

Let Invariant Rationale Discovery Inspire Graph Contrastive Learning

Sihang Li¹ Xiang Wang*² An Zhang³ Ying-Xin Wu⁴ Xiangnan He*⁴ Tat-Seng Chua³

Abstract

Leading graph contrastive learning (GCL) methods perform graph augmentations in two fashions: (1) randomly corrupting the anchor graph, which could cause the loss of semantic information, or (2) using domain knowledge to maintain salient features, which undermines the generalization to other domains. Taking an invariance look at GCL, we argue that a high-performing augmentation should preserve the salient semantics of anchor graphs regarding instance-discrimination. To this end, we relate GCL with invariant rationale discovery, and propose a new framework, Rationale-aware Graph Contrastive Learning (RGCL). Specifically, without supervision signals, RGCL uses a rationale generator to reveal salient features about graph instance-discrimination as the rationale, and then creates rationale-aware views for contrastive learning. This rationale-aware pre-training scheme endows the backbone model with the powerful representation ability, further facilitating the fine-tuning on downstream tasks. On MNIST-Superpixel and MUTAG datasets, visual inspections on the discovered rationales showcase that the rationale generator successfully captures the salient features (*i.e.*, distinguishing semantic nodes in graphs). On biochemical molecule and social network benchmark datasets, the state-of-the-art performance of RGCL demonstrates the effectiveness of rationale-aware views for contrastive learning. Our codes are available at <https://github.com/lsh0520/RGCL>.

¹School of Information Science and Technology, University of Science and Technology of China, Hefei, China ²School of Cyber Science and Technology, University of Science and Technology of China, Hefei, China ³Sea-NExT Joint Lab, National University of Singapore, Singapore ⁴School of Data Science, University of Science and Technology of China, Hefei, China. Correspondence to: Xiang Wang <xiangwang1223@gmail.com>, Xiangnan He <xiangnanhe@gmail.com>.

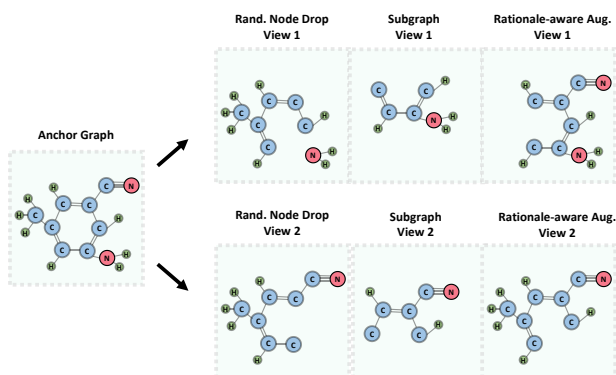


Figure 1. Illustration of graph augmentation in random fashions and rationale-aware fashions. In the molecule graph, the cyano group ($-C\equiv N$) works as the rationale to its toxicity.

1. Introduction

Self-supervised contrastive learning (He et al., 2020; Chen et al., 2020) is attracting a surge of interest in graph neural networks (GNNs) (Wu et al., 2021; Dwivedi et al., 2020), giving rise to graph contrastive learning (GCL). It pre-trains a GNN on a large dataset without relying on handcrafted annotations, which promotes the fine-tuning on downstream tasks (Velickovic et al., 2018; Qiu et al., 2020; You et al., 2020). And GCL has a tremendous impact on real-world applications with abundant unlabeled data, such as the discovery of potential drug-drug interactions in drugs’ pharmacological activity (Wang et al., 2021b).

Inspecting prior studies on GCL, we can systematize the common paradigm as a combination of two modules: (1) graph augmentation, which creates the augmented views of anchor graphs via various approaches, such as node dropping, edge perturbation, and attribute masking; and (2) contrastive learning, which maximizes the agreement for two augmentations of the same anchor, while minimizing the agreement for those of two different anchors. Clearly, graph augmentation is of crucial importance to delineate the discriminative characteristics of graph instances and generate instance-discriminative representations.

State-of-the-art GCL methods conduct graph augmentation either in a random fashion (You et al., 2020; Zhu et al., 2020; Qiu et al., 2020) or in a knowledge-guided fashion (Zhang et al., 2020; Zhu et al., 2021; Wang et al., 2021b). These two lines suffer from inherent limitations correspondingly:

- Graphs usually contain rich saliency information and regularity structure (*e.g.*, functional groups in molecule graphs, coalitions in social networks). Thus, randomly corrupting their properties (*e.g.*, topological structure, node properties, or edge attributes) could lose the discriminative semantics, making the augmented views far from the anchor graphs and misguiding the contrastive learning. Figure 1 showcases random augmentations are likely to pervert the cyano group ($-C\equiv N$), which is the key feature making the molecule hypertoxic, thus easily derailing the discriminative power of the augmented views.
- A research line leverages external domain knowledge to identify the salient semantics of graphs and guide the augmentations. For example, underscoring the functional groups, such as the cyano group for toxicity in Figure 1, is beneficial to molecule representation learning (Rong et al., 2020). Nonetheless, it is prescribed to expensive domain knowledge and easily undermines the generalization performance in unseen domains.

To address these limitations, we revisit the GCL paradigm from the invariance standpoint (Misra & van der Maaten, 2020; Dangovski et al., 2021). Specifically, the contrastive learning module encourages the representation agreement between the anchor graph and its augmented views, *i.e.*, the representations are invariant to the augmentations. Formally, let g be the anchor graph, $A(\cdot)$ be the augmentation function, and $f(\cdot)$ be the encoder network yielding the representations. We can formulate the goal of invariance as $f(A(g)) = f(g)$, where $f(g)$ preserves the salient semantics but contains no information about the way g is augmented.

One naturally raises question: *What augmentations should the representations be invariant to?* To answer it, we relate GCL with invariant rationale discovery (IRD) (Wu et al., 2022; Chang et al., 2020; Li et al., 2022). IRD usually involves two modules: (1) rationale discovery, which can be formulated as a function $R(\cdot)$, extracting features guiding the prediction as the rationale; and (2) prediction, which uses the rationale only for supervised prediction. The prediction is influenced by the rationale, which is invariant regardless of the changes in the rationale’s complement. Considering $f(\cdot)$ as the encoder encapsulating the information to perform prediction, we formulate the aim of IRD as $f(R(g)) = f(g)$. As such, $R(\cdot)$ is supposed to reveal g ’s critical information, which is substantially consistent with the idea of GCL. We might assume that the representations should be invariant to the rationale-aware augmentations that preserve instance-discriminative information.

Motivated by the connection of IRD and GCL, we propose a new framework, Rationale-aware Graph Contrastive Learning (RGCL), to automatically discover rationales as graph augmentations. Specifically, RGCL is a cooperate game between two modules: (1) rationale genera-

tor, which decides fractions to reveal and conceal in the anchor graph, and yields the rationale encapsulating its instance-discriminative information; and (2) contrastive learner, which makes use of rationale-aware views to perform instance-discrimination of graphs. The two players aim for the shared goal of achieving semantically-good representations. On both biochemical molecule and social network benchmark datasets (Wu et al., 2017; Morris et al., 2020), extensive experiments demonstrate the promising performance of RGCL to surpass current state-of-the-art GCL methods (Hu et al., 2020; You et al., 2020; Xu et al., 2021; Suresh et al., 2021).

2. Invariance Look at GCL

We begin with preliminary and related work of GCL and IRD. We then present the invariance view to relate them.

