

Augment to Interpret: Unsupervised and Inherently Interpretable Graph Embeddings

Gregory Scafarto

Madalina Ciortan

Simon Tihon

Quentin Ferré

Euranova, Brussels, Belgium

GREGORY.SCAFARTO@EURANOVA.EU

MADALINA.CIORTAN@EURANOVA.EU

SIMON.TIHON@EURANOVA.EU

QUENTIN.FERRE@EURANOVA.EU

Editors: Berrin Yanıkoğlu and Wray Buntine

Abstract

Unsupervised learning allows us to leverage unlabelled data, which has become abundantly available, and to create embeddings that are usable on a variety of downstream tasks. However, the typical lack of interpretability of unsupervised representation learning has become a limiting factor with regard to recent transparent-AI regulations. In this paper, we study graph representation learning and we show that data augmentation that preserves semantics can be learned and used to produce interpretations. Our framework, which we named INGENIOUS, creates inherently interpretable embeddings and eliminates the need for costly additional post-hoc analysis. We also introduce additional metrics addressing the lack of formalism and metrics in the understudied area of unsupervised-representation-learning interpretability. Our results are supported by an experimental study applied to both graph-level and node-level tasks and show that interpretable embeddings provide state-of-the-art performance on subsequent downstream tasks.¹

Keywords: interpretability; unsupervised learning; graph embedding; contrastive loss

1. Introduction

Graphs are powerful data representations that model objects as well as the relationships between them. Much like images, they are often complex to process. A popular solution is to project such data into a latent space, thus creating vector embeddings, which can later be used by any classical machine learning training and inference pipeline. These embeddings could be learned with supervision, but labelled data is often difficult to obtain. This has resulted in an increased interest in the learning of semantic-preserving embeddings, which can later be used for a variety of downstream tasks. There are numerous ways to learn graph representations (Hamilton et al., 2017) without label supervision, but using graph neural networks (GNN) has become the go-to approach. In particular, graph contrastive learning (GCL) aims at creating graph embeddings without supervision by leveraging augmented views of available samples to extract meaningful information. This approach provided state-of-the-art results when applied to diverse domains ranging from chemistry (Duvenaud et al., 2015) to social sciences (Fan et al., 2019).

1. Our code is available at https://github.com/euranova/Augment_to_Interpret.

Chen et al. (2020) demonstrated that the use of data augmentation in self-supervised contrastive learning can lead to the creation of semantic-preserving embeddings. Similar results have been achieved on graph data by leveraging carefully crafted augmentation (You et al., 2020). Despite promising results, the underlying GNNs act as black-boxes. They are difficult to interpret, debug and trust, which raises the question: *What input features do GNNs focus on when generating representations?* Using edge and node dropping as graph augmentation can lead to the creation of sub-graphs that retain discriminative semantic information while reaching better results (You et al., 2020) than other augmentation schemes. Interestingly, this objective correlates with the aim of GNN interpretability, which is to identify highly influential sparse subsets of nodes and edges with the largest impact on the model’s behaviour (Luo et al., 2020). This observation leads us to the hypothesis that a well-learned augmentation can serve to produce interpretations of the information embedded in GNN representations. Our main contributions follow.

- We propose INGENIOUS (INherently INterpretable Graph and Node Unsupervised embeddings), a framework that generalises over existing approaches based on learned augmentation (Suresh et al., 2021; Gao et al., 2022; Miao et al., 2022) and we introduce new losses and a module to produce useful and interpretable embeddings.
- We show that embeddings produced by INGENIOUS are relevant to a variety of downstream tasks, matching or exceeding results obtained by state-of-the-art approaches.
- We show that carefully designed learned graph augmentation can be used to produce interpretations. We assess their quality along different axes introduced by Nauta et al. (2023): correctness, completeness, continuity, compactness and coherence. To this end, we introduce new metrics.
- We conduct a hyperparameter study, highlighting the importance of sparsity when using augmented views as interpretations.

To our knowledge, this is the first attempt at assessing the link between learned augmentation and interpretability. We characterise it on graph-level and node-level tasks. The structure of the paper is as follows: We first position our work in Section 2 and present our framework in detail in Section 3. In Section 4, we assess the utility and interpretability of the framework by introducing necessary metrics and we perform an in-depth study of the properties of the framework. Finally, we discuss limitations and conclude.

2. Related Work

Previous works on other modalities have shown the superiority of augmentation-based contrastive learning (Chen et al., 2020). However, adapting these works so that they work reliably on graph data raised new challenges, such as defining the right augmentation techniques amongst existing augmentation techniques (Ding et al., 2022). Early methods, including GraphCL (You et al., 2020) and InfoCL (Wang et al., 2022b), focus on random perturbations produced by dropping edges or nodes. Other methods like graphMVP (Liu et al., 2022) and MICRO-Graph (Subramonian, 2021) use domain knowledge to perform better than random augmentation. More recent approaches such as AD-GCL (Suresh et al.,

2021) and MEGA (Gao et al., 2022) use a multilayer perceptron (MLP) as an augmentation learner. They create augmented views by learning to drop low-relevance edges. Similarly, RGCL (Li et al., 2022) adds a repulsive dynamic to the augmented views.

In parallel, considerable work has been done to interpret graph neural networks post training. A backpropagation of gradient-like signals (Simonyan et al., 2014), perturbations of the input graph (Ying et al., 2019; Luo et al., 2020), and the training of a surrogate model (Vu and Thai, 2020) are examples of such techniques. These post-hoc techniques have been shown to be biased and unrepresentative of the true interpretations (Rudin, 2019), leading to an increased interest in inherently interpretable models (Miao et al., 2022; Feng et al., 2022). Furthermore, the aforementioned methods focus on interpreting predictions of models learned with supervision. Zheng et al. (2022) have proposed USIB, the first post-hoc method to interpret unsupervised graph representations, but USIB faces the same faithfulness limitation as other post-hoc techniques.

Given existing limitations and the power of augmentation-based learning, we investigate how augmentation learning can faithfully and intrinsically provide interpretability.

