
AutoGDA: Automated Graph Data Augmentation for Node Classification

Anonymous Author(s)

Anonymous Affiliation

Anonymous Email

Abstract

Graph data augmentation has been used to improve generalizability of graph machine learning. However, by only applying fixed augmentation operations on entire graphs, existing methods overlook the unique characteristics of communities which naturally exist in the graphs. For example, different communities can have various degree distributions and homophily ratios. Ignoring such discrepancy with unified augmentation strategies on the entire graph could lead to sub-optimal performance for graph data augmentation methods. In this paper, we study a novel problem of automated graph data augmentation for node classification from the localized perspective of communities. We formulate it as a bilevel optimization problem: finding a set of augmentation strategies for each community, which maximizes the performance of graph neural networks on node classification. As the bilevel optimization is hard to solve directly and the search space for community-customized augmentations strategy is huge, we propose a reinforcement learning framework AutoGDA that learns the local-optimal augmentation strategy for each community sequentially. Our proposed approach outperforms established and popular baselines on public node classification benchmarks as well as real industry e-commerce networks by up to +12.5% accuracy.

1 Introduction

Data augmentation methods are widely used to improve the generalizability and robustness of machine learning (ML) models [1]. They aim to create plausible variations of existing data without the need of additional human efforts. It has been proved that customized data augmentation, i.e., customizing augmentation strategies for each (batch of) object, are beneficial for ML models [2–4]. For example, customized augmentation strategies [2, 5] have shown improved performance over having a uniform augmentation on the entire dataset. To this end, automated data augmentation methods efficiently seek the optimal customized augmentation strategies for samples/batches [2, 4–6].

Recently, with graph neural networks (GNNs) [7–10] emerging as one of the preferred approaches for learning on graph structured data, graph data augmentation methods [11–17] have shown promising results in improving GNNs. For example, DropEdge [18] randomly removes a fraction of edges in each training epoch to promote GNN’s robustness during test-time inference. The AdaEdge [11] approach iteratively adds (or removes) edges between nodes that are predicted to have the same (or different) labels with high confidence. GAugM and GAugO [13] manipulate the graph structure according to edge probabilities learned by link predictors. Despite the promising improvements on various node classification tasks, existing graph data augmentation approaches *are manually designed for the entire graph and only explore graph properties and characteristics globally.*

It is more involved to apply automated data augmentation on graphs compared to images and text, because of the unique properties of graph data bring a great challenge to the effort. While existing automated augmentation approaches [2, 5] assume that samples are independent and identically distributed (i.i.d.) in the dataset, nodes in the graph are naturally connected and are dependent on each other in a non-Euclidean manner. Therefore, it is not straightforward to apply existing automated augmentation methods for graph data. On the other side, the unique properties of graph data may give us some clue to design new and effective solutions. Nodes in the graph are naturally grouped into

43 communities [19, 20], providing a natural separation of data objects (nodes) for node classification.
 44 Chiang et al. [21] show that nodes from the same community are the most important neighbors for
 45 aggregation-based graph learning algorithms. As communities in graphs such as social networks are
 46 usually disparate in characteristics [22–24] such as density, centrality, homophily, etc., we argue that
 47 data augmentation strategies should be localized (community-specific) to achieve optimal results.
 48 However, how to augment graph data according to the localized characteristics of communities in the
 49 graph remains underexplored.

50 To address the aforementioned challenges, we propose to tackle down the problem of automated
 51 graph data augmentation from the local perspective, i.e., communities in graphs. We first analyze
 52 the disparate characteristics of communities using benchmark datasets. Motivated by observations
 53 and insights, we define automated graph data augmentation as a bilevel optimization problem,
 54 that is, to learn the augmentation strategies that lead to the best node classification performance
 55 of GNNs. As finding the optimal augmentation strategies requests combinatorial optimization, it
 56 is impractical in real world due to huge computational cost. We propose AutoGDA that learns
 57 community-customized augmentation strategies with a reinforcement learning (RL) approach, in-
 58 spired by the auto-augmentation literature in computer vision [2, 5]. Specifically, given communities
 59 in a graph, AutoGDA relies on an RL-agent to sequentially pick up the optimal strategy from several
 60 graph data augmentation operations for each community. The RL-agent in AutoGDA generalizes the
 61 learning and selection of augmentation method from one community to another, and thus automates
 62 and accelerates the process of finding localized augmentation strategies.

63 We conduct extensive experiments across different GNN backbones and datasets to evaluate AutoGDA
 64 against state-of-the-art baselines. We demonstrate that AutoGDA with traditional community detec-
 65 tion algorithms (e.g., the Louvain method [25]) and existing graph data augmentation operations
 66 (DropEdge [18], GAUGM [13], and AttrMask [26]) can achieve consistent performance improvements
 67 over the baselines. Specifically, AutoGDA shows up to 12.5% over the best-performed baseline
 68 method. Moreover, we show that the graph representations learned by AutoGDA are robust against
 69 graph adversarial attacks [27].

70 Our main contributions are as follows.

- 71 • We tackle down the problem of automated graph data augmentation for supervised node classifi-
 72 cation by proposing community-customized augmentations from a localized perspective. To
 73 the best of our knowledge, we are the first to investigate community-customized graph data
 74 augmentation for the task of node classification.
- 75 • We propose AutoGDA, an RL-based framework that automatically learns optimal community
 76 customized graph data augmentation strategies. The AutoGDA framework is flexible on the
 77 augmentation operations and can be easily generalized to heterogeneous graphs.
- 78 • We conduct extensive experiments on eight benchmark datasets (including two real industrial
 79 graph anomaly detection benchmarks) with three widely used GNN backbones to validate
 80 AutoGDA. The experimental results show that (1) AutoGDA consistently outperforms state-of-
 81 the-art graph data augmentation baselines across all datasets, and (2) AutoGDA learns robust
 82 representations that give comparable or better classification performance than state-of-the-art
 83 graph defending methods against adversarial attacks.

84 2 Preliminaries and Problem Definition

85 **Notations** Let $G = (\mathcal{V}, \mathcal{E})$ be an undirected graph of N nodes, where $\mathcal{V} = \{v_1, v_2, \dots, v_N\}$ is the
 86 set of nodes and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the set of edges. We denote the adjacency matrix as $\mathbf{A} \in \{0, 1\}^{N \times N}$,
 87 where $A_{i,j} = 1$ indicates nodes v_i and v_j are connected. We denote the node feature matrix as
 88 $\mathbf{X} \in \mathbb{R}^{N \times F}$, where F is the number of features and \mathbf{x}_i (the i -th row of \mathbf{X}) indicates the feature
 89 vector of node v_i . We denote the node labels for classification as $\mathbf{y} \in \{1, \dots, M\}^N$, where M
 90 is the number of classes. We denote the set of graph communities as $\mathcal{C} = \{C_1, C_2, \dots, C_{N_c}\}$
 91 where N_c is the number of communities and each community C_k is defined by a set of nodes \mathcal{V}_{C_k}
 92 s.t. $\mathcal{V}_{C_i} \cap \mathcal{V}_{C_j} = \emptyset, \forall i, j \in \{1, 2, \dots, N_c\}$ and $i \neq j$. We denote the subgraph containing the
 93 community C_k as $G_{C_k} = (\mathcal{V}_{C_k}, \mathcal{E}_{C_k})$, where $\mathcal{E}_{C_k} \subseteq \mathcal{V}_{C_k} \times \mathcal{V}_{C_k}$ is the set of edges within this
 94 subgraph. With a bit of notation abuse, we use the union symbol to denote the combination of
 95 subgraphs, i.e., $G = \bigcup_{k=1}^{N_c} G_{C_k}$.

96 **Graph Neural Networks** Without the loss of generality, we take the commonly used graph convolutional network (GCN) [7] as an example when explaining GNNs in the following sections. The graph convolution operation of each GCN layer is defined as $\mathbf{H}^{(l)} = \sigma(\tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}} \mathbf{H}^{(l-1)} \mathbf{W}^{(l)})$, where 97 98 99 l is the layer index, $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}$ is the adjacency matrix with added self-loops, $\tilde{\mathbf{D}}$ is the diagonal degree matrix $\tilde{D}_{ii} = \sum_j \tilde{A}_{ij}$, $\mathbf{H}^{(0)} = \mathbf{X}$, $\mathbf{W}^{(l)}$ is the learnable weight matrix at the l -th layer, and 100 101 $\sigma(\cdot)$ denotes a nonlinear activation such as ReLU.