2.1. Graph Contrastive Learning (GCL)

Self-supervised learning has recently achieved promising success in various domains (Kolesnikov et al., 2019; Goyal et al., 2019; Devlin et al., 2018; Brown et al., 2020), thus provoking the interest in GNNs. Its goal is to pre-train a GNN on a large dataset in a self-supervised fashion, so as to enhance the expressive power of GNN and facilitate its fine-tuning on downstream tasks. Such pre-training strategies roughly fall into two research lines. Early studies create augmented views of original graphs and require the learners to predict certain properties from the views, such as graph context (Hu et al., 2020), node attribute (Hu et al., 2020) and edge presence (Hamilton et al., 2017). As such, the resultant representations are covariant to the augmentations (Misra & van der Maaten, 2020; Dangovski et al., 2021).

Recently, studies on GCL have attracted increasing attention, which is the focus of this work. The most commonly-used criterion for GCL is maximizing the mutual information (van den Oord et al., 2018) between two augmented views (Qiu et al., 2020; Velickovic et al., 2018; You et al., 2020). It remains within the confines of the two modules: graph augmentation and contrastive learning.

Graph augmentation. For an anchor graph instance g in the graph set \mathcal{G} , this module employs the augmentation function $A(\cdot)$ to create g ’s views as:

$$g_A = A(g), \quad A \in \mathcal{A}, \quad (1)$$

where $A(\cdot)$ is the function sampled from graph augmentation function set \mathcal{A} to yield the augmented view g_A . We next elaborate the design of $A(\cdot)$ in prior studies.

Most GCL methods instantiate function $A(\cdot)$ in a random fashion by randomly corrupting the anchor’s topological structure, node properties, or edge attributes. For example, GRACE (Zhu et al., 2020) is a hybrid scheme augmenting

on the structure and attribute levels. GCC (Qiu et al., 2020) captures the universal topological properties across multiple graphs. GraphCL (You et al., 2020) systematically studies the impact of combining various random augmentations. Despite promising performance, the intrinsic random nature makes them suffer from potential semantic information loss, thus hardly capturing the salient information. To tackle this problem, domain knowledge is used to identify salient features in graph augmentations. Take the biochemistry domain as an example. MICRO-Graph (Zhang et al., 2020) extracts semantically meaningful motifs (*e.g.*, functional groups of molecules) to construct informative subgraphs. More recently, GraphMVP (Liu et al., 2021b) utilizes the correspondence between 2D topological structures and 3D geometric views of molecules to perform contrastive learning. However, heavily relying on domain knowledge severely undermines their generalization to other domains. In the domain of social networks, GCA (Zhu et al., 2021) exploits node centrality to highlight important connected structures, thus recognizing underlying semantic information. Nevertheless, its intrinsic assumption — node centrality represents saliency — could possibly fail in other domains.

Contrastive learning. Given the augmented views of g , contrastive learning sets up an instance-discrimination task, which can be formulated as follows:

$$\min_f \mathcal{L}_{\text{CL}} = \mathbb{E}_{g \in \mathcal{G}} \mathbb{E}_{A_1, A_2 \sim \mathcal{A}} l_c(f(A_1(g)), f(A_2(g))), \quad (2)$$

where $f(\cdot)$ is the GNN encoder to yield representations, A_1 and A_2 are two functions sampled from graph augmentation function set \mathcal{A} . $l_c(\cdot, \cdot)$ can be instantiated as any contrastive loss function, including NCE (Misra & van der Maaten, 2020), InfoNCE (van den Oord et al., 2018) and NT-Xent (Chen et al., 2020).

Invariance Look. Optimizing Equation (2) encourages the encoder $f(\cdot)$ to produce the same representation for anchor graph g as for its augmented counterparts. That is, the representations are invariant to the augmentation (Misra & van der Maaten, 2020; Dangovski et al., 2021), formally:

$$f(A(g)) = f(g), \quad (3)$$

which essentially enforces the augmented view $g_A = A(g)$ to preserve the critical information in anchor g that is distinguishable from the other anchors.

2.2. Invariant Rationale Discovery (IRD)

Towards interpretability/explainability (Ying et al., 2019; Chen et al., 2018; Wang et al., 2022; 2021a) in supervised learning scenarios, IRD (Chang et al., 2020; Wu et al., 2022; Li et al., 2022) finds a small subset of the input features as rationale, which guides and interprets the prediction. It usually consists of two modules: rationale discovery and classification.

With a slight abuse of notation, we denote the graph, label, rationale and complement variables as the uppercase G , Y , $R(G)$ and $C(G)$, respectively. Correspondingly, the lowercase g , y , $R(g)$ and $C(g)$ represent samples/values of the aforementioned variables.

Rationale Discovery. For a graph instance g , this module extracts a substructure of g as $R(g)$, which is the supporting substructure termed rationale and allows a confident classification alone. Specifically, $R(\cdot)$ is an instantiation of graph augmentation function, preserving critical substructure of g . The rationale can compose salient topological features, node properties, or edge attributes. For example, DIR (Wu et al., 2022) applies an attention network on g 's edges to select the salient edges with top attentive attributions as the rationale of this graph instance.

Classification with Two Principles. Owing to the lack of ground-truth rationales, learning to discover rationale is achieved by implementing function $R(\cdot)$ with a network $r(\cdot)$ and approaching the original input g 's target label y . In general, the oracle rationale needs to satisfy two principles: sufficiency (Yu et al., 2019) and independence (Chang et al., 2020; Wu et al., 2022), which are adopted by the module to push the discovered rationale close to the oracle. By "sufficiency", we mean that the rationale $R(g)$ is sufficient to preserve the critical information of g related to the label y , which is formulated as:

$$p_Y(\cdot | R(g)) = p_Y(\cdot | g), \quad (4)$$

where $p_Y(\cdot | X = x)$ is conditional probability density/mass function conditioned on the random variables corresponding to input x . By "independence", we mean that label variable Y is independent of the complement variable of rationale, $C(G)$, conditioned on the rationale $R(G)$:

$$Y \perp C(G) | R(G), \quad (5)$$

where \perp is probabilistic independence. These two principles shield the classification from the influence of information outside the rationale, so as to enforce the rationale to collect all discriminative features. Moreover, they help to prevent the rationale from degeneration (Yu et al., 2019), where trivial information is collected.

To model these principles, IRD usually devises an encoder network $f(\cdot)$ to generate the rationale representations and a subsequent classifier network $\phi(\cdot)$ to predict g 's label:

$$\begin{aligned} \min_{\phi, f, r} \mathcal{L}_{\text{IRD}} &= \mathbb{E}_{(g, y) \in \mathcal{D}} l_s(\phi(f(R(g))), y) \\ \text{s.t.} \quad & Y \perp C(G) | R(G), \end{aligned} \quad (6)$$

where \mathcal{D} is the dataset involving with the pairs of graph instances and target labels; $l_s(\cdot, \cdot)$ measures the supervised loss like cross-entropy.