3. Materials and Methods

In this section, we present INGENIOUS. We begin with a general overview and then provide a description of the augmentation and learning processes.

3.1. Framework

We propose an inherently interpretable framework that produces graph representations in an unsupervised manner by leveraging learned graph augmentation, as shown in Figure 1. Learned edge-dropping augmentation has been used both in contrastive learning by AD-GCL (Suresh et al., 2021) and MEGA (Gao et al., 2022), two recent state-of-the-art approaches, and in interpretability by GSAT (Miao et al., 2022) and PGExplainer (Luo et al., 2020). Inspired by said contrastive approaches, INGENIOUS is trained with a contrastive loss, following a dual-branch approach as recommended by Chen et al. (2020). INGENIOUS is based on two modules: an embedding module that produces node embeddings (ϕ_u) and graph embeddings (ϕ), and an edge-selection module (θ) that produces semantic-preserving sub-graphs by stochastically dropping uninformative edges. In this study, the embedding module is implemented as a graph neural network (GNN) encoder and the edge-selection module as an MLP, as described in Section 4.1.

The embedding module first generates an embedding ϕ_u for each node u of the input graph. Using these node embeddings, the edge-selection module θ produces a keeping probability $p_{uv} \in [0, 1]$ for each edge (u, v) of the graph. Then, the Gumbel-max reparameterisation trick is used to partially remove edges according to their keeping probability p_{uv} in a differentiable manner. Each edge of an augmented view therefore has a weight $w_{uv} \in [0, 1]$.

This edge-selection process is applied twice to obtain two positively augmented views. A negatively augmented view is also produced by giving each edge the flipped weight $1 - w_{uv}$ from the first positively augmented view. The stochastic aspect of this data augmentation increases the robustness of the model, as it shows more diverse data to the model. Ad-

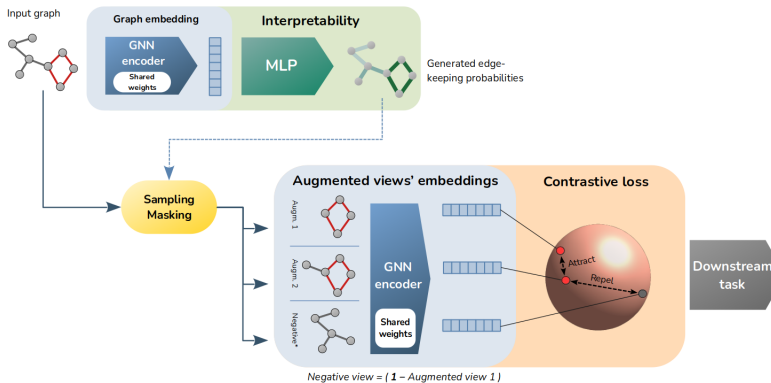


Figure 1: Schema of the INGENIOUS framework, as described in the text.

ditionally, this aspect is needed to apply the dual-branch contrastive loss, since this loss necessitates at least two distinct augmented views for each graph.

Each augmented view of the input graph is then embedded using once again² the embedding module ϕ . Finally, a simclr loss is back-propagated to pull the embeddings of corresponding positively augmented views together while keeping the embeddings of other graphs' positively augmented views distant in the embedding space. The negatively augmented view is also used as a part of the final loss, as described below.

At inference time, the Gumbel-max sampling is no longer used. Keeping probabilities p_{uv} are therefore not used as probabilities anymore but are used directly as weights on edges of the original graph, i.e. $w_{uv} = p_{uv}$, in order for the augmentation to be deterministic. The full-graph embedding is defined as the embedding of the augmented view, to be used in any downstream task. This augmented view serves as the interpretation of the embedding.

3.1.1. EMBEDDING MODULE AND AUGMENTED-GRAPH GENERATION

In this work, we consider attributed graphs $G = (V, E)$ where V is the set of nodes and $E \subset \{(u, v) | u \in V, v \in V\}$ is the set of edges of the graph. By giving an arbitrary order to V , each node can be represented by a single integer $u \in \mathbb{N}$. G can then be described by its node-feature matrix $X \in \mathbb{R}^{|V| \times d}$ and its adjacency matrix $A \in \mathbb{R}^{|V| \times |V|}$, with $|V|$ the total number of nodes in the graph, d the number of node features, X_{ui} the value of the i^{th} feature of the u^{th} node u , and $A_{uv} = 1$ if $(u, v) \in E$ else 0. Node embeddings can be computed by the embedding module as $z_u = \phi_u(A, X)$. Graph embeddings can be computed as $z = \phi(A, X) = Pooling(\{z_u \forall u \in V\})$, with $Pooling$ an aggregation function on sets of vectors, such as the mean or the sum.

Augmented views of a graph G can be produced in two different ways given the embedding z_u of each of its nodes u . The first way, which was described briefly above, is this: The edge-selection module θ produces a keeping probability $p_{uv} = \theta([z_u, z_v])$ for each edge (u, v) , with $[\cdot, \cdot]$ the concatenation operator. Then, edges are sampled following the Gumbel-max

2. When a module includes batch normalisation layers, it is tricky to use this same module for distinct steps in a model, because the input distribution of the batch normalisation layers may be different for each step the module is used for. This aspect was not considered in previous approaches and may have perturbed their results and conclusions. More details are given in supplementary materials.

reparameterization (Gumbel, 1954), giving each edge its weight

$$w_{uv} = \text{Sigmoid}((\text{Logit}(p_{uv}) + \epsilon_{uv})/\tau) \quad (1)$$

with *Logit* the inverse function of the *Sigmoid* function, $\epsilon_{uv} = \log(\alpha_{uv}) - \log(1 - \alpha_{uv})$ a random noise defined using $\alpha_{uv} \sim \text{Uniform}(0, 1)$, and $\tau \in \mathbb{R}^+$ a temperature hyperparameter we set to 1. The higher the value of τ , the more uniform the weight distribution.

The second way is to produce a keeping probability $p_u = \theta(z_u)$ for each node, sample nodes by giving them a weight $w_u = \text{Gumbel}(p_u)$ with *Gumbel* the Gumbel-max trick described above, and then lift the weights from nodes to edges following $w_{uv} = w_u \cdot w_v$.