102 **Graph Data Augmentation** We follow prior literature [13] to classified graph data augmentation methods for node classification into two categories: stochastic operations for original-graph setting and deterministic operations for modified-graph setting. Let $h : G \rightarrow G_m$ be a graph data augmentation operation that generates a variant G_m of the original graph G . In the original-graph setting, h can be stochastic and applying it for T times results with T graph variants G_m , such that $G \cup \{G_m^i\}_{i=1}^T$ is used in training while only G is used for inference. On the other hand, in the modified-graph setting, 103 104 105 106 107 108 h is deterministic and outputs one G_m , such that G_m replaces G for both training and inference.

109 In this work, we consider four typical state-of-the-art graph data augmentation operations for node classification: $\mathcal{A} = \{\text{DROPEdge}, \text{ATTRMASK}, \text{GAUGM_ADD}, \text{GAUGM_RM}\}^1$, which include 110 111 112 both stochastic and deterministic operations and they apply on both graph structure and the node features. It's worth noting that our proposed AutoGDA is not limited to these four augmentation operations and can take any graph data augmentation operations in \mathcal{A} . 113

114 DROPEdge [18] and ATTRMASK [26] are stochastic augmentation operations, where they randomly drop (mask) a given percentage of edges (node attributes) in each training epoch of GNN. DROPEdge implies the graph structure has certain robustness to the edge connectivity and also alleviates the well-known over-smoothing problem of GNNs [18]. ATTRMASK encourages GNNs to recover masked node attributes with their context information in the local neighborhood. On the other hand, GAUGM_ADD and GAUGM_RM [13] are deterministic augmentation operations that deterministically modify the graph structure to promote the graph's homophily and hence improve the model's performance for node classification [13]. GAUGM_ADD, and GAUGM_RM use the predicted edge probabilities for all node pairs by VGAE [28] and add new edges (remove existing edges) with highest (lowest) edge probabilities. 115 116 117 118 119 120 121 122 123

124 Each of the above four augmentation operations has one parameter controlling the percentage magnitude of the augmentation, resulting with a 100^4 searching space when finding the optimal strategy for each dataset. As we separate the graph into multiple communities and aim to search for the optimal community customized augmentation strategy, the searching space would be $100^{4 \times N_c}$, where N_c is the number of communities. As N_c gets larger, the searching space becomes infeasible for traditional parameter searching methods such as grid search. Therefore, a new efficient approach for automated graph augmentation search is desired. 125 126 127 128 129 130

131 **Problem Definition** Following the definition of previous literature [13, 18] on graph data augmentation for node classification, the problem of finding a hand-crafted graph data augmentation strategy can be defined as follows. Given the graph data G , a set of graph data augmentation operations \mathcal{A} , and GNN model $f : G \rightarrow \hat{y}$, find the best strategy of applying \mathcal{A} on G such that the node classification performance of f on G is maximized. 132 133 134 135

136 In this work, as we propose to find the best augmentation strategy for each community in the graph, the problem of automated graph data augmentation is then defined as: **given** the graph data G , graph communities \mathcal{C} in G (the community detection method for generating \mathcal{C} is treated as a hyperparameter), a set of graph data augmentation operations \mathcal{A} , and GNN model $f : G \rightarrow \hat{y}$, **find** the best strategy of applying \mathcal{A} on the subgraphs of each community $C_k \in \mathcal{C}$ such that the node classification performance of f on G is maximized. The main difference between our problem definition and the previous literature is the use of graph communities \mathcal{C} to find the best data augmentation strategies. 137 138 139 140 141 142

¹To differentiate from the baseline methods (normal font), we use the small caps font to denote the augmentation operations.




Figure 1: Graph communities detected by the Louvain method on the PubMed dataset show diverse distribution on different characteristics of graph structure.

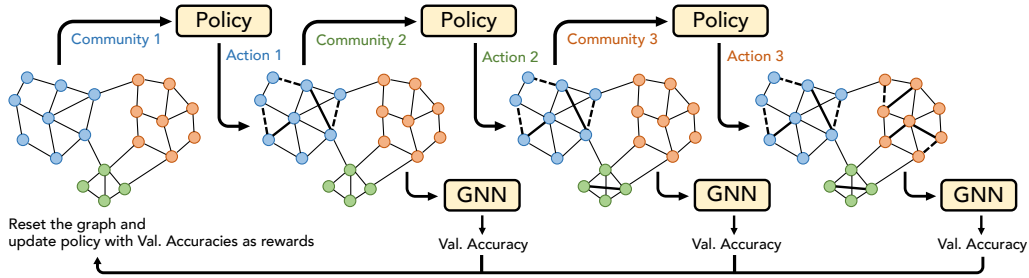


Figure 2: Overview of one iteration of our proposed AutoGDA on an example graph with three communities. In each step, the policy network takes the observation of one graph community as input and outputs the augmentation strategy for it. The GNN is then fine-tuned with the augmented graph and the validation accuracy is used as the reward to update the policy network.

143 **3 Automated Graph Data Augmentation**

144 Section 3.1 shows our motivation of customizing graph data augmentation strategies for different
 145 communities. Section 3.2 models automated graph data augmentation as a bilevel optimization
 146 problem. Section 3.3 presents the reinforcement learning-based framework AutoGDA.

147 **3.1 Motivation**

148 Community structures naturally exist in graphs [19, 20] and graph community detection has been
 149 extensively studied in the past few decades [29]. Graph community detection methods (e.g., the
 150 Louvain method [25]) separates the set of nodes in the graph into disjoint subsets such that the quality
 151 of the communities, which is usually measured by modularity [20], are maximized. Thus, the nodes
 152 within the same community are more densely connected and also more important to each other for
 153 node classification [21, 30] comparing with the nodes in different communities.

154 Our idea is based on the observation that different communities in the same graph mostly show
 155 disparate data distribution, which was also shown in previous literature [22–24]. The structure of
 156 communities commonly varies in terms of density, centrality, etc. For example, Figure 1 shows the
 157 characteristics of communities (detected by Louvain) on the PubMed graph [7]. Figures 1(a), 1(b),
 158 and 1(c) show the distributions of averaged degree, number of triangles, and centrality, respectively.
 159 Figure 1d presents the homophily ratio [31] of community subgraphs, where the homophily ratio
 160 is calculated by the fraction of edges which connect nodes that have the same class label. Previous
 161 works [11, 13] have shown that the graph homophily is strongly correlated with node classification
 162 performance of GNNs, because semi-supervised graph learning methods are mainly based on the
 163 homophily assumption [32]. From Figure 1 we observe that the communities disparate with different
 164 distributions under different measurements.

165 Moreover, for certain deterministic graph data augmentation methods such as GAUGM [13]. The
 166 minority communities may be ignored during the augmentation process. For example, GAUGM [13]
 167 modifies the graph structure according to the edge probabilities given an trained edge predictor,
 168 in which it adds missing edges with highest probabilities and remove existing edges with lowest
 169 probabilities. However, it is possible that all of the modification would happen in only one or few
 170 graph communities, which shows the strongest homophily patterns.

171 With the above observations on the disparate characteristics of graph communities, we argue that
 172 the state-of-the-art methods that apply the augmentation operations on the whole graph may not be
 173 the best practice of graph data augmentation. Auto-augmentation literature in computer vision [2, 5]
 174 has shown that customizing augmentation operations for data objects/batches is more effective than
 175 using the same strategy for the entire dataset. Although it is infeasible to learn the best operation
 176 for each data object (node) for node classification due to the dependent nature of graph data, graph
 177 data augmentation could benefit from having customized augmentation strategy for each community.
 178 The next sub-section formulates the problem of automated graph data augmentation via community
 179 customization.

180 3.2 Bilevel Optimization Formulation

181 We formulate the problem of automated graph data augmentation in a similar way to the auto-
 182 augmentation problem in vision tasks [2, 5]: it aims to find a set of graph data augmentation
 183 operations for each community in the graph, which maximizes the performance of a graph neural
 184 network model on the task of (semi-)supervised node classification.

185 Let the graph data augmentation policy network be defined as $g_\theta : G \rightarrow \{0, 1, \dots, 99\}^{|\mathcal{A}|}$, which is
 186 a multi-layer perceptron (MLP) that is parameterized by θ . The policy takes a (sub)graph as input
 187 and outputs the augmentation strategy for this (sub)graph, which in our case is the four percentage
 188 magnitudes for $\mathcal{A} = \{\text{DROPEdge}, \text{ATTRMASK}, \text{GAUGM_ADD}, \text{GAUGM_RM}\}$. Let $aug(g_\theta(G))$
 189 be a function that applies the augmentation strategy $g_\theta(G)$ on (sub)graph G , then the automated
 190 graph data augmentation process for subgraph G_{C_k} is formulated as

$$G'_{C_k} = aug(g_\theta(G_{C_k})) \quad (1)$$

191 where G'_{C_k} denotes the subgraph G_{C_k} after augmentation.