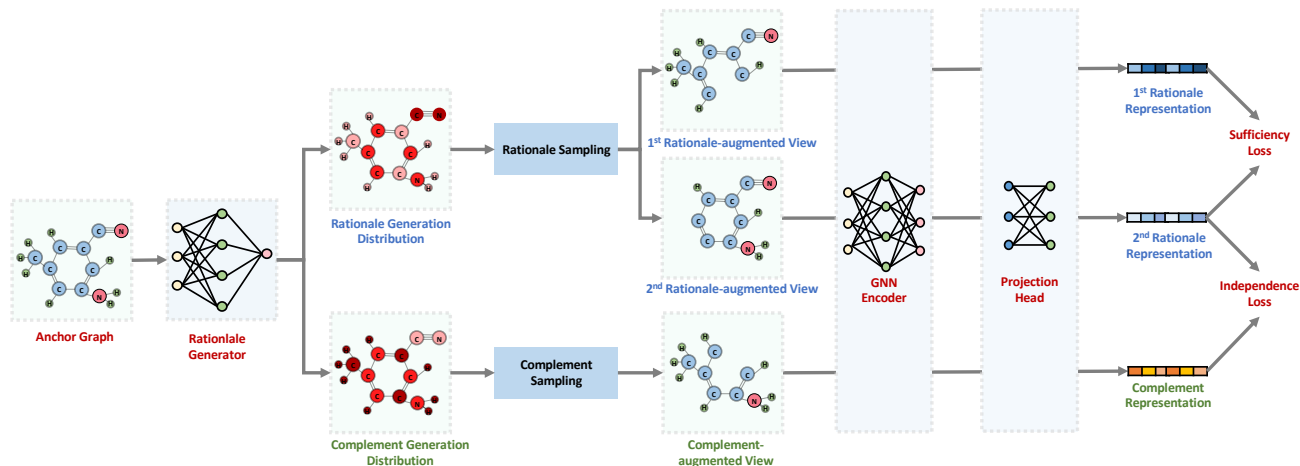


Figure 2. RGCL framework. A rationale generator identifies a discriminative subset of nodes in the original graph. Darker color indicates higher sampling probability in the rationale/complement generation distribution. Two rationales and one complement are generated according to Equation (8) and Equation (9). The rationale generator, shared GNN-encoder and shared projector are jointly optimized via minimizing both the sufficiency loss (cf. Equation (16)) and independence loss (cf. Equation (17)).

Invariance Look. We also take an invariance look at IRD. First, the sufficiency principle (cf. Equation (4)) enforces the encoder $f(\cdot)$ to refine the same information for rationale $R(g)$ as for its original graph g . Second, the independence principle (cf. Equation (5)) naturally makes classification insensitive to the rationale’s complement. For any specific g , we are finding a rationale $R(g)$ as follows:

$$f(R(g)) = f(g), \quad Y \perp C(G)|R(G). \quad (7)$$

2.3. Connecting GCL and IRD

Jointly analyzing the invariance views on GCL and IRD (cf. Equations (3) and (7)), we find the rationale discovery in IRD suitable to address the limitations of graph augmentations in existing GCL frameworks, answering *What augmentations should the representations be invariant to?*. Specifically, when conducting instance-discrimination in GCL, we could extract the rationale from the anchor graph, which latches on instance-discriminative information. We then reveal the rationale that the instance-discrimination task should be invariant to in the augmented views, while concealing its complement.

3. Methodology

Inspired by IRD, we propose a new pre-training scheme, Rationale-aware Graph Contrastive Learning (RGCL), which discovers rationales to garner instance-discriminative semantics, and then performs contrastive learning on the rationale-aware views. We next present the pipeline of RGCL as shown in Figure 2: rationale-aware graph augmentation, representation learning, and contrastive learning.

3.1. Rationale-aware Graph Augmentation

Prior study (You et al., 2020) has shown that, compared with edge perturbation and attribute adjusting, node dropping benefits downstream tasks across different categories of graph datasets. Thus, given an anchor graph, we focus on identifying a subset of salient nodes with edges between them as its rationale.

Towards this end, we need to assess the attribution of each node to instance-discrimination. Moreover, to maintain the diversity of augmented views, it is not sagacious to model the node contribution and rationale generation in a deterministic manner. Hence, we adopt the idea of probabilistic sampling upon an approximation: given an anchor graph $G = g$, its rationale $R(G)$ follows a probability distribution $P_R(R(G)|G = g)$, which summarizes the probability of each node being salient:

$$\begin{aligned} P_R(R(G) = R(g)|G = g) \\ = \prod_{v \in \mathcal{V}_R} p(v|g) \prod_{v \in \mathcal{V}_C} (1 - p(v|g)), \end{aligned} \quad (8)$$

where \mathcal{V} and \mathcal{V}_R are the node sets of g and its rationale $R(g)$, respectively; $\mathcal{V}_C = \mathcal{V} \setminus \mathcal{V}_R$ is the node sets of the complement $C(g)$; $p(v|g)$ denotes the probability of v being included into $R(g)$, reflecting how semantically important it is. Analogously, we can define the distribution of the rationale’s complement $C(G)$ as:

$$\begin{aligned} P_C(C(G) = C(g)|G = g) \\ = \prod_{v \in \mathcal{V}_C} (1 - p(v|g)) \prod_{v \in \mathcal{V}_R} p(v|g), \end{aligned} \quad (9)$$

where $1 - p(v|g)$ measures how deficient node v is to accomplish instance-discrimination.

We now present how to implement the foregoing process. Here we hire a rationale generator network $r(\cdot)$ to parameterize the probability distribution function $p(\cdot|g)$:

$$\mathbf{P} = r(g), \quad (10)$$

where $r(\cdot)$ is a GNN-MLP combined encoder that takes the anchor graph g as input and yields normalized node attribution scores $\mathbf{P} \in \mathbb{R}^{|\mathcal{V}| \times 1}$, where the v -th element in \mathbf{P} corresponds to aforementioned probability $p(v|g)$.

Further, we sample rationale-aware views from distribution $P_R(\cdot|G = g)$ to obtain rationale-aware views:

$$R(g) \sim P_R(\cdot|G = g) \quad \text{s.t.} \quad |\mathcal{V}_R| = \rho \cdot |\mathcal{V}|, \quad (11)$$

where based on normalized nodes' attribution scores \mathbf{P} , we sample $\rho \cdot |\mathcal{V}_g|$ nodes from the original graph g , while keeping the edges between sampled nodes. Similarly, rationale complement views are obtained as:

$$C(g) \sim P_C(\cdot|G = g) \quad \text{s.t.} \quad |\mathcal{V}_C| = \rho \cdot |\mathcal{V}|. \quad (12)$$

Note that $C(g)$ is not the node difference between g and $R(g)$, but a stochastic complement sampled from the distribution $P_C(\cdot|G = g)$. As a result, each node within $R(g)$ (or $C(g)$) is assigned with its probability $p(v|g)$ (or $1 - p(v|g)$), which illustrates how crucial (or trivial) it is to conduct instance-discrimination in GCL.

3.2. Rationale-aware Representation Learning

After sampling from these two distributions, we have $R(g)$ and $C(g)$, while discarding the remaining nodes. For the rationale-augmented view $R(g)$, we then associate it with the attribution vector $\mathbf{P}_R \in \mathbb{R}^{|\mathcal{V}_R| \times 1}$, where we keep node attribute scores in \mathbf{P} corresponds to the node set of $R(g)$. The complement view $C(g)$ is processed similarly with the attribution vector $\mathbf{P}_C \in \mathbb{R}^{|\mathcal{V}_C| \times 1}$.

Having established the rationale-augmented view $R(g)$ with the probability vector \mathbf{P}_R , we feed them into the GNN backbone $f(\cdot)$ (*i.e.*, the target model being pre-trained) to generate the rationale-aware representation:

$$\mathbf{x}_R = f(R(g)) = \text{Pooling}(\text{GNN}(R(g)) \odot \mathbf{P}_R), \quad (13)$$

where $f(\cdot)$ is a combination of the base encoder $\text{GNN}(\cdot)$ and the pooling layer $\text{Pooling}(\cdot)$, and yields the d' -dimensional rationale representation \mathbf{x}_R . To be more specific, $\text{GNN}(\cdot)$ outputs $\mathbf{X}_R \in \mathbb{R}^{|\mathcal{V}_R| \times d'}$ that collects the representations of nodes within $R(g)$. Subsequently, we apply the element-wise product between \mathbf{X}_R and \mathbf{P}_R , and then employ the pooling function $\text{Pooling}(\cdot)$ to compress the node representations into the rationale representation. Note that, when fine-tuning on downstream tasks, we disable the

rationale discovery module and discard \mathbf{P}_R in $f(\cdot)$, *i.e.*, $\mathbf{x} = \text{Pooling}(\text{GNN}(g))$.