Although both could be used for both graph and node embeddings, in our approach we use the first way for graph embeddings and the second way for node embeddings, since it works best empirically. These edge weights are then used to weigh messages during the message-passing operation of the GNN encoder. At train time, independent samplings are used to generate two distinct positively augmented views. The negatively augmented view is obtained by flipping the first positively augmented view, that is, each edge of the negatively augmented view has a weight $w_{uv} = 1 - w_{uv}^+$, with w_{uv}^+ the weight of the edge (u, v) in the first positively augmented view. Finally, the two positively augmented views and the negatively augmented view are embedded by the embedding module, which produces embeddings z_1^+ , z_2^+ and z^- . At inference time, the deterministic process described previously is used.

3.1.2. LOSSES

In this section, we define the INGENIOUS loss. It is composed of four losses with diverse objectives.

The first loss is defined as the dual-branch approach of the simclr loss (Chen et al., 2020), which leverages two augmented views rather than an anchor and an augmented view, as proposed by Zhu et al. (2021). This is the main loss of the framework as it structures the embedding space. Specifically, given augmented-view-embedding triplets coming from N distinct graphs and represented as $\{Z_1^+, Z_2^+, Z^-\} \subset \mathbb{R}^{N \times D}$, with D the embedding dimension, we first normalise each embedding so that it has an L2-norm of 1. Then, the loss is defined as

$$\mathcal{L}_{\text{simclr}}(Z_1^+, Z_2^+) = -\frac{1}{2N} \sum_{i=1}^N \left(\log(l_i(Z_1^+, Z_2^+)) + \log(l_i(Z_2^+, Z_1^+)) \right) \quad (2)$$

$$l_i(Z_1^+, Z_2^+) = \frac{\exp(\text{sim}(Z_{1i}^+, Z_{2i}^+)/\beta)}{\sum_{k=1}^N \left(\exp(\text{sim}(Z_{1i}^+, Z_{2k}^+)/\beta) + \mathbb{I}_{[k \neq i]} \exp(\text{sim}(Z_{1i}^+, Z_{1k}^+)/\beta) \right)} \quad (3)$$

with $\text{sim}(\cdot, \cdot)$ a similarity operation such as the dot product, $\mathbb{I}_{[\text{condition}]}$ an indicator function with a value of 1 when the condition is met, else 0, and β a hyperparameter we set to 0.07 as in Chen et al. (2020).

The second loss, which we call the *negative* loss, complements the first loss in structuring the latent space. It leverages the negatively augmented views of the normalised triplets:

$$\mathcal{L}_-(Z_1^+, Z^-) = \sum_{i=1}^N \mathcal{L}_{\text{simclr}}(Z_i^*, Z_i^*) \quad (4)$$

with $Z_i^* = \{Z_{1i}^+, Z_i^-\} \in \mathbb{R}^{2 \times D}$. The motivation behind this idea stems from the observation that the simclr loss has an attracting effect between corresponding augmented-view embeddings and a repelling effect between everything else. If corresponding augmented-view embeddings are always an embedding and itself, there is no attracting effect anymore as something cannot attract itself. Thus, only the repelling effect of the simclr loss has an effect, which is a repelling effect between the positively-augmented-view embedding and the negatively-augmented-view embedding.

The third loss is the information loss from GSAT. It is used as regularisation. This loss minimises the mutual information between the positively augmented view and its original graph anchor following the graph information bottleneck (GIB) principle (Miao et al., 2022). It favours sparser augmented views as it prevents the model from collapsing and selecting all edges as important. The information loss is defined as:

$$\mathcal{L}_{info}(G, P) = \frac{1}{|E|} \sum_{(u,v) \in E} p_{uv} \log \frac{p_{uv}}{r} + (1 - p_{uv}) \log \frac{1 - p_{uv}}{1 - r} \quad (5)$$

with $P = \{p_{uv} \forall (u, v) \in E\}$ the keeping probabilities of all edges, $G = (V, E)$ a batch of graphs as one big graph, $|E|$ the total number of edges and $r \in [0, 1]$ a hyperparameter that represents the prior of the probability to keep any given edge. Similarly to GSAT, the hyperparameter r is initially set to 0.9 and gradually decays to the tuned value of 0.7. As proposed in GSAT, the weights w_{uv} are used in practice instead of the probabilities p_{uv} .

The fourth and last loss is the watchman loss. Wang et al. (2022a) have shown that edge-dropping augmentation can degrade the quality of the latent space. Additionally, we have observed that learned-augmentation-based methods sometimes have unstable training. Preliminary experiments supporting this claim are given in supplementary materials. The watchman loss mitigates these issues by forcing the embeddings to contain graph-level information. In detail, a watchman module ψ uses augmented-view embeddings Z_1^+ to predict $\Lambda_k(G_i) \forall i$, the top k eigenvalues of the Laplacian matrix of each graph G_i composing the batch G . The loss is defined as

$$\mathcal{L}_{Wm}(G, Z_1^+) = \frac{1}{N} \sum_{i=1}^N MSE(\psi(Z_1^+), \Lambda_k(G_i)) \quad (6)$$

with $MSE(\cdot, \cdot)$ the squared euclidean distance between two vectors. This loss is meaningful for graph-embedding tasks only, as node embeddings have no reason to contain information about the full graph.

Finally, the INGENIOUS loss is defined on a batch G of N graphs as:

$$\mathcal{L} = \mathcal{L}_{simclr}(Z_1^+, Z_2^+) + \mathcal{L}_-(Z_1^+, Z^-) + \mathcal{L}_{info}(G, W) + \lambda_{Wm} \cdot \mathcal{L}_{Wm}(G, Z_1^+) \quad (7)$$

with $W = \{w_{uv} \forall (u, v) \in E\}$ the edge weights of the first positively augmented view and λ_{Wm} a hyperparameter we set to 0 for node-embedding tasks and to 0.3 for graph-embedding tasks, based on preliminary empirical results.