192 We denote the GNN model as $f_\omega : G \rightarrow \hat{\mathbf{y}}$, which is parameterized by ω . It takes the graph as
 193 input and outputs the predicted node labels $\hat{\mathbf{y}}$. Let the \mathbf{y}_{tr} and \mathbf{y}_{val} be the node labels for training set
 194 and validation set, respectively. The objective of obtaining the best augmentation policy (solving for
 195 θ) could be described as a bilevel optimization problem [33]. The inner level is the model weight
 196 optimization, which is solving for the optimal ω_θ given a fixed augmentation policy (g_θ):

$$\omega_\theta = \arg \min_{\omega} \mathcal{L} \left(f_\omega \left(\bigcup_{k=1}^{N_c} aug(g_\theta(G_{C_k})) \right), \mathbf{y}_{tr} \right), \quad (2)$$

197 where \mathcal{L} denotes the loss function (cross entropy).

198 The outer level is the augmentation policy optimization, which is optimizing the policy parameter θ
 199 using the result of the inner level problem. Here we take the validation performance (accuracy) as the
 200 optimization objective. Then we have the problem formulated as below:

$$\theta^* = \arg \max_{\theta} ACC \left(f_{\omega_\theta} \left(\bigcup_{k=1}^{N_c} aug(g_\theta(G_{C_k})) \right), \mathbf{y}_{val} \right), \quad (3)$$

where $\omega_\theta = \arg \min_{\omega} \mathcal{L}(\omega, \theta)$ (Eq. (2)),

201 where θ^* denotes the parameter of the optimal policy, and $ACC(f(G), \mathbf{y}_{val})$ denotes the validation
 202 accuracy.

203 3.3 AutoGDA Framework

204 As the graph data augmentation operations \mathcal{A} modifies the graph structure and also affects the training
 205 of the GNN model, when applying the augmentation operations community by community, the
 206 graph data augmentation can be formulated as a sequential process on the graph. Figure 2 illustrates
 207 the sequential process: in each step of the iteration, the policy network takes the observation of a
 208 community and outputs the action containing the set of augmentation magnitudes which will be
 209 applied on this community. We finetune the GNN model after applying the augmentations and take
 210 the validation accuracy as rewards.

211 **Parameter Sharing** As the solving of bilevel optimization problems is extremely time consuming
 212 due the repeated solving of the inner loop [33], we utilize the weight sharing scheme for automated

213 augmentation proposed by Tian et al. [5]. At the start of each episode, we reset the current graph to
 214 the original graph, pretrain the GNN model on the original graph (without optimizing the outer loop
 215 or applying any data augmentation operation), and obtain $\bar{\omega}$. That is,

$$\bar{\omega} = \arg \min_{\omega} \mathcal{L}(f_{\omega}(G), \mathbf{y}_{tr}) \quad (4)$$

216 In each step of this episode, instead of training a new GNN model from scratch to get ω_{θ} , we load
 217 the parameters $\bar{\omega}$ from pretraining and finetune the GNN model for only a small number of epochs
 218 with the given actions (augmentation strategy) to get $\bar{\omega}_{\theta}$. Therefore, the outer level for optimizing the
 219 augmentation policy parameters becomes

$$\theta^* = \arg \max_{\theta} ACC \left(f_{\bar{\omega}_{\theta}} \left(\bigcup_{k=1}^{N_c} aug(g_{\theta}(G_{C_k})) \right), \mathbf{y}_{val} \right). \quad (5)$$

220 **Reinforcement Learning (RL) Environment** The set of graph data augmentation operations \mathcal{A}
 221 contains both stochastic and deterministic augmentation operations, where the stochastic operations
 222 (DROPEdge, ATTRMASK) affect the GNN model in training and the deterministic operations
 223 (GAUGM_ADD, GAUGM_RM) directly modifies the graph structure. Therefore, as we apply the
 224 augmentation strategy on each community, the node/graph representation obtained by GNN trained
 225 with the augmentations also changes. This process forms a Markov decision process, whose length is
 226 equal to the number of communities.

227 In the RL environment, we take the current graph with the given augmentation strategy as the state
 228 and use the graph representation of one community as the observation for one step. That is, for each
 229 graph community C_k , the observation is the pooled graph representation of the subgraph G_{C_k} , i.e.,
 230 the element-wise mean of the node representations for all nodes in \mathcal{V}_{C_k} . As the output dimension of
 231 GNN’s last layer is the number of unique labels, which is usually very small, we take the output of the
 232 GNN’s second last layer as the node representations. The policy network g_{θ} takes the observation (i.e.,
 233 graph representation of the community) as input and outputs the magnitudes of different augmentation
 234 operations for the community. Note that the augmentation operation would not be applied if its
 235 magnitude is zero.

236 For optimizing our proposed method, we opt for simplicity and employ the widely-used Proximal
 237 Policy Optimization (PPO) [34] algorithm. We use the validation performance of the GNN model
 238 after finetuning in step as the reward to the RL policy.

239 **Summary** Algorithm 1 summarizes
 240 the whole process of AutoGDA. In
 241 each episode of the policy optimization
 242 stage, we first pretrain the GNN
 243 model by Eq. (4) and reset the graph.
 244 The subgraphs G'_{C_k} in line 3 are for
 245 tracking the augmented communities.
 246 Then for each community, we first obtain
 247 and apply its augmentations strategy
 248 given by the policy (line 5), then
 249 load the pretrained parameters $\bar{\omega}$ for
 250 the GNN model, finetune it with the
 251 current augmentation strategies, and
 252 use the validation accuracy as reward
 253 to update the policy. After the policy
 254 network is sufficiently trained, we get
 255 the final graph data augmentation strategy
 256 for the whole graph (all communities)
 257 by the trained policy g_{θ^*} . Finally,
 258 we reset GNN and train it with
 259 $\bigcup_{k=1}^{N_c} aug(g_{\theta^*}(G_{C_k}))$ to get the predicted labels $\hat{\mathbf{y}}_{test}$.

Algorithm 1: AutoGDA

Input : $G, \mathcal{C}, \mathbf{y}_{tr}, \mathbf{y}_{val}$
 /* Policy Optimization */
 1 **for** *episode in range*($n_episodes$) **do**
 2 Pretrain GNN and obtain $\bar{\omega}$ by Eq. (4);
 3 $\{G'_{C_1}, \dots, G'_{C_{N_c}}\} = \{G_{C_1}, \dots, G_{C_{N_c}}\}$;
 4 **for** k in $\{1, 2, \dots, N_c\}$ **do**
 5 $G'_{C_k} = aug(g_{\theta}(G_{C_k}))$; // Eq. (1)
 6 Load $\bar{\omega}$;
 7 Finetune GNN with $\bigcup_{k=1}^{N_c} G'_{C_k}$ and get $\bar{\omega}_{\theta}$;
 8 Use val. ACC to update θ ; // Eq. (5)
 9 **end**
 10 **end**
 11 $\theta^* = \theta$;
 /* Inference */
 12 $G^* = \bigcup_{k=1}^{N_c} aug(g_{\theta^*}(G_{C_k}))$;
 13 Reset and train GNN with G^* and get $\hat{\mathbf{y}}_{test}$;
Output : $\hat{\mathbf{y}}_{test}$

260 For the edges which connect nodes that belong to different communities, we make them an optional
 261 special community of edges (C_{N_c+1}) in AutoGDA. In the case we use this community, it only

Table 1: Node classification accuracy across GNN architectures and public benchmarks. Bolded are the best performance and the comparable ones (within the standard deviation of the best performance).