Moreover, we resort to a projection head $h(\cdot)$ to map the graph representation into another latent space where the contrastive learning is conducted, aiming to enforce mutual information between the anchor and rationale to a tighter lower bound (Chen et al., 2020). Formally, aforementioned process is shown as:

$$\mathbf{r} = h(\mathbf{x}_R) = L_2(\text{MLP}(\mathbf{x}_R)), \quad (14)$$

where $h(\cdot)$ is instantiated by an MLP with l_2 normalized outputs. Similarly, we can get the representation and the projection of the complement-aware view $C(g)$:

$$\mathbf{x}_C = f(C(g)), \quad \mathbf{c} = h(\mathbf{x}_C). \quad (15)$$

3.3. Rationale-aware Contrastive Learning

Given the rationale-aware graph augmentations, we move on to achieving the sufficiency and independence principles (*cf.* Equations (4) and (5)) in contrastive learning.

Toward this end, for an anchor graph g , we randomly sample two rationales from its rationale generation distribution, view them as the positive pair $(R_1(g), R_2(g))$, and establish their projected representations \mathbf{r}_1^+ and \mathbf{r}_2^+ via Equation (14). Following the previous work (Chen et al., 2020), the rationales of other anchors are treated as the negative views of g , and summarized their representations into \mathcal{R}_g^- . We then model the sufficiency principle as minimizing the following contrastive loss:

$$l_{\text{su}}(g) = -\log \frac{\exp(\mathbf{r}_1^{+\top} \mathbf{r}_2^+ / \tau)}{\sum_{\mathbf{r}^- \in \mathcal{R}_g^-} \exp(\mathbf{r}_1^{+\top} \mathbf{r}^- / \tau)}, \quad (16)$$

where τ is a temperature hyperparameter. It encourages the agreement between the positive rationale-aware views of the same anchor graph, while enforcing the divergence between different anchors. Through minimizing $l_{\text{su}}(g)$, rationale generator is required to refine the crucial information in the anchor graph about instance-discrimination, which has a causal relation with semantics (*cf.* Equations (4)).

Going beyond the negative views derived from other anchors' rationales, we further sample one complement $R^c(g)$ from its complement generation distribution, and cast it as the additional negative views of g . And we formalize the independence principle as the minimization of the contrastive loss below:

$$l_{\text{in}}(g) = -\log \frac{\exp(\mathbf{r}_1^{+\top} \mathbf{r}_2^+ / \tau)}{\exp(\mathbf{r}_1^{+\top} \mathbf{r}_2^+ / \tau) + \sum_{\mathbf{c} \in \mathcal{C}} \exp(\mathbf{r}_1^{+\top} \mathbf{c} / \tau)}, \quad (17)$$

where set \mathcal{C} summarizes all complement representations appearing in the minibatch data. The minimization of $l_{\text{in}}(g)$

pushes away the representation of the complement, *i.e.*, \mathbf{c} , from that of the rationale, *i.e.*, \mathbf{r} , making the captured rationale stable regardless of changes of its complement, which consists with the independence principle in Equation (5).

Finally, our objective function conflates these two types of losses, depicting the cooperative game between the rationale generator network $r(\cdot)$ and the target backbone model $f(\cdot)$:

$$\min_{r,f,h} \mathcal{L}_{\text{RGCL}} = \mathbb{E}_{g \in \mathcal{G}} [l_{\text{su}}(g) + \lambda \cdot l_{\text{in}}(g)], \quad (18)$$

where λ is the hyperparameter to control the tradeoff between $l_{\text{su}}(g)$ and $l_{\text{in}}(g)$. After optimization, we throw away the projection head $h(\cdot)$ when fine-tuning on downstream tasks. It is worth mentioning that our RGCL is a backbone-agnostic self-supervised graph learning framework, which can be applicable to different backbone models.

4. Experiments

In this section, extensive experiments are conducted to answer two research questions:

- **RQ1:** How effective is the rationale generator in capturing semantically important nodes to construct contrastive views invariant to augmentations?
- **RQ2:** Whether rationale-aware views promote GCL, further improving the performance of pre-trained backbone model on downstream tasks?

4.1. Effectiveness of Rationale Generator (RQ1)

To justify the effectiveness of our rationale generator, we conducted experiments on two datasets from different domains.

Real world biochemical molecule dataset. Following the settings of previous works (Hu et al., 2020; You et al., 2020), we use Zinc-2M – 2 million unlabeled molecule graphs sampled from the ZINC15 database (Sterling & Irwin, 2015) to pre-train the backbone model and rationale generator. Then the pre-trained rationale generator is applied on a real world molecule benchmark dataset – Mutag. Figure 3a presents the rationales labeled by two chemistry experts (the first row in yellow) and those discovered by our RGCL (the second row in green), to interpret “What part of the molecule graphs are critical to the instance-discriminative property?”. Note that the rationale discovered by our rationale generator is the probabilities of each node, so we highlight the top- k nodes in green, where k is the number of nodes highlighted by the experts in corresponding graphs. Treating the expert-driven rationales as the ground truth, the precision of RGCL-discovered rationales is 78.9%. This comparison justifies the reliability of our rationale discovery. Details of experiment settings are included in Appendix C.

Table 1. Classification accuracy (%) comparison on MNIST-Superpixel. Numbers are from You et al. (2020). **Bold** indicates the best performance while underline indicates the second best.

DATASETS	1% MNIST	10% MNIST
NO PRE-TRAIN	60.39 ± 1.95	79.71 ± 0.65
AUG.	67.43 ± 0.36	83.99 ± 2.19
GAE	57.58 ± 2.07	86.67 ± 0.93
INFOMAX	63.24 ± 0.78	83.34 ± 0.24
GRAPHCL	<u>83.41 ± 0.33</u>	<u>93.11 ± 0.17</u>
RGCL (OURS)	83.84 ± 0.41	94.92 ± 0.12

Superpixel dataset. Following the setting of You et al. (2020), we formulate the pre-training dataset by omitting the labels of the train set in MNIST-Superpixel dataset and utilize it to pre-train our backbone encoder and rationale generator jointly. After pre-training, finetuning is performed in a supervised manner with 1% and 10% data samples with labels randomly chosen from the train set. Then we randomly pick two cases to visualize the rationale generated by the rationale generator, which is shown in Figure 3b. And Table 1 shows the performance gain compared with multiple strong baselines. Due to the limitation of space, experiment settings including the briefs of MNIST-Superpixel and baselines are summarized in Appendix B.

We have the following observations on MNIST-Superpixel experiment results in Figure 3b and Table 1:

(1) Our rationale generator is capable of discovering semantic nodes (*e.g.* dark nodes in the original superpixel graph) and the generated rationale helps to preserve these semantic nodes in generated rationale-augmented views, thus providing correct guidance for the following contrastive learning process. Consequently, RGCL outperforms GraphCL, a typical GCL framework performing augmentations in a random fashion and other baselines.

