pairs of nodes with different labels
 $L_{\mathbf{C}}(\mathbf{X}) = -\eta \frac{1}{|\mathbf{C}|} \sum_{(v_i, v_j) \in \mathbf{C}} D(\mathbf{X}[v_i], \mathbf{X}[v_j])$
 $L_{\text{dis}}(\mathbf{X}) = \eta \frac{1}{|\mathbf{D}_{\text{dis}}|} \sum_{(v_i, v_j) \in \mathbf{D}_{\text{dis}}} D(\mathbf{X}[v_i], \mathbf{X}[v_j])$
 $L_{\text{con}}(\mathbf{X}) = \eta \frac{1}{|\mathbf{D}_{\text{con}}|} \sum_{(v_i, v_j) \in \mathbf{D}_{\text{con}}} D(\mathbf{X}[v_i], \mathbf{X}[v_j])$
 $L(\mathbf{X}) = L_{\mathbf{C}}(\mathbf{X}) + L_{\text{dis}}(\mathbf{X}) + L_{\text{con}}(\mathbf{X})$
Return $\arg \min_{\mathbf{X}} L(\mathbf{X})$

Data availability

GRAPHXAI data resource²⁸ is hosted on Harvard Dataverse under a persistent identifier <https://doi.org/10.7910/DVN/KULOS8>. We have deposited different a number of SHAPEGEN -generated datasets and real-world graphs at this repository.

Code availability

Project website for GRAPHXAI is at <https://zitniklab.hms.harvard.edu/projects/GraphXAI>. The code to reproduce results, documentation, and tutorials are available in GRAPHXAI's Github repository at <https://github.com/mims-harvard/GraphXAI>. The repository contains Python scripts to generate and evaluate explanations using performance metrics and also visualize explanationa. In addition, the repository contains information and Python scripts to build new versions of GRAPHXAI as the underlying primary resources get updated and new data become available.

Received: 19 August 2022; Accepted: 17 January 2023;

Published online: 18 March 2023

References

1. Agarwal, C., Lakkaraju, H. & Zitnik, M. Towards a unified framework for fair and stable graph representation learning. In *UAI* (2021).
2. Sanchez-Lengeling *et al.* Evaluating attribution for graph neural networks. *NeurIPS* (2020).
3. Giunchiglia, V., Shukla, C. V., Gonzalez, G. & Agarwal, C. Towards training GNNs using explanation directed message passing. In *The First Learning on Graphs Conference* (2022).
4. Morselli Gysi, D. *et al.* Network medicine framework for identifying drug-repurposing opportunities for covid-19. *Proceedings of the National Academy of Sciences* (2021).
5. Zitnik, M., Agrawal, M. & Leskovec, J. Modeling polypharmacy side effects with graph convolutional networks. In *Bioinformatics* (2018).
6. Baldassarre, F. & Azizpour, H. Explainability techniques for graph convolutional networks. In *ICML Workshop on LRGR* (2019).
7. Faber, L. *et al.* Contrastive graph neural network explanation. In *ICML Workshop on Graph Representation Learning and Beyond* (2020).
8. Huang, Q., Yamada, M., Tian, Y., Singh, D. & Chang, Y. Graphlime: Local interpretable model explanations for graph neural networks. *IEEE Transactions on Knowledge and Data Engineering* (2022).
9. Lucic, A., Ter Hoeve, M. A., Tolomei, G., De Rijke, M. & Silvestri, F. Cf-gnnexplainer: Counterfactual explanations for graph neural networks. In *AISTATS* (PMLR, 2022).
10. Luo, D. *et al.* Parameterized explainer for graph neural network. In *NeurIPS* (2020).
11. Pope, P. E., Kolouri, S., Rostami, M., Martin, C. E. & Hoffmann, H. Explainability methods for graph convolutional neural networks. In *CVPR* (2019).
12. Schlichtkrull, M. S., De Cao, N. & Titov, I. Interpreting graph neural networks for nlp with differentiable edge masking. In *ICLR* (2021).
13. Vu, M. N. & Thai, M. T. PGM-Explainer: probabilistic graphical model explanations for graph neural networks. In *NeurIPS* (2020).
14. Ying, R., Bourgeois, D., You, J., Zitnik, M. & Leskovec, J. GNNExplainer: generating explanations for graph neural networks. In *NeurIPS* (2019).
15. Agarwal, C. *et al.* Probing GNN explainers: A rigorous theoretical and empirical analysis of GNN explanation methods. In *AISTATS* (2022).
16. Faber, L., K. Moghaddam, A. & Wattenhofer, R. When comparing to ground truth is wrong: On evaluating GNN explanation methods. In *KDD* (2021).
17. Hu, W. *et al.* Open Graph Benchmark: datasets for machine learning on graphs. In *NeurIPS* (2020).
18. Baruah, T. *et al.* GNNMark: a benchmark suite to characterize graph neural network training on gpus. In *ISPASS* (2021).
19. Du, Y. *et al.* GraphGT: machine learning datasets for graph generation and transformation. In *NeurIPS Datasets and Benchmarks* (2021).
20. Freitas, S. *et al.* A large-scale database for graph representation learning. In *NeurIPS Datasets and Benchmarks* (2021).
21. Zheng, Q. *et al.* Graph robustness benchmark: Benchmarking the adversarial robustness of graph machine learning. In *NeurIPS Datasets and Benchmarks* (2021).
22. Huang, K. *et al.* Therapeutics Data Commons: Machine learning datasets and tasks for drug discovery and development. In *NeurIPS Datasets and Benchmarks* (2021).
23. Huang, K. *et al.* Artificial intelligence foundation for therapeutic science. *Nature Chemical Biology* **18**, 1033–1036 (2022).
24. Wang, Z., Yin, H. & Song, Y. Benchmarking the combinatorial generalizability of complex query answering on knowledge graphs. In *NeurIPS Datasets and Benchmarks* (2021).
25. Liu, M. *et al.* DIG: A turnkey library for diving into graph deep learning research. *JMLR* (2021).
26. Fey, M. & Lenssen, J. E. Fast graph representation learning with PyTorch Geometric. *ICLR 2019 (RLGM Workshop)* (2019).
27. Wang, M. *et al.* Deep Graph Library: Towards efficient and scalable deep learning on graphs. In *ICLR workshop on representation learning on graphs and manifolds* (2019).