GNN	Method	Cora	CiteSeer	PubMed	Flickr
GCN	Original	81.5±0.4	70.3±0.5	79.0±0.3	61.2±0.4
	+AdaEdge	81.9±0.7	72.8±0.7	79.8±0.4	61.2±0.5
	+DropEdge	82.0±0.8	71.8±0.2	79.3±0.3	61.4±0.7
	+FLAG	80.2±0.3	68.1±0.5	78.5±0.2	62.3±0.4
	+GAugM	83.5±0.4	72.3±0.4	80.2±0.3	68.2±0.7
	+GAugO	83.6±0.5	73.3±1.1	79.3±0.4	62.2±0.3
	+AutoGDA	84.4±0.3	73.0±0.4	81.6±0.5	71.4±0.5
GSAGE	Original	81.3±0.5	70.6±0.5	78.3±0.6	57.4±0.5
	+AdaEdge	81.5±0.6	71.3±0.8	78.5±0.2	57.7±0.7
	+DropEdge	81.6±0.5	70.8±0.5	78.7±0.7	58.4±0.7
	+FLAG	79.2±0.9	67.9±1.4	77.4±0.3	48.5±0.6
	+GAugM	83.2±0.4	71.2±0.4	78.7±0.3	65.2±0.4
	+GAugO	82.0±0.5	72.7±0.7	79.4±0.9	56.3±0.6
	+AutoGDA	83.2±0.5	72.5±0.4	80.0±0.5	73.4±0.6
GAT	Original	83.0±0.7	72.5±0.7	79.0±0.3	46.9±1.6
	+AdaEdge	82.0±0.6	71.1±0.8	79.1±0.6	48.2±1.0
	+DropEdge	81.9±0.6	71.0±0.5	78.9±0.6	50.0±1.6
	+FLAG	79.6±0.6	67.7±0.7	78.2±0.5	48.9±1.1
	+GAugM	82.1±1.0	71.5±0.5	79.0±0.5	63.7±0.9
	+GAugO	82.2±0.8	71.6±1.1	78.5±0.8	51.9±0.5
	+AutoGDA	84.8±0.2	73.2±0.4	79.8±0.6	65.1±0.9

disjoint with others in edges (i.e., $\mathcal{V}_{CN_c+1} \cap (\cup_{k=1}^{N_c} \mathcal{V}_{Ck}) = \emptyset$ and $\mathcal{E}_{CN_c+1} = \mathcal{E} \setminus (\cup_{k=1}^{N_c} \mathcal{E}_{Ck})$), so we only apply DROPEdge and GAUG_RM on it as the other two augmentation operations are covered by other communities.

4 Experiments

In this section, we evaluate the performance of the proposed AutoGDA across different GNN backbones and datasets, and over alternative graph data augmentation methods.

4.1 Experimental Setup

Datasets We evaluate with 4 public benchmark datasets across domains: citation networks with strong homophily (Cora, CiteSeer, PubMed [7]) and social networks that exhibits heterophily [31] (Flickr [35]). We also evaluate with 2 real industry application benchmarks ECom20K and ECom43K for the task of graph anomaly detection. To validate the possibility of extending AutoGDA on heterogeneous graphs, we also evaluate with 2 heterogeneous graph benchmarks for the task of entity classification: AIFB and MUTAG [36]. Details for the datasets are provided in Appendix B.

Baselines We evaluate the AutoGDA and baselines using 3 widely used GNN architectures: GCN [7], GraphSAGE [8], and GAT [37]. We compare the node classification performance of AutoGDA with that achieved by standard GNN, as well as four state-of-the-art graph data augmentation baselines: AdaEdge [11], DropEdge [18], GAugM [13], and GAugO [13]. To evaluate the robustness of AutoGDA under graph adversarial attacks, we evaluate AutoGDA and state-of-the-art graph defending methods (GNN-Jaccard [38] and ElasticGNN [39]) against graph adversarial attacks: DICE [40] and Metattack [27]. Moreover, to evaluate the generalizability of AutoGDA on heterogeneous graph, we also test AutoGDA and baseline methods with R-GCN [41] on the heterogeneous graph benchmarks.

4.2 Experimental Results

We show comparative results of AutoGDA and baseline methods in Tables 1 and 2 and Figure 3. Table 1 is organized per GNN architecture (row), per dataset (column), and different methods (within-row). Table 2 is organized per dataset (row), per graph adversarial attack method (column), per attack level (within-column), and different methods (within-row). For customer privacy concern,

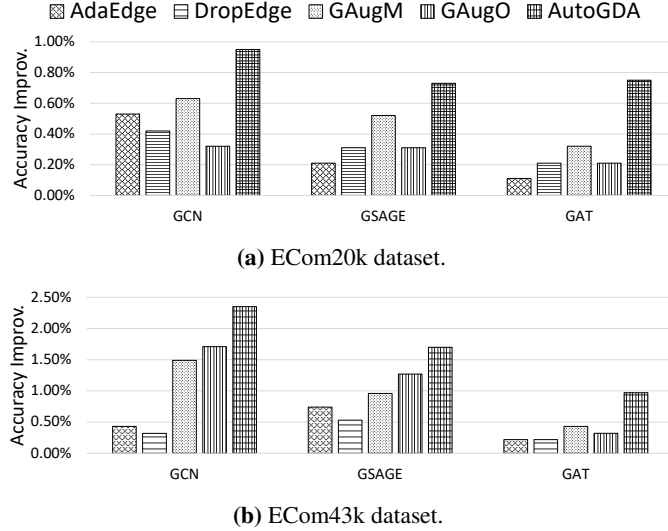


Figure 3: Relative improvements over GNNs for accuracy on two real-world industry anomaly detection datasets.

Table 2: Node classification accuracy against different levels of adversarial attacks. Bolded are the best performance and the comparable ones (within the standard deviation of the best performance).


Dataset	Attack Method Pt. Rate	DICE [40]			Metattack [27]		
		10%	30%	50%	10%	30%	50%
Cora	GCN [7]	78.4±0.6	73.6±0.9	66.8±1.2	70.2±0.9	32.6±2.0	16.6±0.8
	GAugO [13]	78.8±0.4	74.0±0.3	67.3±0.5	72.5±1.1	57.1±0.7	40.6±1.1
	GNN-Jaccard [38]	73.1±0.5	66.4±0.5	66.9±0.4	72.1±0.7	51.6±0.5	38.4±0.7
	ElasticGNN [39]	79.8±0.8	74.5±0.8	67.6±1.4	74.3±1.2	49.0±1.3	35.4±1.4
	AutoGDA	80.2±0.7	74.7±0.3	67.9±0.9	74.5±1.0	63.2±1.9	53.7±1.2
CiteSeer	GCN [7]	65.8±1.1	60.1±0.8	56.5±0.8	38.4±1.0	15.9±0.8	10.7±1.9
	GAugO [13]	67.0±0.5	60.9±0.4	56.5±0.9	50.2±0.9	34.3±0.6	29.1±1.2
	GNN-Jaccard [38]	66.5±1.3	59.8±0.7	54.1±1.0	43.7±1.2	27.0±0.7	19.8±0.4
	ElasticGNN [39]	66.7±1.4	59.7±0.8	56.3±1.4	47.5±1.3	31.8±0.3	23.5±4.3
	AutoGDA	67.8±1.1	62.1±1.3	57.6±1.1	61.5±1.6	52.6±1.3	47.8±1.7
PubMed	GCN [7]	73.9±0.3	67.1±0.3	63.6±0.7	67.2±0.4	41.3±1.5	27.5±1.7
	GAugO [13]	74.6±0.4	67.8±0.3	64.8±0.5	71.6±0.9	51.8±0.7	40.7±1.0
	GNN-Jaccard [38]	73.7±0.4	67.3±0.5	64.2±0.4	68.2±0.4	41.9±0.9	29.1±1.1
	ElasticGNN [39]	75.5±0.6	68.7±0.7	65.4±0.7	71.6±0.5	49.8±0.4	40.1±0.8
	AutoGDA	75.1±0.8	68.6±0.5	65.6±0.4	73.1±1.2	55.9±1.8	42.2±1.6

288 relative improvements instead of the performances are reported in Figure 3. We bold the best
 289 and comparable performances. In short, our proposed AutoGDA consistently achieve the best or
 290 comparable performances in all combinations of GNN architectures and datasets.

291 **Effectiveness on graph data augmentation** From Table 1 and Figure 3 we observe that our proposed
 292 AutoGDA achieves improvements over all three GNN architectures (averaged across datasets):
 293 AutoGDA improves 5.0% (GCN), 6.1% (GraphSAGE), and 7.5% (GAT), respectively. From the
 294 dataset perspective, AutoGDA also achieves improvements over all 6 datasets (averaged across GNN
 295 architectures): AutoGDA improves 2.7%, 2.2%, 2.1%, 27.6%, 0.8%, and 1.7% respectively for each
 296 dataset (Cora, CiteSeer, PubMed, Flickr, ECom20k, ECom43k). Notably, AutoGDA is especially
 297 effective on social networks (Flickr), which are naturally more heterophily. Although GAugM [13]
 298 outperformed all other baselines with large margin, AutoGDA still significantly improved from
 299 GAugM by combining the advantages of different augmentations and learning the best combined
 300 strategy. Finally, we note that AutoGDA outperforms all graph data augmentation baselines: specifi-
 301 cally, AutoGDA improves 5.6%, 5.4%, 6.1% 2.0%, and 4.7% respectively over AdaEdge, DropEdge,

Table 3: Ablation experiments using GCN on PubMed.