4. Experiments

In this section, we evaluate INGENIOUS both in terms of utility and interpretability by comparing its performance to the performance of its unsupervised competitors AD-GCL and

MEGA on graph-classification tasks and we conduct an ablation study on the loss terms. However, these competitors have not been extended to handle node-classification tasks, and to the best of our knowledge, there is no competitor that uses learned augmentation for these tasks. For both graph and node classification tasks, we use GSAT not as a direct competitor, as it is a supervised model, but as an upper bound on the obtainable utility.

4.1. Experimental Setup

Datasets For graph classification, we use two synthetic datasets (*BA2Motifs* (Luo et al., 2020), *SPMotifs.5* (Wu et al., 2022)) and one real-world dataset (*Mutag* (Morris et al., 2020)). For node classification, we use two synthetic datasets (*Tree-grid*, *Tree-cycle* (Ying et al., 2019)) and one real-world dataset (*Cora* (McCallum et al., 2000)). All of them are used with a batch size of 256 except for *Cora* which is used with a batch size of 128. For *Mutag* and *BA2Motifs*, we randomly split the data into train (80%), validation (10%) and test (10%) sets. For the other datasets, we use the same split as in their original papers.

Model For the embedding module ϕ , we use a GIN (Xu et al., 2019) model with 3 layers. For graph-classification tasks, a final pooling layer is added to the model. For the edge-selection module θ , we use a 2-layer MLP. For the watchman module ψ , we use a 3-layer MLP. We use a learning rate of 1e-3 and a dropout ratio of 0.3, and we train the whole on 150 epochs. For all approaches, the metrics defined in Sections 4.2 and 4.3 are measured on the model that obtains the best downstream ACC (see Section 4.2) measured on the validation set, amongst all epochs, as in GSAT. All results are averaged on 3 random seeds.

4.2. Utility

The utility of embeddings learned without supervision is traditionally (Suresh et al., 2021) approximated by the performance of simple models trained on these embeddings and their labels across various downstream tasks. Similarly, we feed the representations to be evaluated to a linear model and we train the latter on the classification task associated with the available labels. The final metric is the accuracy of the model, which we call **downstream ACC**. As visible in Figure 2, the accuracy of INGENIOUS is similar to the accuracy of its competitors (AD-GCL and MEGA) in the case of graph classification. Furthermore, it is close to the supervised upper bound (GSAT) despite being unsupervised.

4.3. Interpretability

The idea of interpreting representations can be defined as finding informative parts of the input that led to its representation. Like most graph-interpretability methods (Yuan et al., 2020), INGENIOUS focuses on topology and produces sub-graphs as interpretations. Examples of such interpretations are visible in Figure 3. Interpretability and how to measure it has been extensively studied in the supervised setting (Nauta et al., 2023), but little work has focused on interpretability for unsupervised representation learning. Monotony and localisation metrics have been introduced by (Wickstrøm et al., 2023; Zheng et al., 2022), but most of the major categories introduced by the taxonomy proposed by Nauta et al. (2023) have not been covered yet. Building on previous work, we introduce five metrics -

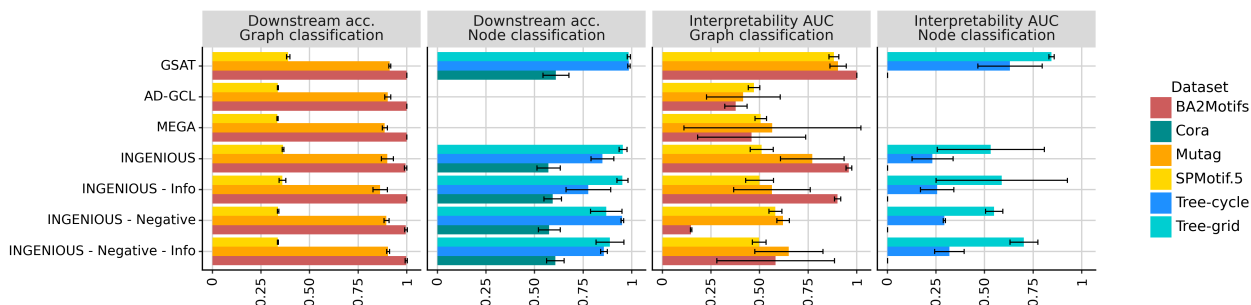


Figure 2: Downstream ACC and interpretability AUC. INGENIOUS outperforms its unsupervised competitors, AD-GCL and MEGA.

Faithfulness, Opposite faithfulness, Wasserstein distance, Bi-modality and Interpretability AUC -, all based on desirable characteristics for interpretations.

Correctness and Completeness Correctness (how much the interpretation truly describes the behaviour of the model) and completeness (how much of the model behaviour is described by the interpretation) (Nauta et al., 2023) are studied by evaluating whether an interpretation aligns with the edges that truly impact the embedding. The canonical approach to evaluating it is the deletion (resp. insertion) experiment introduced by RISE (Petsiuk et al., 2018), which involves removing elements in order (resp. opposite order) of importance and studying the impact on model outputs. Usually, this impact is measured as a difference in predicted scores. However, in the unsupervised setting, outputs are not predicted scores, but embeddings. Therefore, we propose to extend the metric by measuring this impact as the Euclidean distance between the embedding of the original graph and the embeddings of the perturbed graphs.

To evaluate correctness, we use an adapted faithfulness metric. Given a graph, we delete its edges in decreasing order of importance weights given by the edge-selection module. To achieve this, we group edges of similar importance into bins. In our experiments, we use 30 bins. Then, we successively remove bins of edges from the graph in decreasing order of importance, going from the full graph to a graph without any edge. The perturbed

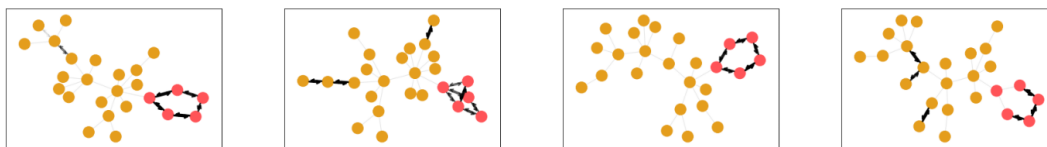


Figure 3: Illustrative examples of interpretations produced by INGENIOUS on four *BA2Motifs* graphs selected randomly. Nodes included in the class-specific motifs are shown in red. Edge importance weights produced by INGENIOUS are represented through edge thickness. Therefore, thick edges between red nodes indicate that INGENIOUS focuses on the informative motifs.