(2) The nodes’ color distribution in the visualized rationale graph follows a dark-light-dark pattern from center to edge, which attracts our attention and an intuitive explanation is provided as follows. Semantic nodes in the center of the graph are captured by the rationale generator and assigned the highest probability, so we call this zone semantic zone. However, sampling nodes close to semantic ones when generating augmented views could possibly lead to semantic information change, for example, adding a node in the center of a θ superpixel graph may change the semantic information from θ to δ and adding a node at the up-left part of a I superpixel graph makes semantic information change from I to 7 . Therefore we name this zone near the semantic zone as confusion zone and nodes in this area are supposed to be assigned the lowest sampling probability in order to avoid semantic information change. Nodes at the edge part of the graph contain no semantic information but including some of them in the augmented views benefits




Figure 3. (a) The first row presents the rationales labeled by chemistry experts and the second row presents those discovered by our rationale generator. (b) From top to bottom are original MNIST data, MNIST-Superpixel graph and the rationale generated by the rationale generator in RGCL. The darker color indicates higher sampling probability in rationale graphs.

Table 2. Transfer learning ROC-AUC scores (%) on downstream graph classification tasks compared with state-of-the-art methods where statistics are from You et al. (2020) except GraphLoG and AD-GCL, whose results are reproduced on our platform. **Bold** indicates the best performance while underline indicates the second best on each dataset.

DATASET	BBBP	Tox21	ToxCast	SIDER	CLINTOX	MUV	HIV	BACE	AVG.	GAIN
NO PRE-TRAIN	65.8±4.5	74.0±0.8	63.4±0.6	57.3±1.6	58.0±4.4	71.8±2.5	75.3±1.9	70.1±5.4	67.0	-
ATTRMASKING	64.3±2.8	76.7±0.4	64.2±0.5	61.0±0.7	71.8±4.1	74.7±1.4	77.2±1.1	79.3±1.6	71.1	4.1
CONTEXT PRED	68.0±2.0	<u>75.7±0.7</u>	<u>63.9±0.6</u>	60.9±0.6	65.9±3.8	<u>75.8±1.7</u>	<u>77.3±1.0</u>	<u>79.6±1.2</u>	70.9	3.9
GRAPHCL	69.68±0.67	73.87±0.66	62.40±0.57	60.53±0.88	75.99±2.65	69.80±2.66	78.47±1.22	75.38±1.44	70.77	3.77
GRAPHLOG*	71.04±1.86	74.65±0.60	62.32±0.51	57.86±1.44	78.72±2.58	74.95±1.96	75.12±1.98	82.6±1.25	72.16	5.16
AD-GCL*	<u>68.26±1.03</u>	73.56±0.72	63.10±0.66	59.24±0.86	<u>77.63±4.21</u>	74.94±2.54	75.45±1.28	75.02±1.88	70.90	3.90
RGCL (OURS)	71.42±0.66	75.20±0.34	63.33±0.17	61.38±0.61	83.38±0.91	76.66±0.99	<u>77.90±0.80</u>	76.03±0.77	73.16	6.16

the diversity of the augmented views, making the backbone model more robust. Hence, these nodes are supposed to be sampled with moderate probabilities and the edge part is named as background zone. During the contrastive loss optimization process, the rationale generator is trained to discover the semantic nodes and push away data samples from each other to prevent confusion. Consequently, from the center part to edge part, the semantic zone - confusion zone - background zone follows a dark-light-dark pattern.

Summary. The overall results of these two experiments answered **RQ1** and demonstrated the effectiveness of our proposed rationale generator.

4.2. Performance on Downstream Tasks (RQ2)

Transfer learning. We first pre-train a backbone model on Zinc-2M (Sterling & Irwin, 2015) and then finetune it on 8 benchmark multi-task binary classification datasets in biochemistry domain, which are contained in MoleculeNet (Wu et al., 2018). Note that the downstream datasets are split using scaffold split in order to simulate the real world case — out-of-distribution — and examine the generalization ability of the pre-trained model. Details of experiment settings are included in Appendix C.

We adopt AttrMasking (Hu et al., 2020), ContextPred (Hu et al., 2020), GraphCL (You et al., 2020), GraphLoG (Xu

et al., 2021) and AD-GCL (Suresh et al., 2021), which are state-of-the-art pre-training paradigms in this area, as our baselines. Finetuning procedure is repeated for 10 times with different random seeds and we evaluate the mean and standard deviation of ROC-AUC scores on each downstream dataset, which is consistent with our baselines. Further, the average and its gain compared with the training-from-scratch model are listed as well. All transfer learning performance on downstream tasks are presented in Table 2.

With the guidance of rationale to construct semantic information preserved views, our RGCL framework achieves best performance on 4 out of 8 datasets and highest average gain compared with existing baselines. Compared with GraphCL (You et al., 2020), our framework RGCL leverages an rationale generator to construct rationale-aware views and a mixture loss function in Equation (18) to ensure both sufficiency and independency for generated rationales. So we go a step further to compare the performance of GraphCL with RGCL to demonstrate the crucial role of rationale-augmented views in GCL. Rationale-augmented views in RGCL preserve more semantic information in anchor graphs, thus making it outperform GraphCL in 7 out of 8 downstream datasets except HIV, and raising the average performance gain from 3.77 to 6.16. As for the performance deterioration on HIV dataset, one possible reason is that the rationale probability distribution in HIV is different from

Table 3. Unsupervised representation learning classification accuracy (%) on TU datasets. The compared numbers are from You et al. (2020) except AD-GCL, whose statistics are reproduced on our platform. **Bold** indicates the best performance while underline indicates the second best on each dataset.

DATASET	NCII	PROTEINS	DD	MUTAG	COLLAB	RDT-B	RDT-MSK	IMDB-B	AVG.
NO PRE-TRAIN	65.40±0.17	72.73±0.51	75.67.4±0.29	87.39±1.09	65.29±0.16	76.86 ±0.25	48.48±0.28	69.37±0.37	70.15
INFOGRAPH	76.20±1.06	74.44±0.31	72.85±1.78	89.01±1.13	70.05±1.13	82.50±1.42	53.46±1.03	73.03±0.87	74.02
GRAPHCL	77.87±0.41	74.39±0.45	78.62±0.40	86.80±1.34	71.36±1.15	89.53±0.84	55.99±0.28	71.14±0.44	75.71
AD-GCL	73.91±0.77	73.28±0.46	75.79±0.87	88.74±1.85	72.02±0.56	90.07±0.85	54.33±0.32	70.21±0.68	74.79
OURS	78.14±1.08	75.03±0.43	78.86±0.48	87.66±1.01	70.92±0.65	90.34±0.58	56.38±0.40	71.85±0.84	76.15

Table 4. Ablation study for RGCL on downstream transfer learning datasets.

DATASET	BBBP	Tox21	ToxCast	SIDER	CLINTOX	MUV	HIV	BACE	AVG.
RGCL w/o RV	70.08±0.86	74.01±0.56	62.23±0.38	60.38±0.52	77.29±1.81	70.12±1.28	77.42±0.94	75.08±1.22	70.83
RGCL w/o I	70.64±0.92	75.02±0.48	63.12±0.56	61.20±0.56	84.26±1.55	73.64±1.08	77.08±1.21	75.88±0.94	72.60
RGCL	71.42±0.66	75.20±0.34	63.33±0.17	61.38±0.61	83.38±0.91	76.66±0.99	77.90±0.80	76.03±0.77	73.16

that in the pre-training dataset ZINC15, making the pre-trained model fail to capture discriminative features in HIV dataset. Hence it is challenging to train a robust rationale generator with great generalization ability, which we will explore in the future.