	PubMed
DropEdge	79.3±0.3
AutoGDA with DROPEdge (single community)	79.3±0.4
AutoGDA with DROPEdge	79.4±0.2
GAugM	80.2±0.3
AutoGDA with GAUGM (single community)	80.2±0.3
AutoGDA with GAUGM	81.2±0.4
AutoGDA (single community)	80.4±0.3
AutoGDA	81.6±0.5

**Figure 4:** AutoGDA learns diverse augmentation strategies for different communities in the PubMed graph.

302 FLAG, GAugM, and GAugO (averaged across datasets and GNNs). We reason that learning cus-
 303 tomized augmentation for each graph community and combining several state-of-the-art graph data
 304 augmentation operations both contributed to the performance improvement of AutoGDA.

305 **Robustness against graph adversarial attacks** From Table 2 we observe the proposed AutoGDA
 306 is able to effectively learn robust representation under graph adversarial attacks. Although the
 307 recently baseline ElasticGNN [39] also achieved good performance against Random Injection and
 308 DICE [40], AutoGDA outperformed ElasticGNN with large margins under Metattack [27]. In short,
 309 our proposed AutoGDA achieved the best performance for 21 out of 27 combinations of datasets and
 310 graph adversarial attack methods.

311 **Necessity of community adaptive augmentations** Table 3 shows ablation experiments on PubMed
 312 using GCN. We compare the performances of AutoGDA using only one augmentation operations
 313 versus using all augmentation operations in \mathcal{A} , and we also compare the performances of AutoGDA
 314 with single community (viewing the whole graph as the only community) versus our default setting
 315 (using Louvain method for community detection). We note that when using only one augmentation
 316 operation (combining GAUGM_ADD and GAUGM_RM as GAUGM) and under single community
 317 setting, AutoGDA performs parameter search on the existing graph data augmentation methods. We
 318 also observe from Table 3 that the default AutoGDA (with multiple graph communities) consistently
 319 outperform the single community setting, showing the community customized augmentation strategy
 320 is crucial for the performance improvement.

321 **Case Study** Figure 4 showcases the learned augmentation strategies for seven different communities
 322 on PubMed dataset by our proposed AutoGDA with GCN. We observe that our propose AutoGDA is
 323 able to learn diverse augmentations strategies for different communities in the graph.

324 5 Conclusions

325 In this paper, we studied the problem of community-customized data augmentation for node classifi-
 326 cation. Our work showed that different communities require different augmentation strategies for the
 327 best node classification performance due to their disparate characteristics. We proposed an automated
 328 graph data augmentation framework AutoGDA that adopted reinforcement learning to automatically
 329 learn the optimal augmentation strategy for each community. Through extensive experiments on
 330 benchmark graph datasets from multiple domains, our proposed approach AutoGDA achieved up to
 331 12.5% accuracy improvement over the state-of-the-art graph data augmentation baselines.

References

- 332
- 333 [1] Connor Shorten and Taghi M Khoshgoftaar. A survey on image data augmentation for deep
334 learning. *Journal of Big Data*, 6(1):1–48, 2019. 1
- 335 [2] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment:
336 Learning augmentation strategies from data. In *Proceedings of the IEEE/CVF Conference on*
337 *Computer Vision and Pattern Recognition*, pages 113–123, 2019. 1, 2, 5, 15
- 338 [3] Sungbin Lim, Ildoo Kim, Taesup Kim, Chiheon Kim, and Sungwoong Kim. Fast autoaugment.
339 *Advances in Neural Information Processing Systems*, 32:6665–6675, 2019. 15
- 340 [4] Daniel Ho, Eric Liang, Xi Chen, Ion Stoica, and Pieter Abbeel. Population based augmentation:
341 Efficient learning of augmentation policy schedules. In *International Conference on Machine*
342 *Learning*, pages 2731–2741. PMLR, 2019. 1, 15
- 343 [5] Keyu Tian, Chen Lin, Ming Sun, Luping Zhou, Junjie Yan, and Wanli Ouyang. Improving
344 auto-augment via augmentation-wise weight sharing. In *arXiv:2009.14737*, 2020. 1, 2, 5, 6, 15
- 345 [6] Chen Lin, Minghao Guo, Chuming Li, Xin Yuan, Wei Wu, Junjie Yan, Dahua Lin, and Wanli
346 Ouyang. Online hyper-parameter learning for auto-augmentation strategy. In *Proceedings of*
347 *the IEEE/CVF International Conference on Computer Vision*, pages 6579–6588, 2019. 1, 15
- 348 [7] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional
349 networks. In *arXiv:1609.02907*, 2016. 1, 3, 4, 7, 8, 15, 16
- 350 [8] Will Hamilton, Zhitaoy Ying, and Jure Leskovec. Inductive representation learning on large
351 graphs. In *Advances in Neural Information Processing Systems*, pages 1024–1034, 2017. 7, 15
- 352 [9] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. A
353 comprehensive survey on graph neural networks. *IEEE transactions on neural networks and*
354 *learning systems*, 32(1):4–24, 2020.
- 355 [10] Ziwei Zhang, Peng Cui, and Wenwu Zhu. Deep learning on graphs: A survey. *IEEE Transactions*
356 *on Knowledge and Data Engineering*, 2020. 1
- 357 [11] Deli Chen, Yankai Lin, Wei Li, Peng Li, Jie Zhou, and Xu Sun. Measuring and relieving the
358 over-smoothing problem for graph neural networks from the topological view. In *Proceedings*
359 *of the AAAI Conference on Artificial Intelligence*, volume 34, pages 3438–3445, 2020. 1, 4, 7,
360 15, 17, 18
- 361 [12] Kezhi Kong, Guohao Li, Mucong Ding, Zuxuan Wu, Chen Zhu, Bernard Ghanem, Gavin
362 Taylor, and Tom Goldstein. Flag: Adversarial data augmentation for graph neural networks. In
363 *arXiv:2010.09891*, 2020. 15, 18
- 364 [13] Tong Zhao, Yozen Liu, Leonardo Neves, Oliver Woodford, Meng Jiang, and Neil Shah. Data
365 augmentation for graph neural networks. In *Proceedings of the AAAI Conference on Artificial*
366 *Intelligence*, volume 35, pages 11015–11023, 2021. 1, 2, 3, 4, 7, 8, 15, 16, 17, 18
- 367 [14] Hyeonjin Park, Seunghun Lee, Sihyeon Kim, Jinyoung Park, Jisu Jeong, Kyung-Min Kim,
368 Jung-Woo Ha, and Hyunwoo J Kim. Metropolis-hastings data augmentation for graph neural
369 networks. *Advances in Neural Information Processing Systems*, 34:19010–19020, 2021.
- 370 [15] Zhengzheng Tang, Ziyue Qiao, Xuehai Hong, Yang Wang, Fayaz Ali Dharejo, Yuanchun
371 Zhou, and Yi Du. Data augmentation for graph convolutional network on semi-supervised
372 classification. In *arXiv:2106.08848*, 2021. 15
- 373 [16] Tong Zhao, Bo Ni, Wenhao Yu, Zhichun Guo, Neil Shah, and Meng Jiang. Action sequence
374 augmentation for early graph-based anomaly detection. In *Proceedings of the 30th ACM*
375 *International Conference on Information & Knowledge Management*, pages 2668–2678, 2021.
376 15
- 377 [17] Tong Zhao, Gang Liu, Stephan Günnemann, and Meng Jiang. Graph data augmentation for
378 graph machine learning: A survey. *arXiv preprint arXiv:2202.08871*, 2022. 1, 15
- 379 [18] Yu Rong, Wenbing Huang, Tingyang Xu, and Junzhou Huang. Dropedge: Towards deep
380 graph convolutional networks on node classification. In *International Conference on Learning*
381 *Representations*, 2019. 1, 2, 3, 7, 15, 17, 18
- 382 [19] Michelle Girvan and Mark EJ Newman. Community structure in social and biological networks.
383 *Proceedings of the national academy of sciences*, 99:7821–7826, 2002. 2, 4