Table 1: Faithfulness and Wasserstein results. A higher faithfulness gap means the interpretations are more faithful, and a higher Wasserstein gap means the embedding space is more continuous in terms of interpretations.

Method	Faithfulness Gap			Wasserstein Gap		
	BA2Motifs	Mutag	SPMotifs.5	BA2Motifs	Mutag	SPMotifs.5
GSAT	.34 ± .00	.21 ± .08	.06 ± .02	.04 ± .00	.02 ± .00	<u>.04 ± .00</u>
AD-GCL	.01 ± .13	.03 ± .09	.08 ± .06	.10 ± .01	.01 ± .00	.03 ± .00
MEGA	.06 ± .02	.39 ± .12	.68 ± .03	.04 ± .01	.01 ± .00	.01 ± .00
INGENIOUS	.60 ± .02	<u>.41 ± .18</u>	.45 ± .03	<u>.06 ± .01</u>	.01 ± .00	.05 ± .01
INGENIOUS - Info	.68 ± .02	.48 ± .03	<u>.48 ± .06</u>	.05 ± .02	.01 ± .00	.05 ± .01
INGENIOUS - Negative	.13 ± .02	.18 ± .00	.22 ± .06	.02 ± .01	.01 ± .00	.02 ± .01
INGENIOUS - Negative - Info	<u>.61 ± .10</u>	.15 ± .09	.32 ± .00	.03 ± .01	.01 ± .00	.03 ± .00

graphs are then embedded using the embedding module on its own, taking the importance weights w_{uv} of the remaining edges into account during the message-passing operation of the module. The final metric is the area under the curve (AUC) of dissimilarity scores across different percentages of edges removed. For graph classification, the dissimilarity score is the Euclidean distance between the embeddings of the pruned graphs and the embedding of the initial graph. For node classification, we perform this experiment for each node, taking only the k -neighbourhood (k being the depth of the model) of that node into consideration, and we take the average dissimilarity score. In order to ease comparisons between models, these dissimilarity scores are normalised by scaling, to be equal to 1 when all edges are removed. Removing important edges should change the embedding. Hence, removing them first should result in a larger AUC.

To evaluate completeness, we use the opposite faithfulness. It seeks to assess the complementary statement to faithfulness: Are the edges marked as unimportant truly unimportant? To measure it, we apply the same procedure as for the faithfulness metric, except edges are removed in increasing order of predicted importance. In theory, removing unimportant edges should not impact the embedding. The AUC should therefore be small.

To summarise the results intuitively, we focus on the faithfulness gap, that is, the difference in value between the faithfulness and the opposite faithfulness. A larger difference is better. Results are presented in Table 1. Faithfulness and opposite faithfulness results are presented in supplementary material. The model we propose has a clear advantage, achieving a high faithfulness gap in all cases, even higher than the supervised GSAT model. These results demonstrate the superior correctness and completeness of INGENIOUS.

Continuity Continuity refers to the smoothness of the interpretability function (Nauta et al., 2023). We seek to evaluate whether inputs for which the model response is similar, that is, inputs with similar embeddings, also have similar interpretations. However, much like faithfulness, no metrics have been defined in the literature to evaluate this in an unsupervised setting. To evaluate it anyway, we need to define *(dis)similarity*, both between embeddings and between interpretations. While embeddings lie in a Euclidean space, where the L2-norm can be used as a dissimilarity measure, the interpretations lie in the space of graphs. There is no consensual method to compare any two graphs or sub-graphs. Indeed,

a large variety of methods have been proposed (McCabe et al., 2021). Amongst them, we choose to use the node-degree-based Wasserstein distance as a dissimilarity measure, as it is amongst the simplest to compute and the most intuitive. In line with the literature on graph interpretability and augmentation, which focuses solely on graph topology, this measure disregards node features. It also has the advantage of enabling its use with any graph dataset. The continuity metric we propose could be adapted with any other similarity metric with minor changes, but the adaptation would not significantly alter its theoretical analysis. The study of possible adaptations is therefore left for future work.

Formally, given a graph G , we obtain its interpretation G^{exp} by applying a hard-threshold. We keep a fixed percentage of the most important edges, arbitrarily set to 10%. However, for datasets for which there is a prior on the number of important edges for a downstream task, we use that prior number instead of 10%.

The Wasserstein distance (WD) between two graphs G_1 and G_2 serves as the dissimilarity between their interpretations and is defined as $d_W(G_1, G_2) = W_1(\nu(G_1^{exp}), \nu(G_2^{exp}))$, with $\nu(G)$ the distribution of node degrees in G and W_1 the 1-Wasserstein metric. Let $\pi(G)$ be the local neighbourhood of a graph G , defined as the 10% closest graphs of the dataset in the embedding space according to the Euclidean distance. To evaluate the continuity of interpretations, we compare two values:

- $W_1^{local} = \mathbb{E}_i \left[\mathbb{E}_{G_j \in \pi(G_i)} [d_W(G_i, G_j)] \right]$, the expected WD between a graph and another graph in the local neighbourhood of the first one;
- $W_1^{global} = \mathbb{E}_i [\mathbb{E}_j [d_W(G_i, G_j)]]$, the expected WD between any two graphs.

To save time, the expected values are computed on a subset of 6 batches of graphs, i.e. 1536 graphs. A big difference between these two values would be a good indicator of continuity, as it would mean that graphs that are similar in the embedding space have interpretations that are more similar to each other than to the interpretations of other graphs. To summarise the results intuitively, we focus on the difference $W_1^{global} - W_1^{local}$; a larger gap indicates a stronger continuity. We report these differences in Table 1. We see the advantage provided by the negative term of the INGENIOUS loss to produce an embedding space with continuous interpretations, as visible by the higher scores of INGENIOUS and INGENIOUS - Info compared to INGENIOUS - Negative and INGENIOUS - Negative - Info, on *BA2Motifs* and *SPMotifs.5*. As opposed to *Mutag*, these datasets have clear distinct motifs by class and therefore most benefit from the repulsive effect of the negative loss.