Unsupervised learning. Furthermore, we follow the settings of InfoGraph (Sun et al., 2020) to evaluate RGCL in the unsupervised representation learning where the encoded graph representations are fed into a non-linear SVM classifier. We compare RGCL with four self-supervised learning baselines on TU datasets (Morris et al., 2020): untrained GIN (Xu et al., 2019), InfoGraph (Sun et al., 2020), GraphCL (You et al., 2020) and AD-GCL (Suresh et al., 2021) with the default settings in Sun et al. (2020). Summaries of datasets and baselines are provided in Appendix D. Table 3 shows that RGCL outperforms the baselines on 5 out of 8 cases and achieves the best average performance.

Summary. Extensive experiments on transfer learning and unsupervised representation learning demonstrate the state-of-the-art performance of RGCL, answering **RQ2**: rationale-aware views benefit GCL, endowing the pre-trained model with better transferability and generalization ability.

4.3. Ablation Study

In this section, ablation studies are conducted to demonstrate the importance of rationale-views and independence term l_{in} in loss function (cf. Equation (18)). We first replace our rationale-aware data augmentation module with one in a random fashion (i.e., randomly discarding certain nodes to construct augmented views). As illustrated in Table 4, the performance of this new GCL framework without rationale views (RGCL w/o RV), decreases sharply compared with RGCL, demonstrating the crucial role of rationale-aware views in guiding contrastive learner. Note that RGCL w/o RV can be viewed as GraphCL (You et al., 2020) but a pooling layer with attention mechanism. Then we ablate the final loss function (e.g. eliminating independence term l_{in} in it).




Figure 4. Sensitivity w.r.t. hyperparameter τ , λ and ρ .

The performance drop of RGCL without independence term (RGCL w/o I) demonstrates the strength of independence principle in IRD, which is regulating the information outside the rationale.

4.4. Hyperparameter sensitivity

RGCL’s sensitivity w.r.t. hyperparameters τ , λ and ρ is evaluated by presenting the average performances on both transfer learning and unsupervised settings as in Figure 4. And we have following observations:

- As verified in Wang & Liu (2021), τ plays a crucial role in contrastive learning. A too small value of τ (e.g., 0.05) hurts RGCL’s performance; continuously increasing its value (i.e., over 0.2) leads to performance drop as well. Thus we suggest to tune τ around 0.1 carefully on both MoleculeNet (Wu et al., 2018) and TU (Morris et al., 2020) datasets. Moreover, according to our results, different datasets may desire different settings of τ .
- The independence term l_{in} works as a regularizer in optimization process, so the performance drops when overemphasizing it (i.e., setting a large value for λ). Nevertheless,

a proper setting of regularizer (*i.e.*, $\lambda = 0.1$) prevents the model from degeneration, benefiting contrastive learning.

- Hyperparameter ρ controls the scale of the augmented views. We suggest to tune it around a comparatively large value (*e.g.*, 80%), because a too small value enlarges the distribution gap between augmented views (*i.e.*, graphs of small scales) and anchor graphs (*i.e.*, graphs of original scales), undermining the representative ability of pre-trained model on downstream tasks.

5. Conclusion

In this paper, we revisit graph contrastive learning (GCL) from the standpoint of invariant rationale discovery (IRD) and propose a novel graph contrastive learning framework, Rationale-aware Graph Contrastive Learning (RGCL) for GNN pre-training. In our framework, a GNN-MLP combined rationale generator requires no domain knowledge nor manually designed features to generate rationales - discovering semantic information related nodes in graph instances. Under the correct guidance of the rationale, semantic information preserved views are generated to conduct contrastive learning. Visual inspections on MNIST-Superpixel dataset demonstrate the effectiveness of rationale generator by revealing its capability of discovering discriminative nodes. Extensive experiments conducted on biochemical molecule and social network benchmark datasets show that rationale-aware views promote pre-trained models' transferability and generalization ability, thus achieving state-of-the-art performance compared with baselines.

In future, we will explore the fine-tuning of rationale generator on downstream tasks, such that the pre-training scheme in GCL endows the rationales with a local view of self-discriminative patterns and the fine-tuning will enhance the rationales' global view of class-wise discriminative patterns. It is promising to probe into explainability of GCL. Moreover, we would like to investigate the retrospective and introspective learning of rationale discovery, which in turn guide the discrimination tasks and improve the generalization ability of the backbone models (Liu et al., 2021a).

6. Acknowledgements

This work is supported by the National Key Research and Development Program of China (2021ZD0111802) and National Natural Science Foundation of China (U19A2079, U21B2026, U1936210). This research is also supported by CCD Key Lab of Ministry of Culture and Tourism and Sea-NExT Joint Lab.

References

- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. Language models are few-shot learners. In *NeurIPS*, 2020.
- Chang, S., Zhang, Y., Yu, M., and Jaakkola, T. S. Invariant rationalization. In *ICML*, volume 119, pp. 1448–1458, 2020.
- Chen, J., Song, L., Wainwright, M. J., and Jordan, M. I. Learning to explain: An information-theoretic perspective on model interpretation. In *ICML*, volume 80, pp. 882–891, 2018.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. E. A simple framework for contrastive learning of visual representations. In *ICML*, volume 119, pp. 1597–1607, 2020.
- Dangovski, R., Jing, L., Loh, C., Han, S., Srivastava, A., Cheung, B., Agrawal, P., and Soljagic, M. Equivariant contrastive learning. *CoRR*, abs/2111.00899, 2021.
- Devlin, J., Chang, M., Lee, K., and Toutanova, K. BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805, 2018.
- Dwivedi, V. P., Joshi, C. K., Laurent, T., Bengio, Y., and Bresson, X. Benchmarking graph neural networks. *CoRR*, 2020.
- Goyal, P., Mahajan, D., Gupta, A., and Misra, I. Scaling and benchmarking self-supervised visual representation learning. In *ICCV*, pp. 6391–6400, 2019.
- Hamilton, W., Ying, Z., and Leskovec, J. Inductive representation learning on large graphs. *NeurIPS*, 30, 2017.
- He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pp. 9729–9738, 2020.
- Hu, W., Liu, B., Gomes, J., Zitnik, M., Liang, P., Pande, V. S., and Leskovec, J. Strategies for pre-training graph neural networks. In *ICLR*, 2020.
- Kipf, T. N. and Welling, M. Variational graph auto-encoders. *CoRR*, abs/1611.07308, 2016.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. In *ICLR (Poster)*. OpenReview.net, 2017.