- 384 [20] Mark EJ Newman and Michelle Girvan. Finding and evaluating community structure in
385 networks. *Physical review E*, 69(2):026113, 2004. 2, 4
- 386 [21] Wei-Lin Chiang, Xuanqing Liu, Si Si, Yang Li, Samy Bengio, and Cho-Jui Hsieh. Cluster-gcn:
387 An efficient algorithm for training deep and large graph convolutional networks. In *Proceedings*
388 *of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*,
389 pages 257–266, 2019. 2, 4
- 390 [22] Yong-Yeol Ahn, James P Bagrow, and Sune Lehmann. Link communities reveal multiscale
391 complexity in networks. *nature*, 466(7307):761–764, 2010. 2, 4
- 392 [23] Mohammad Hamdaqa, Ladan Tahvildari, Neil LaChapelle, and Brian Campbell. Cultural scene
393 detection using reverse louvain optimization. *Science of Computer Programming*, 95:44–72,
394 2014.
- 395 [24] Filippo Radicchi, Claudio Castellano, Federico Cecconi, Vittorio Loreto, and Domenico Parisi.
396 Defining and identifying communities in networks. *Proceedings of the national academy of*
397 *sciences*, 101(9):2658–2663, 2004. 2, 4
- 398 [25] Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. Fast
399 unfolding of communities in large networks. *Journal of statistical mechanics: theory and*
400 *experiment*, 2008(10):P10008, 2008. 2, 4, 17, 19
- 401 [26] Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, and Yang Shen.
402 Graph contrastive learning with augmentations. *Advances in Neural Information Processing*
403 *Systems*, 33:5812–5823, 2020. 2, 3, 15
- 404 [27] Daniel Zügner and Stephan Günnemann. Adversarial attacks on graph neural networks via meta
405 learning. In *International Conference on Learning Representations*, 2019. 2, 7, 8, 9
- 406 [28] Thomas N Kipf and Max Welling. Variational graph auto-encoders. In *arXiv:1611.07308*, 2016.
407 3, 15
- 408 [29] Fragkiskos D Malliaros and Michalis Vazirgiannis. Clustering and community detection in
409 directed networks: A survey. *Physics reports*, 533(4):95–142, 2013. 4
- 410 [30] Namkyeong Lee, Junseok Lee, and Chanyoung Park. Augmentation-free self-supervised
411 learning on graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36,
412 pages 7372–7380, 2022. 4
- 413 [31] Jiong Zhu, Yujun Yan, Lingxiao Zhao, Mark Heimann, Leman Akoglu, and Danai Koutra. Be-
414 yond homophily in graph neural networks: Current limitations and effective designs. *Advances*
415 *in Neural Information Processing Systems*, 33:7793–7804, 2020. 4, 7
- 416 [32] Wenzheng Feng, Jie Zhang, Yuxiao Dong, Yu Han, Huanbo Luan, Qian Xu, Qiang Yang,
417 Evgeny Kharlamov, and Jie Tang. Graph random neural network for semi-supervised learning
418 on graphs. In *arXiv:2005.11079*, 2020. 4, 15
- 419 [33] Benoît Colson, Patrice Marcotte, and Gilles Savard. An overview of bilevel optimization.
420 *Annals of operations research*, 153(1):235–256, 2007. 5
- 421 [34] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal
422 policy optimization algorithms. In *arXiv:1707.06347*, 2017. 6, 17
- 423 [35] Xiao Huang, Jundong Li, and Xia Hu. Label informed attributed network embedding. In
424 *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, pages
425 731–739, 2017. 7, 16
- 426 [36] Petar Ristoski, Gerben Klaas Dirk de Vries, and Heiko Paulheim. A collection of benchmark
427 datasets for systematic evaluations of machine learning on the semantic web. In *International*
428 *Semantic Web Conference*, pages 186–194, 2016. 7, 16
- 429 [37] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua
430 Bengio. Graph attention networks. In *arXiv:1710.10903*, 2017. 7, 15, 16
- 431 [38] Huijun Wu, Chen Wang, Yuriy Tyshetskiy, Andrew Docherty, Kai Lu, and Liming Zhu. Ad-
432 versarial examples on graph data: Deep insights into attack and defense. In *arXiv:1903.01610*,
433 2019. 7, 8
- 434 [39] Xiaorui Liu, Wei Jin, Yao Ma, Yaxin Li, Hua Liu, Yiqi Wang, Ming Yan, and Jiliang Tang.
435 Elastic graph neural networks. In *International Conference on Machine Learning*, pages
436 6837–6849, 2021. 7, 8, 9

- 437 [40] Marcin Waniek, Tomasz P Michalak, Michael J Wooldridge, and Talal Rahwan. Hiding
438 individuals and communities in a social network. volume 2, pages 139–147, 2018. 7, 8, 9
- 439 [41] Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max
440 Welling. Modeling relational data with graph convolutional networks. In *European semantic*
441 *web conference*, pages 593–607. Springer, 2018. 7, 16, 18
- 442 [42] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. Deepwalk: Online learning of social repre-
443 sentations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge*
444 *discovery and data mining*, pages 701–710, 2014. 15
- 445 [43] Daixin Wang, Peng Cui, and Wenwu Zhu. Structural deep network embedding. In *Proceedings*
446 *of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*,
447 pages 1225–1234, 2016.
- 448 [44] Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. Line: Large-
449 scale information network embedding. In *Proceedings of the 24th international conference on*
450 *world wide web*, pages 1067–1077, 2015. 15
- 451 [45] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks
452 on graphs with fast localized spectral filtering. In *Advances in neural information processing*
453 *systems*, pages 3844–3852, 2016. 15
- 454 [46] Mikael Henaff, Joan Bruna, and Yann LeCun. Deep convolutional networks on graph-structured
455 data. In *arXiv:1506.05163*, 2015.
- 456 [47] Ruoyu Li, Sheng Wang, Feiyun Zhu, and Junzhou Huang. Adaptive graph convolutional neural
457 networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
458 15
- 459 [48] Ron Levie, Federico Monti, Xavier Bresson, and Michael M Bronstein. Cayleynets: Graph
460 convolutional neural networks with complex rational spectral filters. *IEEE Transactions on*
461 *Signal Processing*, 67(1):97–109, 2018.
- 462 [49] Yao Ma, Xiaorui Liu, Tong Zhao, Yozen Liu, Jiliang Tang, and Neil Shah. A unified view on
463 graph neural networks as graph signal denoising. In *arXiv:2010.01777*, 2020. 15
- 464 [50] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. Spectral networks and locally
465 connected networks on graphs. In *arXiv:1312.6203*, 2013. 15
- 466 [51] Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodola, Jan Svoboda, and
467 Michael M Bronstein. Geometric deep learning on graphs and manifolds using mixture model
468 cnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*,
469 pages 5115–5124, 2017. 15
- 470 [52] Hongyang Gao, Zhengyang Wang, and Shuiwang Ji. Large-scale learnable graph convolutional
471 networks. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge*
472 *Discovery & Data Mining*, pages 1416–1424, 2018.
- 473 [53] Mathias Niepert, Mohamed Ahmed, and Konstantin Kutzkov. Learning convolutional neural
474 networks for graphs. In *International conference on machine learning*, pages 2014–2023, 2016.
475 15
- 476 [54] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure
477 Leskovec. Graph convolutional neural networks for web-scale recommender systems. In
478 *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery &*
479 *Data Mining*, pages 974–983, 2018. 15
- 480 [55] Keyulu Xu, Chengtao Li, Yonglong Tian, Tomohiro Sonobe, Ken-ichi Kawarabayashi, and
481 Stefanie Jegelka. Representation learning on graphs with jumping knowledge networks. In
482 *arXiv:1806.03536*, 2018. 15
- 483 [56] Guohao Li, Matthias Muller, Ali Thabet, and Bernard Ghanem. Deepgcns: Can gcns go as
484 deep as cnns? In *Proceedings of the IEEE International Conference on Computer Vision*, pages
485 9267–9276, 2019. 15
- 486 [57] Vikas Verma, Meng Qu, Alex Lamb, Yoshua Bengio, Juho Kannala, and Jian Tang. Graph-
487 mix: Regularized training of graph neural networks for semi-supervised learning. In
488 *arXiv:1909.11715*, 2019. 15