Compactness and readability Previous work has shown that sparsity is an important factor in human understanding (Nauta et al., 2023), as it is hindered by human cognitive capacity limitations, and that interpretations should be sparse to not overwhelm the user. Moreover, for inherently interpretable methods as INGENIOUS, sparsity is related to faithfulness as non-sparse interpretations would mean that all features are used by the model to produce the embeddings. Nevertheless, we argue that sparsity, alone, can be misleading, since it can favour interpretations with globally lower weights. An ideal interpretation should be bi-modal, with weights of 1 for important edges and 0 for unimportant ones. We, however, relax this assumption and consider that an interpretation is easy to read if the distribution of weights has two modes. To estimate this, we approximate the density

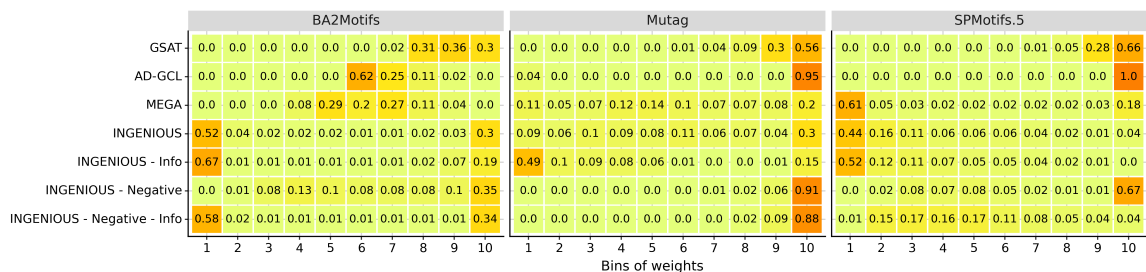


Figure 4: Density function of edge weights. Optimal interpretations should have two density peaks and maximal separation. INGENIOUS results in importance values that are much sparser and closer to extreme values.

function of edge weights by using a histogram with ten bins between 0 and 1. We also define the **sparsity index** as the average of the edge importance weights.

Figure 4 shows that on simple datasets such as *BA2Motifs*, all methods produce sparse and bimodal interpretations, even without regularisation. However, more complex datasets such as *Mutag* and *SPMotifs.5* highlight the need for regularisation. On *Mutag* for example, adding the regularisation lowers the proportion of scores between 0.9 and 1 from 90% to 30%. Both AD-GCL and GSAT tend to collapse on high value and narrow-support densities, as opposed to MEGA which has importance scores closer to extreme values.

Coherence Coherence assesses the alignment with domain knowledge. Sometimes, in the supervised setting, human prior can be used to define expected interpretations that can be compared with generated interpretations. In the literature, this is referred to as *the clicking game* (Wickstrøm et al., 2023). However, it relies on the prior that the model focuses on the same features as the domain expert would, which is a strong assumption (Petsiuk et al., 2018). In an unsupervised setting, it is even harder to define relevant expected interpretations for embeddings. A way to approach this is to use synthetic datasets only; we use *BA2Motifs*, *SPMotifs.5*, *Tree-grid* and *Tree-cycle*. In these datasets, base graphs are generated from the same distribution. Then, different motifs are added to each random graph (Wu et al., 2022; Ying et al., 2019; Luo et al., 2020) depending on the dataset. For graph classification datasets, the motif also depends on the label of the graph. Therefore, we hypothesise that the base graph can be considered as noise and that all the distinctive information of each graph is in the motifs.

Under this hypothesis, any good unsupervised embedding method should rely on the motifs only. This is partially validated by the good accuracy shown in Section 4.2, as it is impossible to have good accuracy on the synthetic datasets if the embeddings contain no information about the motifs. The motifs can therefore be used as expected interpretations.

To measure how close a produced interpretation is to the expected one, we consider an edge-classification problem. An edge is labelled as important if it is part of the expected interpretation. The predicted probability of an edge being important is given by the produced importance weight of that edge. The area under the ROC curve is reported in Figure 2. In this paper, we call this metric the **interpretability AUC**.

In terms of interpretability AUC, INGENIOUS outperforms its unsupervised competitors. The supervised approach GSAT is better as expected. The ablation study shows that

both regularisation losses are needed to reach a better interpretability AUC, although their advantage is less significant on node classification tasks. This could find an explanation in the link between interpretability and sparsity. We study this link in Section 4.4.

4.4. In-Depth Analysis of the Sparsity

Impact of the sparsity on faithfulness In addition to being a desirable characteristic for interpretations, sparsity impacts contrastive learning by increasing the difference between the augmented views. Figure 5 shows a significant positive correlation between opposite faithfulness (red dots) and sparsity, while there is a negative correlation between faithfulness (blue dots) and sparsity. This suggests there may be a link between higher sparsity and better faithfulness gap. This conclusion applies to all graph-classification tasks. For node-classification tasks, the sparsity has a smaller range of values and presents no correlation with the faithfulness but still a small correlation with the opposite faithfulness. Further results with a random baseline are presented in supplementary materials. Overall, the faithfulness gap is improved by the sparsity.

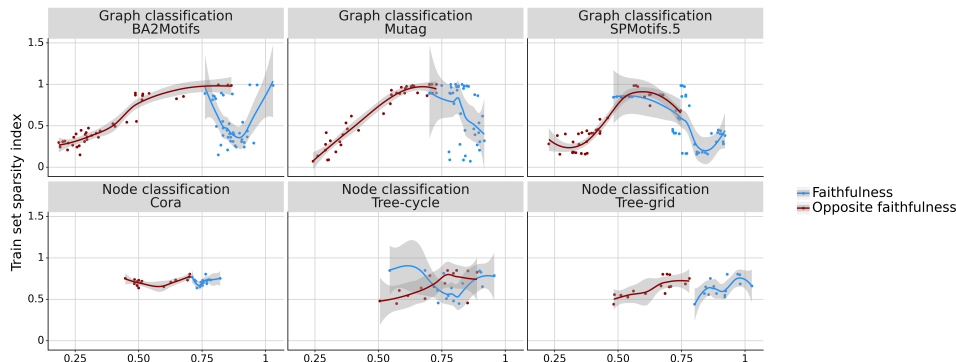


Figure 5: (Opposite) faithfulness with respect to the sparsity index. A lower sparsity results in a lower opposite faithfulness (lower is better) and a marginally higher faithfulness, which means a better faithfulness gap. The grey area represents the 95% confidence interval for a smoothed conditional mean.