- Kolesnikov, A., Zhai, X., and Beyer, L. Revisiting self-supervised visual representation learning. In *CVPR*, pp. 1920–1929, 2019.
- Li, Y., Wang, X., Xiao, J., Ji, W., and seng Chua, T. Invariant grounding for video question answering. In *CVPR*, 2022.
- Liu, E. Z., Haghgoo, B., Chen, A. S., Raghunathan, A., Koh, P. W., Sagawa, S., Liang, P., and Finn, C. Just train twice: Improving group robustness without training group information. In *ICML*, volume 139, pp. 6781–6792, 2021a.
- Liu, S., Wang, H., Liu, W., Lasenby, J., Guo, H., and Tang, J. Pre-training molecular graph representation with 3d geometry. *CoRR*, abs/2110.07728, 2021b.
- Misra, I. and van der Maaten, L. Self-supervised learning of pretext-invariant representations. In *CVPR*, pp. 6706–6716, 2020.
- Monti, F., Boscaini, D., Masci, J., Rodola, E., Svoboda, J., and Bronstein, M. M. Geometric deep learning on graphs and manifolds using mixture model cnns. In *CVPR*, pp. 5115–5124, 2017.
- Morris, C., Kriege, N. M., Bause, F., Kersting, K., Mutzel, P., and Neumann, M. Tudataset: A collection of benchmark datasets for learning with graphs. *CoRR*, abs/2007.08663, 2020.
- Qiu, J., Chen, Q., Dong, Y., Zhang, J., Yang, H., Ding, M., Wang, K., and Tang, J. Gcc: Graph contrastive coding for graph neural network pre-training. In *SIGKDD*, pp. 1150–1160, 2020.
- Rong, Y., Bian, Y., Xu, T., Xie, W., Wei, Y., Huang, W., and Huang, J. Self-supervised graph transformer on large-scale molecular data. In *NeurIPS*, 2020.
- Sterling, T. and Irwin, J. J. Zinc 15–ligand discovery for everyone. *Journal of chemical information and modeling*, 55(11):2324–2337, 2015.
- Sun, F., Hoffmann, J., Verma, V., and Tang, J. Infograph: Unsupervised and semi-supervised graph-level representation learning via mutual information maximization. In *ICLR*, 2020.
- Suresh, S., Li, P., Hao, C., and Neville, J. Adversarial graph augmentation to improve graph contrastive learning. *CoRR*, abs/2106.05819, 2021.
- van den Oord, A., Li, Y., and Vinyals, O. Representation learning with contrastive predictive coding. *CoRR*, abs/1807.03748, 2018.
- Velickovic, P., Fedus, W., Hamilton, W. L., Liò, P., Bengio, Y., and Hjelm, R. D. Deep graph infomax. *CoRR*, abs/1809.10341, 2018.
- Wang, F. and Liu, H. Understanding the behaviour of contrastive loss. In *CVPR*, pp. 2495–2504, 2021.
- Wang, X., Wu, Y., Zhang, A., He, X., and seng Chua, T. Towards multi-grained explainability for graph neural networks. In *NeurIPS*, 2021a.
- Wang, X., Wu, Y., Zhang, A., Feng, F., He, X., and Chua, T.-S. Reinforced causal explainer for graph neural networks. *TPAMI*, 2022.
- Wang, Y., Min, Y., Chen, X., and Wu, J. Multi-view graph contrastive representation learning for drug-drug interaction prediction. In *WWW*, pp. 2921–2933, 2021b.
- Wu, Y., Wang, X., Zhang, A., He, X., and Chua, T.-S. Discovering invariant rationales for graph neural networks. In *ICLR*, 2022.
- Wu, Z., Ramsundar, B., Feinberg, E. N., Gomes, J., Geniesse, C., Pappu, A. S., Leswing, K., and Pande, V. S. Moleculenet: A benchmark for molecular machine learning. *CoRR*, abs/1703.00564, 2017.
- Wu, Z., Ramsundar, B., Feinberg, E. N., Gomes, J., Geniesse, C., Pappu, A. S., Leswing, K., and Pande, V. Moleculenet: a benchmark for molecular machine learning. *Chemical science*, 9(2):513–530, 2018.
- Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., and Yu, P. S. A comprehensive survey on graph neural networks. *TNNLS*, 32(1):4–24, 2021.
- Xu, K., Hu, W., Leskovec, J., and Jegelka, S. How powerful are graph neural networks? In *ICLR*, 2019.
- Xu, M., Wang, H., Ni, B., Guo, H., and Tang, J. Self-supervised graph-level representation learning with local and global structure. In *ICML*, volume 139, pp. 11548–11558, 2021.
- Ying, Z., Bourgeois, D., You, J., Zitnik, M., and Leskovec, J. Gnnexplainer: Generating explanations for graph neural networks. In *NeurIPS*, pp. 9240–9251, 2019.
- You, Y., Chen, T., Sui, Y., Chen, T., Wang, Z., and Shen, Y. Graph contrastive learning with augmentations. *NeurIPS*, 33:5812–5823, 2020.
- Yu, M., Chang, S., Zhang, Y., and Jaakkola, T. S. Rethinking cooperative rationalization: Introspective extraction and complement control. In *EMNLP/IJCNLP*, pp. 4092–4101, 2019.

Zhang, S., Hu, Z., Subramonian, A., and Sun, Y. Motif-driven contrastive learning of graph representations. *CoRR*, abs/2012.12533, 2020.

Zhu, Y., Xu, Y., Yu, F., Liu, Q., Wu, S., and Wang, L. Deep graph contrastive representation learning. *CoRR*, abs/2006.04131, 2020.

Zhu, Y., Xu, Y., Yu, F., Liu, Q., Wu, S., and Wang, L. Graph contrastive learning with adaptive augmentation. In *WWW*, pp. 2069–2080, 2021.

A. Rationale-aware Graph Contrastive Learning (RGCL) algorithm

Algorithm 1 RGCL algorithm

Initialize: dataset $\{g_m : m = 1, 2, \dots, M\}$, encoder $f(\cdot)$, rationale generator $r(\cdot)$, projector $h(\cdot)$, sampling ratio ρ and hyperparameter λ .

for sampled minibatch of data $\{g_n : n = 1, 2, \dots, N\}$ **do**

for $n = 1$ **to** N **do**

$\mathbf{P}_n = r(g_n)$

$R'(g_n) \sim P_R(R(G)|G = g_n), \quad R''(g_n) \sim P_R(R(G)|G = g_n)$ # Equation (11)

$R^c(g_n) \sim P_C(R^c(G)|G = g_n)$ # Equation (12)

$\mathbf{r}'_n = h(f(R'(g_n))), \quad \mathbf{r}''_n = h(f(R''(g_n)))$

$\mathbf{c}_n = h(f(R^c(g_n)))$

end for

for $n = 1$ **to** N **do**

$\mathcal{R}_n^- = \{\mathbf{r}'_i, \mathbf{r}''_i : i = 1, 2, \dots, n-1, n+1, \dots, N\}$

$\mathcal{C} = \{\mathbf{c}_i : i = 1, 2, \dots, N\}$

$l_{\text{su}}^n = -\log \frac{\exp(\mathbf{r}'_n \top \mathbf{r}''_n / \tau)}{\sum_{\mathbf{r}^- \in \mathcal{R}_n^-} \exp(\mathbf{r}'_n \top \mathbf{r}^- / \tau)}$ # Equation (16)

$l_{\text{in}}^n = -\log \frac{\exp(\mathbf{r}'_n \top \mathbf{r}''_n / \tau)}{\exp(\mathbf{r}'_n \top \mathbf{r}''_n / \tau) + \sum_{\mathbf{c} \in \mathcal{C}} \exp(\mathbf{r}'_n \top \mathbf{c} / \tau)}$ # Equation (17)

end for

$\mathcal{L}_{\text{RGCL}} = \frac{1}{N} \sum_{n=1}^N (l_{\text{su}}^n + \lambda \cdot l_{\text{in}}^n)$ # Equation (18)

 Update $f(\cdot)$, $r(\cdot)$, and $h(\cdot)$ to minimize $\mathcal{L}_{\text{RGCL}}$

end for

return Encoder $f(\cdot)$

B. MNIST-superpixel Settings

Datasets We summarize the statistics of MNIST-superpixel dataset in Table 5. We refer to Monti et al. (2017) for more detailed information if needed.

Table 5. Statistics for MNIST-superpixel dataset.

DATASETS	CATEGORY	GRAPHS#	AVG. N#	AVG. DEGREE
MNIST	SUPERPIXEL GRAPHS	70,000	70.57	8

Baselines In semi-supervised MNIST-superpixel experiment, we adopt baselines including: (1) training the model from scratch (*aka.* **No Pre-Train**) and that with graph augmentations but without contrastive learning process (*aka.* **Aug.**). Following the settings of previous work (You et al., 2020), we also take an adjacency information reconstruction method **GAE** (Kipf & Welling, 2016), a local & global mutual information maximization method **Infograph** (Sun et al., 2020) and a contrastive learning method with a random fashion **GraphCL** (You et al., 2020) for comparison.