28. Agarwal, C., Queen, O., Lakkaraju, H. & Zitnik, M. Evaluating explainability for graph neural networks. *Harvard Dataverse* <https://doi.org/10.7910/DVN/KULOS8> (2022).
29. Simonyan, K. *et al.* Deep inside convolutional networks: Visualising image classification models and saliency maps. In *ICLR* (2014).
30. Sundararajan, M., Taly, A. & Yan, Q. Axiomatic attribution for deep networks. In *ICML* (2017).
31. Yuan, H., Yu, H., Wang, J., Li, K. & Ji, S. On explainability of graph neural networks via subgraph explorations. In *ICML* (2021).
32. Xu, K., Hu, W., Leskovec, J. & Jegelka, S. How powerful are graph neural networks? In *ICLR* (2019).
33. Kipf, T. N. & Welling, M. Semi-supervised classification with graph convolutional networks. In *ICLR* (2017).
34. Taha, A. A. & Hanbury, A. Metrics for evaluating 3d medical image segmentation: analysis, selection, and tool. In *BMC Medical Imaging* (2015).
35. Kazius, J. *et al.* Derivation and validation of toxicophores for mutagenicity prediction. In *Journal of Medicinal Chemistry* (2005).
36. Yuan, H., Yu, H., Gui, S. & Ji, S. Explainability in graph neural networks: A taxonomic survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2022).
37. Hunter, J. D. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering* (2007).
38. Hagberg, A. A., Schult, D. A. & Swart, P. J. Exploring network structure, dynamics, and function using networkx. In *Proceedings of the 7th Python in Science Conference* (2008).
39. Wu, Z. *et al.* A comprehensive survey on graph neural networks. In *IEEE Transactions on Neural Networks and Learning Systems* (2020).
40. Barabási, A.-L. & Albert, R. Emergence of scaling in random networks. In *Science* (1999).
41. Guyon, I. Design of experiments of the nips 2003 variable selection benchmark. In *NIPS 2003 workshop on feature extraction and feature selection* (2003).
42. Pedregosa, F. *et al.* Scikit-learn: Machine learning in Python. In *JMLR* (2011).
43. McCallum, A. K. *et al.* Automating the construction of internet portals with machine learning. In *Information Retrieval* (2000).
44. Sen, P. *et al.* Collective classification in network data. In *AI magazine* (2008).
45. Wang, K. *et al.* Microsoft academic graph: When experts are not enough. In *Quantitative Science Studies* (2020).
46. Zhu, J. *et al.* Beyond homophily in graph neural networks: Current limitations and effective designs. In *NeurIPS* (2020).
47. Jin, W. *et al.* Node similarity preserving graph convolutional networks. In *WSDM* (2021).
48. Albert, R. & Barabási, A.-L. Statistical mechanics of complex networks. *Reviews of modern physics* (2002).
49. Sterling, T. & Irwin, J. J. Zinc 15–ligand discovery for everyone. *Journal of Chemical Information and Modeling* (2015).

Acknowledgements

C.A., O.Q., and M.Z. gratefully acknowledge the support by NSF under Nos. IIS-2030459 and IIS-2033384, US Air Force Contract No. FA8702-15-D-0001, and awards from Harvard Data Science Initiative, Amazon Research, Bayer Early Excellence in Science, AstraZeneca Research, and Roche Alliance with Distinguished Scientists. H.L. was supported in part by NSF under Nos IIS-2008461 and IIS-2040989, and research awards from Google, JP Morgan, Amazon, Bayer, Harvard Data Science Initiative, and D3 Institute at Harvard. O.Q. was supported, in part, by Harvard Summer Institute in Biomedical Informatics. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funders.

Author contributions

C.A., O.Q., H.L. and M.Z. contributed new analytic tools and wrote the manuscript. C.A. and O.Q. retrieved, processed, and harmonized datasets. C.A. and O.Q. implemented the synthetic dataset generator and performed the analyses for technical validation of the new resource. M.Z. conceived the study.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to M.Z.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023