Impact of hyperparameters on sparsity This section deals with the impact of the temperature of the Gumbel-max sampling and the impact of the edge-keeping-probability prior r . A high temperature drives the sampling probability closer to a uniform distribution, which increases the diversity of augmented views. Figure 6 shows that sparsity decreases with an increasing temperature value until it reaches a plateau. On the synthetic dataset *BA2Motifs*, a temperature sweet-spot for interpretability AUC is visible at a temperature equal to 20; too high or too low of a temperature is harmful. However, for the other experiments of this paper, to ensure the comparability of our results with those of GSAT, we use the same temperature value as in GSAT, i.e. $\tau = 1$, even though this choice might not result in optimal performances.

The first column of Figure 7 shows that r is positively correlated with the sparsity of INGENIOUS. Lower r leads to a better faithfulness gap, meaning more complete and correct interpretations. Sparser augmented views lead the model to use fewer edges to

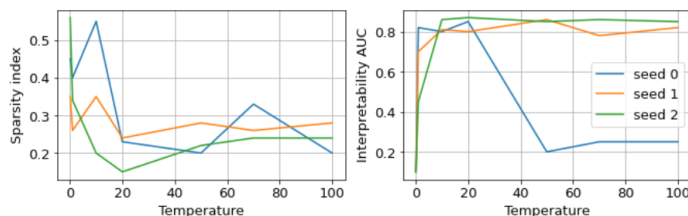


Figure 6: Sparsity and interpretability AUC on *BA2Motifs* with different temperatures. A sweet spot is visible at 20; too high or too low temperatures worsen the results.

produce embeddings, which results in a slight drop in downstream ACC. This phenomenon is less marked on simple datasets where all methods reach a perfect downstream ACC. This trade-off between downstream ACC and interpretability has already been documented in previous work (Jacovi and Goldberg, 2020). Similarly, the sparsity positively impacts the continuity of the interpretability function, as can be seen by the increasing Wasserstein gap. This impact, however, is less strong on *Mutag*. In *INGENIOUS*, the sparsity can be controlled to an extent through regularisation and hyperparameter tuning.

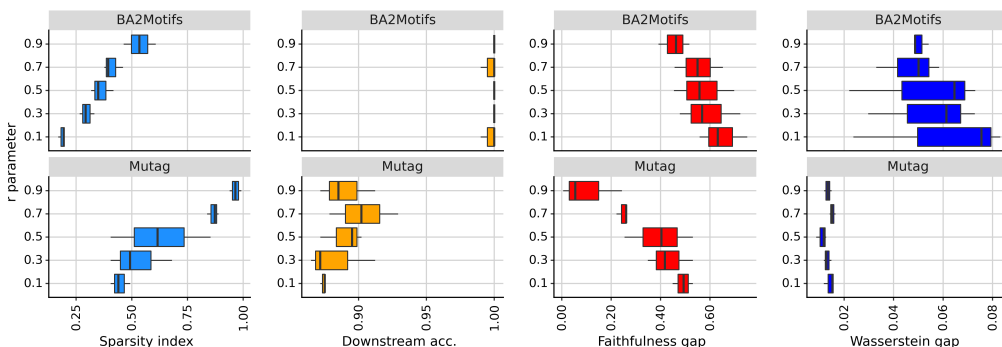


Figure 7: Analysis of the hyperparameter r . A lower r hyperparameter can improve the sparsity, resulting in interpretations that are more correct, complete and robust, with only a slight impact on the downstream ACC.

5. Limitations and Future Work

Like most GNN-interpretability methods (Luo et al., 2020; Ying et al., 2019; Yuan et al., 2020), we limit our interpretability analysis to topology. However, early results in supplementary materials show the possibility of extending *INGENIOUS* to node features. As defined by Doshi-Velez and Kim (2018), we limit our study to a functionally grounded evaluation (which does not require human subjects). Both human-grounded and application-grounded evaluations, which are more specific and costly, are left as future work.

6. Conclusion

We unify previous works on augmentation-based graph representation learning to introduce a new unsupervised embedding approach. Our method, *INGENIOUS*, introduces new losses

for better regularisation. We show their positive impact through an ablation study. Aiming for interpretability, we introduce new metrics to evaluate the quality of interpretations of graph representations learned without supervision, building upon previous work in the supervised setting. Thanks to the approach and metrics we propose, we show that it is possible to produce inherently interpretable embeddings that are also useful for various downstream tasks and that augmented views, when sparse, can be used as interpretations. *INGENIOUS* compares favourably with state-of-the-art unsupervised models in terms of utility of the embeddings and correctness, completeness, continuity and readability of the interpretations. By tuning a reduced set of two hyperparameters, the user can manage the sparsity/utility trade-off. Our work paves the way for novel approaches of intrinsically interpretable unsupervised embeddings for modalities as complex as graph data.