- 489 [58] Yang Gao, Hong Yang, Peng Zhang, Chuan Zhou, and Yue Hu. Graph neural architecture
490 search. In *IJCAI*, volume 20, pages 1403–1409, 2020. 15
- 491 [59] Kaixiong Zhou, Qingquan Song, Xiao Huang, and Xia Hu. Auto-gnn: Neural architecture
492 search of graph neural networks. In *arXiv:1909.03184*, 2019. 15
- 493 [60] Kaize Ding, Zhe Xu, Hanghang Tong, and Huan Liu. Data augmentation for deep graph
494 learning: A survey. *arXiv preprint arXiv:2202.08235*, 2022. 15
- 495 [61] Gang Liu, Tong Zhao, Jiabin Xu, Tengfei Luo, and Meng Jiang. Graph rationalization with
496 environment-based augmentations. In *Proceedings of the 28th ACM SIGKDD International
497 Conference on Knowledge Discovery & Data Mining*, 2022.
- 498 [62] Xiaotian Han, Zhimeng Jiang, Ninghao Liu, and Xia Hu. G-mixup: Graph data augmentation
499 for graph classification. *International Conference on Machine Learning*, 2022. 15
- 500 [63] Yiwei Wang, Yujun Cai, Yuxuan Liang, Henghui Ding, Changhu Wang, Siddharth Bhatia, and
501 Bryan Hooi. Adaptive data augmentation on temporal graphs. *Advances in Neural Information
502 Processing Systems*, 34, 2021. 15
- 503 [64] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov.
504 Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine
505 learning research*, 15(1):1929–1958, 2014. 15
- 506 [65] Zhan Gao, Subhrajit Bhattacharya, Leiming Zhang, Rick S Blum, Alejandro Ribeiro, and
507 Brian M Sadler. Training robust graph neural networks with topology adaptive edge dropping.
508 In *arXiv:2106.02892*, 2021. 15
- 509 [66] Dongsheng Luo, Wei Cheng, Wenchao Yu, Bo Zong, Jingchao Ni, Haifeng Chen, and Xiang
510 Zhang. Learning to drop: Robust graph neural network via topological denoising. In *Proceedings
511 of the 14th ACM International Conference on Web Search and Data Mining*, pages 779–787,
512 2021.
- 513 [67] Cheng Zheng, Bo Zong, Wei Cheng, Dongjin Song, Jingchao Ni, Wenchao Yu, Haifeng Chen,
514 and Wei Wang. Robust graph representation learning via neural sparsification. In *International
515 Conference on Machine Learning*, pages 11458–11468. PMLR, 2020. 15
- 516 [68] Thomas Kipf, Ethan Fetaya, Kuan-Chieh Wang, Max Welling, and Richard Zemel. Neural
517 relational inference for interacting systems. In *International Conference on Machine Learning*,
518 pages 2688–2697. PMLR, 2018. 15
- 519 [69] Aditya Grover, Aaron Zweig, and Stefano Ermon. Graphite: Iterative generative modeling of
520 graphs. In *International conference on machine learning*, pages 2434–2444. PMLR, 2019.
- 521 [70] Luca Franceschi, Mathias Niepert, Massimiliano Pontil, and Xiao He. Learning discrete
522 structures for graph neural networks. In *International conference on machine learning*, pages
523 1972–1982. PMLR, 2019. 15
- 524 [71] Yingxue Zhang, Soumyasundar Pal, Mark Coates, and Deniz Ustebay. Bayesian graph con-
525 volutional neural networks for semi-supervised classification. In *Proceedings of the AAAI
526 Conference on Artificial Intelligence*, volume 33, pages 5829–5836, 2019. 15
- 527 [72] Yanqiao Zhu, Weizhi Xu, Jinghao Zhang, Yuanqi Du, Jieyu Zhang, Qiang Liu, Carl Yang, and
528 Shu Wu. A survey on graph structure learning: Progress and opportunities. *arXiv e-prints*,
529 pages arXiv–2103, 2021. 15
- 530 [73] Yu Chen, Lingfei Wu, and Mohammed Zaki. Iterative deep graph learning for graph neural
531 networks: Better and robust node embeddings. *Advances in neural information processing
532 systems*, 33:19314–19326, 2020. 15
- 533 [74] Yiwei Wang, Wei Wang, Yuxuan Liang, Yujun Cai, Juncheng Liu, and Bryan Hooi. Nodeaug:
534 Semi-supervised node classification with data augmentation. In *Proceedings of the 26th ACM
535 SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 207–217,
536 2020. 15
- 537 [75] Tong Zhao, Gang Liu, Daheng Wang, Wenhao Yu, and Meng Jiang. Learning from counter-
538 factual links for link prediction. In *International Conference on Machine Learning*, pages
539 26911–26926. PMLR, 2022. 15
- 540 [76] Xinyu Zhang, Qiang Wang, Jian Zhang, and Zhao Zhong. Adversarial autoaugment. In
541 *arXiv:1912.11188*, 2019. 15

- 542 [77] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical
543 automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF*
544 *Conference on Computer Vision and Pattern Recognition Workshops*, pages 702–703, 2020. 15
- 545 [78] Wei Jin, Xiaorui Liu, Xiangyu Zhao, Yao Ma, Neil Shah, and Jiliang Tang. Automated
546 self-supervised learning for graphs. In *arXiv:2106.05470*, 2021. 15
- 547 [79] Yuning You, Tianlong Chen, Yang Shen, and Zhangyang Wang. Graph contrastive learning
548 automated. In *arXiv:2106.07594*, 2021.
- 549 [80] Susheel Suresh, Pan Li, Cong Hao, and Jennifer Neville. Adversarial graph augmentation to
550 improve graph contrastive learning. *Advances in Neural Information Processing Systems*, 34,
551 2021. 15
- 552 [81] Youzhi Luo, Michael McThrow, Wing Yee Au, Tao Komikado, Kanji Uchino, Koji Maruhash,
553 and Shuiwang Ji. Automated data augmentations for graph classification. *arXiv preprint*
554 *arXiv:2202.13248*, 2022. 15
- 555 [82] Leo A Goodman. Snowball sampling. *The annals of mathematical statistics*, pages 148–170,
556 1961. 16
- 557 [83] James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *Journal*
558 *of machine learning research*, 13(2), 2012. 18
- 559 [84] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna:
560 A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM*
561 *SIGKDD international conference on knowledge discovery & data mining*, pages 2623–2631,
562 2019. 18
- 563 [85] Amine Abou-Rjeili and George Karypis. Multilevel algorithms for partitioning power-law
564 graphs. In *Proceedings 20th IEEE International Parallel & Distributed Processing Symposium*,
565 pages 10–pp. IEEE, 2006. 19

566 A Related Work

567 **Graph Neural Networks** GNNs enjoy widespread use in modern graph-based machine learning
 568 due to their flexibility to incorporate node features, custom aggregations and inductive operation,
 569 unlike earlier works which were based on embedding lookups [42–44]. In recent years, many spectral
 570 GNN variants [7, 45–49] were proposed following the initial idea of convolution based on spectral
 571 graph theory [50]. As spectral GNNs usually requires (expensive) operations on the full adjacency
 572 matrix, spatial GNNs which perform graph convolution with neighborhood aggregation became
 573 prominent [8, 37, 51–53] owing to their scalability and flexibility [54]. Several other works also
 574 proposed advanced architectures which add residual connections to facilitate deep GNN training
 575 [55, 56]. GraphMix [57] proposed to regularize the GNN model with a fully connected network.
 576 More recently, graph neural architecture search (NAS) methods [58, 59] utilizing reinforcement
 577 learning were proposed to learn the optimal GNN architecture.

578 **Graph Data Augmentation** As GNNs have emerged as a rising approach for learning with graph data,
 579 Graph Data Augmentation (GDA) [17, 60–62] for GNNs were proposed and studied in recent years.
 580 Due to the complex, non-Euclidean structure of graphs, most GDA work focused on manipulating
 581 the graph structure [13, 63, 63]. DropEdge [18] randomly drops a fraction of edges during each
 582 training epoch, in a way similar to Dropout [64]. Following DropEdge, several works [65–67]
 583 proposed methods that learns to drop instead of dropping at random. Graph structure learning
 584 methods [47, 68–70] can also be viewed as graph data augmentation as they learn from graphs whose
 585 structures are partially or totally unknown. AdaEdge [11] and BGCN [71] are iterative methods
 586 that updates the graph structure with the prediction of GNNs. Zhao et al. [13] showed that graph
 587 homophily critically affects message passing-based GNNs and proposed GAugM and GAugO that
 588 manipulate the graph structure with the edge probabilities given by VGAE [28] to augment the
 589 graph data. Kong et al. [12] and Tang et al. [15] proposed augmentation methods that operates on
 590 the node features. Graph structural learning methods [70, 72, 73] search for better graph structure
 591 that augments the initial graph structure, with the goal of optimizing the graph for downstream
 592 tasks. Aside from the methods that directly use GDA for semi-supervised node classification, several
 593 works that used GDA in self-supervised graph learning were also proposed. NodeAug [74] and
 594 Grand [32] used augmentation with self-supervised consistency loss as an additional term to the
 595 cross entropy loss. GraphCL [26] used augmentation in self-supervised graph contrastive learning for
 596 graph classification. Eland [16] proposed action sequence augmentation for graph anomaly detection.
 597 Zhao et al. [75] studied counterfactual data augmentation for link prediction.