C. Transfer Learning Settings

Datasets We utilize MoleculeNet (Wu et al., 2017) as downstream tasks for transfer learning and summarize the statistics which are from You et al. (2020) in Table 6.

Baselines In our implementation, we choose various state-of-the-art self-supervised pre-training frameworks as baselines for transferability comparison:

- **Attribute Masking** (Hu et al., 2020) Attribute Masking learns the regularities of the node/edge attributes distributed over graph structure to capture inherent domain knowledge.
- **Context Prediction** (Hu et al., 2020) Context Prediction predicts subgraphs’ surrounding graph structures to pre-train a backbone GNN so that it maps nodes appearing in similar structural contexts to nearby representations.

Table 6. Statistics for transfer learning MoleculeNet datasets.

DATASETS	CATEGORY	UTILIZATION	GRAPHS#	AVG. N#	AVG. DEGREE
ZINC-2M	BIOCHEMICAL MOLECULES	PRE-TRAINING	2,000,000	26.62	57.72
BBBP	BIOCHEMICAL MOLECULES	FINETUNING	2,039	24.06	51.90
Tox21	BIOCHEMICAL MOLECULES	FINETUNING	7,831	18.57	38.58
ToxCast	BIOCHEMICAL MOLECULES	FINETUNING	8,576	18.78	38.52
SIDER	BIOCHEMICAL MOLECULES	FINETUNING	1,427	33.64	70.71
CLINTox	BIOCHEMICAL MOLECULES	FINETUNING	1,477	26.15	55.76
MUV	BIOCHEMICAL MOLECULES	FINETUNING	93,087	24.23	52.55
HIV	BIOCHEMICAL MOLECULES	FINETUNING	41,127	25.51	54.93
BACE	BIOCHEMICAL MOLECULES	FINETUNING	1,513	34.08	73.71

- **GraphCL** (You et al., 2020) GraphCL learns unsupervised representations of graph data via contrastive learning with graph augmentations in a random fashion.
- **GraphLoG** (Xu et al., 2021) GraphLoG uses the hierarchical prototypes to capture the global semantic clusters while preserving the local similarities.
- **AD-GCL** (Suresh et al., 2021) AD-GCL optimizes adversarial graph augmentation strategies used in GCL to avoid capturing redundant information.

D. Unsupervised Learning Settings

Datasets We summarize the statistics of TU-datasets (Morris et al., 2020) for unsupervised learning in Table 7.

Table 7. Statistics for unsupervised learning TU-datasets.

DATASETS	CATEGORY	GRAPHS#	AVG. N#	AVG. DEGREE
NCII	BIOCHEMICAL MOLECULES	4,110	29.87	1.08
PROTEINS	BIOCHEMICAL MOLECULES	1,113	39.06	1.86
DD	BIOCHEMICAL MOLECULES	1,178	284.32	715.66
MUTAG	BIOCHEMICAL MOLECULES	188	17.93	19.79
COLLAB	SOCIAL NETWORKS	5,000	74.49	32.99
RDT-B	SOCIAL NETWORKS	2,000	429.63	1.15
RDT-M	SOCIAL NETWORKS	2,000	429.63	497.75
IMDB-B	SOCIAL NETWORKS	1,000	19.77	96.53

Baselines To demonstrate the effectiveness of our proposed framework, we adopt multiple strong unsupervised learning baselines, such as Infograph (Sun et al., 2020), GraphCL (You et al., 2020) and AD-GCL (Suresh et al., 2021) as our baselines. Note that the briefs of aforementioned baselines can be found in Appendix B and C.

E. Model Structure and Hyperparameters

To make a fair comparison, we follow the backbone model settings in You et al. (2020). Our model architectures, the mainbody of which includes GCN (Kipf & Welling, 2017) and GIN (Xu et al., 2019), and corresponding hyperparameters are summarized in Table 8.

Note that $\{\cdot\}$ indicates the grid search ranges we tune these hyperparameters in and **bold** numbers are our final settings.

F. Complexity Analysis

We make a comparison on complexity between RGCL and GraphCL in this section. In terms of space complexity, besides the backbone encoder, although RGCL introduces another rationale generator to capture rationales in anchor graphs, space complexity will not be significantly increased since the rationale generator is a much smaller network compared with the backbone encoder and its scale can be adjusted if necessary.

Table 8. Model architectures and hyperparameters

EXPERIMENTS	TRANSFER LEARNING	UNSUPERVISED LEARNING	MNIST-SUPERPIXEL
BACKBONE GNN TYPE	GIN	GIN	GIN
BACKBONE NEURON#	[300,300,300,300,300]	[32,32,32]	[110,110,110,110]
RATIONALE GEN. GNN TYPE	GCN	GIN	GCN
RATIONALE GEN. GNN NEURON#	[128,64,32]	[32,32]	[16,16]
RATIONALE GEN. MLP NEURON#	[32,1]	[32,1]	[16,1]
PROJECTOR NEURON#	[300,300]	[32,32]	[110,110]
POOLING LAYER	GLOBAL MEAN POOL	GLOBAL ADD POOL	GLOBAL ADD POOL
LEARNING RATE lr	{0.0001, 0.001 , 0.01}	{0.001, 0.01 , 0.1}	{0.001, 0.005 , 0.01}
SAMPLING RATIO ρ	{0.6, 0.7, 0.8 , 0.9}	{0.7, 0.8, 0.9 }	{0.7, 0.8 , 0.9}
TEMPERATURE τ	{0.05, 0.1 , 0.2, 0.3, 0.5}	{0.1, 0.2 , 0.3}	{0.1, 0.2, 0.3, 0.5 , 0.7}
λ IN EQUATION 18	{0.05, 0.1 , 0.3, 0.5, 1.0}	{ 0.1 , 0.5, 1.0}	{0.1, 0.2 , 0.5}
TRAINING EPOCHS	{60, 80 , 100}	{10, 20 , 40}	{60, 100 , 150}

Then we move on to time complexity analysis of pre-training process. Suppose the average numbers of nodes and edges per graph in pre-training datasets are $|\mathcal{V}|$ and $|\mathcal{E}|$, respectively. Let B denote the batch size, ρ denote the sampling ratio, L_B denote the number of GNN layers in backbone encoder, L_R denote that in rationale generator and d denote the latent space dimension where contrastive loss is calculated. The time complexity comparison of processing a minibatch of training data is presented in Table 9.

Table 9. The comparison of time complexity between RGCL and GraphCL

FRAMEWORKS	GRAPHCL	RGCL
DATA AUGMENTATION	$O(2B\rho \mathcal{V} \log \mathcal{V})$	$O(3B\rho \mathcal{V} \log \mathcal{V})$
RATIONALE GENERATOR PROPAGATION	—	$O((\mathcal{E} ^2 + \mathcal{V})L_R B)$
BACKBONE ENCODER PROPAGATION	$O((2 \mathcal{E} ^2 + \mathcal{V})L_B B)$	$O((3 \mathcal{E} ^2 + \mathcal{V})L_B B)$
CONTRASTIVE LOSS	$O(B^2 d)$	$O(2B^2 d)$

As illustrated in Figure 2, RGCL has a three-tower structure while GraphCL has a two-tower one. The extra one stems from the calculation of the complements of rationales to regulate the information outside the rationale. And in realistic implementation on our platform (GeForce RTX 2080 Ti and Intel(R) Core(TM) i9-9900X), it takes about 15 hours to finish GraphCL pre-training while 20 hours are needed for RGCL, which roughly matches our theoretical analysis above.