References

- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *Proceedings of the 37th International Conference on Machine Learning, ICML'20*. JMLR.org, 2020.
- Kaize Ding, Zhe Xu, Hanghang Tong, and Huan Liu. Data augmentation for deep graph learning: A survey. *SIGKDD Explor. Newsl.*, 24(2):61–77, dec 2022. ISSN 1931-0145.
- Finale Doshi-Velez and Been Kim. Considerations for evaluation and generalization in interpretable machine learning. *Explainable and interpretable models in computer vision and machine learning*, pages 3–17, 2018.
- David K Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alan Aspuru-Guzik, and Ryan P Adams. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015.
- Wenqi Fan, Yao Ma, Qing Li, Yuan He, Eric Zhao, Jiliang Tang, and Dawei Yin. Graph neural networks for social recommendation. In *The World Wide Web Conference, WWW '19*, page 417–426, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450366748.
- Aosong Feng, Chenyu You, Shiqiang Wang, and Leandros Tassiulas. Kergnns: Interpretable graph neural networks with graph kernels. *CoRR*, abs/2201.00491, 2022.
- Hang Gao, Jiangmeng Li, Wenwen Qiang, Lingyu Si, Fuchun Sun, and Changwen Zheng. Bootstrapping informative graph augmentation via a meta learning approach. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 3001–3007. International Joint Conferences on Artificial Intelligence Organization, 7 2022. Main Track.
- E. J. Gumbel. Statistical theory of extreme values and some practical applications. lectures by emit j. gumbel. national bureau of standards, washington, 1954. 51 pp. diagrams. 40 cents. *The Aeronautical Journal*, 58(527):792–793, 1954.

- William L. Hamilton, Rex Ying, and Jure Leskovec. Representation learning on graphs: Methods and applications. *IEEE Data Eng. Bull.*, 40(3):52–74, 2017.
- Alon Jacovi and Yoav Goldberg. Towards faithfully interpretable NLP systems: How should we define and evaluate faithfulness? In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4198–4205, Online, July 2020. Association for Computational Linguistics.
- Sihang Li, Xiang Wang, An Zhang, Xiangnan He, and Tat-Seng Chua. Let invariant rationale discovery inspire graph contrastive learning. In *ICML, 2022*.
- Shengchao Liu, Hanchen Wang, Weiyang Liu, Joan Lasenby, Hongyu Guo, and Jian Tang. Pre-training molecular graph representation with 3d geometry. In *International Conference on Learning Representations, 2022*.
- Dongsheng Luo, Wei Cheng, Dongkuan Xu, Wenchao Yu, Bo Zong, Haifeng Chen, and Xiang Zhang. Parameterized explainer for graph neural network. *Advances in Neural Information Processing Systems*, 33, 2020.
- Stefan McCabe, Leo Torres, Timothy LaRock, Syed Arefinul Haque, Chia-Hung Yang, Harrison Hartle, and Brennan Klein. netrd: A library for network reconstruction and graph distances. *Journal of Open Source Software*, 6:2990, June 2021.
- Andrew Kachites McCallum, Kamal Nigam, Jason Rennie, and Kristie Seymore. Automating the construction of internet portals with machine learning. *Information Retrieval*, 3: 127–163, 2000.
- Siqi Miao, Mia Liu, and Pan Li. Interpretable and generalizable graph learning via stochastic attention mechanism. *International Conference on Machine Learning, 2022*.
- Christopher Morris, Nils M. Kriege, Franka Bause, Kristian Kersting, Petra Mutzel, and Marion Neumann. Tudataset: A collection of benchmark datasets for learning with graphs. In *ICML 2020 Workshop on Graph Representation Learning and Beyond (GRL+ 2020)*, 2020.
- Meike Nauta, Jan Trienes, Shreyasi Pathak, Elisa Nguyen, Michelle Peters, Yasmin Schmitt, Jörg Schlötterer, Maurice van Keulen, and Christin Seifert. From anecdotal evidence to quantitative evaluation methods: A systematic review on evaluating explainable AI. *ACM Computing Surveys*, feb 2023.
- Vitali Petsiuk, Abir Das, and Kate Saenko. Rise: Randomized input sampling for explanation of black-box models. In *British Machine Vision Conference, 2018*.
- Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1:206–215, 05 2019.
- Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Workshop Track Proceedings, 2014*.

- Arjun Subramonian. Motif-driven contrastive learning of graph representations. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(18):15980–15981, May 2021.
- Susheel Suresh, Pan Li, Cong Hao, and Jennifer Neville. Adversarial graph augmentation to improve graph contrastive learning. In *Advances in Neural Information Processing Systems*, 2021.
- Minh N. Vu and My T. Thai. Pgm-explainer: Probabilistic graphical model explanations for graph neural networks. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- Haonan Wang, Jieyu Zhang, and Wei Huang. Augmentation-free graph contrastive learning with performance guarantee. *ArXiv*, abs/2204.04874, 2022a.
- Haoqing Wang, Xun Guo, Zhi-Hong Deng, and Yan Lu. Rethinking minimal sufficient representation in contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16041–16050, 2022b.
- Kristoffer K. Wickstrøm, Daniel J. Trosten, Sigurd Løkse, Ahcène Boubekki, Karl øyvind Mikalsen, Michael C. Kampffmeyer, and Robert Jenssen. Relax: Representation learning explainability. *Int. J. Comput. Vision*, 131(6):1584–1610, mar 2023. ISSN 0920-5691.
- Ying-Xin Wu, Xiang Wang, An Zhang, Xiangnan He, and Tat-Seng Chua. Discovering invariant rationales for graph neural networks. *CoRR*, abs/2201.12872, 2022.
- Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019.
- Rex Ying, Dylan Bourgeois, Jiaxuan You, Marinka Zitnik, and Jure Leskovec. GNN explainer: A tool for post-hoc explanation of graph neural networks. *CoRR*, abs/1903.03894, 2019.
- Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, and Yang Shen. Graph contrastive learning with augmentations. In *Advances in Neural Information Processing Systems*, volume 33, pages 5812–5823. Curran Associates, Inc., 2020.
- Hao Yuan, Haiyang Yu, Shurui Gui, and Shuiwang Ji. Explainability in graph neural networks: A taxonomic survey. *IEEE transactions on pattern analysis and machine intelligence*, PP, 2020.
- Qinghua Zheng, Jihong Wang, Minnan Luo, Yaoliang Yu, Jundong Li, Lina Yao, and Xiaojun Chang. Towards explanation for unsupervised graph-level representation learning, 2022.
- Yanqiao Zhu, Yichen Xu, Qiang Liu, and Shu Wu. An Empirical Study of Graph Contrastive Learning. *arXiv.org*, September 2021.