598 **Automated Data Augmentation** Several automated data augmentation approaches [2–4, 6, 76]
 599 have been proposed in CV in the past few years. These methods seek to find the optimal data
 600 augmentation policies for each given dataset automatically. Cubuk et al. [2] formulated the automated
 601 data augmentation problem as a discrete search problem and proposed a reinforcement learning
 602 framework to search the best augmentation operations via proxy tasks (i.e., a smaller model). Several
 603 works [3, 4, 6] were then proposed to improve the efficiency of automated data augmentation. Cubuk
 604 et al. [77] showed that the models with different number of parameters benefits from different
 605 magnitude of augmentation operations and proposed RandAug that searches for one augmentation
 606 magnitude that is used for all operations. Tian et al. [5] proposed the Augmentation-Wise Weight
 607 Sharing method that enables a fast evaluation process on the original model while not sacrificing
 608 the efficiency. Recently, several automated augmentation methods [78–80] have been proposed for
 609 self-supervised graph representation learning. Luo et al. [81] studied automated data augmentation for
 610 the task of graph classification. Nevertheless, automated data augmentation is rather under-explored
 611 for (semi-)supervised node classification tasks.

612 B Additional Dataset Details

613 In this section, we provide some additional, relevant dataset details. The statistics of all datasets are
 614 provided in Tables 4 and 5.

615 **Citation networks.** Cora, CiteSeer and PubMed are citation networks that are commonly used as
 616 benchmarks in GNN-related prior works [7, 11, 13, 18, 37]. In these citation networks, the nodes are
 617 published papers; the features are preprocessed (e.g., bag-of-words) vectors of the corresponding

Table 4: Summary statistics and experimental setup for the homogeneous datasets.

	Cora	CiteSeer	PubMed	Flickr	ECom20k	ECom43k
# Nodes	2,708	3,327	19,717	7,575	20,799	43,117
# Edges	5,278	4,552	44,338	239,738	47,661	117,469
# Features	1,433	3,703	500	12,047	132	132
# Classes	7	6	3	9	2	2
# Training nodes	140	120	60	757	2,275	4,209
# Validation nodes	500	500	500	1,515	2,275	4,209
# Test nodes	1,000	1,000	1,000	5,303	6,825	12,627

Table 5: Summary statistics and experimental setup for the heterogeneous datasets.

	AIFB	MUTAG
# Nodes	8,285	23,644
# Relations	45	23
# Edges	29,043	74,227
# Classes	4	2
# Training nodes	218	112
# Validation nodes	54	28
# Test nodes	68	36

618 paper title and/or abstract; the edges represent the citation relation between papers; the labels are the
619 category of each paper.

620 **Social networks.** Flickr is an online social network platform, where users can also follow each
621 other as well as posting images and videos. The user-specified list of interest tags are used as user
622 features and the groups that users joined are used as labels [35].

623 **E-commerce networks.** ECom20k and ECom43k are e-commerce networks that were constructed
624 with customer purchase/review records from a leading international e-commerce website. The
625 network contains four types of nodes: customers, sellers, products, and reviews, in which customer
626 purchases products from sellers and leave reviews to products. They are two graphs constructed with
627 records in different time periods. The task for these two datasets is abusive customer detection and
628 the customers with golden labels are split into train/validation/test sets for node classification. The
629 node features contains original node attributes as well as the one-hot encoded vectors of the node
630 types. As the golden labels are limited and severely biased, the datasets are sampled with snowball
631 sampling [82] with labeled abusive customers to ensure the relative independence of the graph. The
632 node attributes used is anonymized and do not contain any personally identifiable information.

633 **Heterogeneous networks.** AIFB and MUTAG are commonly used Resource Description Frame-
634 work (RDF) formatted datasets [36, 41]. Each graph contains multiple types of entities (nodes) and
635 multiple types of relations (edges). In each dataset, the task is to classify the target properties of one
636 group of entities. The datasets do not contain raw node features.

637 **Validation Method.** For Cora, Citeseer, PubMed, and Flickr, we follow the commonly used semi-
638 supervised setting in most GNN literature [7, 13, 37]. For ECom20k and ECom43k, we use 20/20/60%
639 for train/validation/test splitting. For AIFB and MUTAG, we use the official splitting provided in the
640 DGL package².

641 C Implementation Details

642 Our code package can be found at <https://tinyurl.com/autogda> for anonymous reviewing, and
643 it will be made publicly available.

²<https://www.dgl.ai/>

Table 6: Node classification accuracy and total hyperparameter searching time on PubMed dataset.

	Accuracy	Search Time (mins)
AdaEdge	79.8±0.4	283.5
DropEdge	79.3±0.3	48.1
FLAG	78.5±0.2	23.6
GAugM	80.2±0.3	53.2
GAugO	79.3±0.4	149.4
Different hyperparameter searching method for community-specific augmentations		
Random Search	81.8±0.6	196.2
Bayesian Search	81.3±0.5	126.5
AutoGDA	81.6±0.5	74.6

644 All the experiments in this work were conducted on either an AWS EC2 P4 Instance³ or a G4dn
645 Instance⁴. The P4 instance is equipped with 48 Intel Cascade Lake processor cores (96 vCPUs), 1.1
646 TB of RAM, and 8 Nvidia A100 GPU cards (40 GB of RAM each). The G4dn instance is equipped
647 with 48 Intel Cascade Lake vCPUs, 192 GB of RAM, and 4 Nvidia T4 GPU cards (16 GB of RAM
648 each). To ensure fair comparison for the search time experiments, all the experiments in Table 6
649 are conducted on the same G4dn instance. Our method are implemented with Python 3.8.5 with
650 PyTorch. A list of used packages is included in `requirements.txt` in the code package. The
651 commands for reproducing our results can be found in `README.md` in the code package. Note that
652 although the EC2 instances are equipped with multiple GPU cards, AutoGDA only need one GPU to
653 run all the experiments.

654 We report test accuracy averaged over 20 runs along with respective standard deviations. For baseline
655 methods with same datasets in their original paper, we directly use their reported performances –
656 numbers reported in the original paper are more preferred than the reproduced results in other papers.

657 C.1 Baseline methods

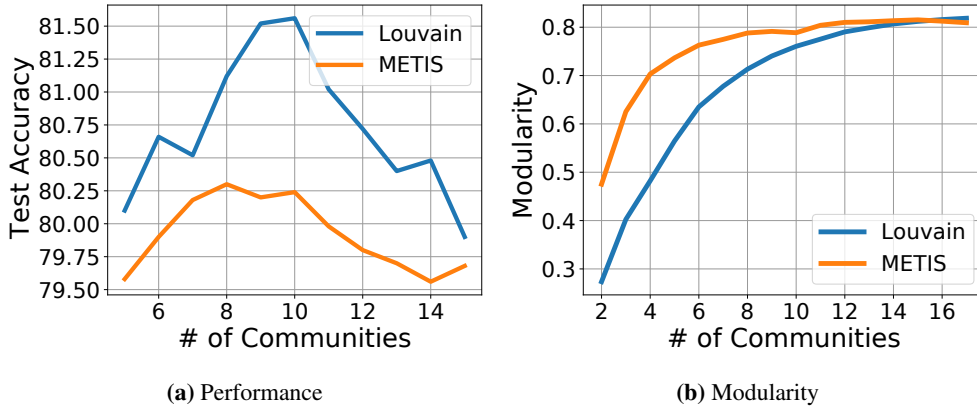
658 For the new reproduced results in this work, all original GNN architectures are implemented in DGL⁵
659 with Adam optimizer. For a fair comparison, we use hidden size of 128 for all methods. For baseline
660 methods, we implemented AdaEdge [11] and DropEdge [18] with PyTorch and DGL, and used the
661 official code packages⁶ from the authors for GAugM and GAugO [13]. The hyperparameters for all
662 baseline methods are tuned according to with the same range as in the proposed AutoGDA, and the
663 hyperparameters for GAugM and GAugO [13] are tuned with the script by their authors.

664 C.2 AutoGDA variants

665 As the same as for the baselines, all GNNs in AutoGDA are implemented with DGL and we used
666 hidden size of 128. We use the public PPO [34] implementation in the `stable-baselines3`
667 package⁷ and implement our RL environment with `gym`⁸. All parameters for PPO algorithm are set
668 as default from `stable-baselines3`. We use the Louvain method [25] as the default community
669 detection method for all experiments in Section 4 except the ones in the sensitivity analysis (Figure 5).
670 The number of communities is treated as a hyperparameter and determined with the help of modularity
671 measurement as described in the sensitivity analysis in Section 4.

Table 7: Node classification accuracy of different graph data augmentation methods on heterogeneous graphs.

	AIFB	MUTAG
RGCN	93.1±1.4	68.4±3.1
RGCN+AdaEdge	93.9±1.1	69.4±3.4
RGCN+DropEdge	94.1±2.7	70.6±2.1
RGCN+AutoGDA	95.6±2.2	72.2±1.8

**Figure 5:** AutoGDA is robust to the choices of community detection algorithms as well as the number of communities. The number of communities can be decided with the modularity measurement.

672 D Additional Experimental Results

673 **Efficiency of the proposed AutoGDA .** Table 6 shows the node classification performance and
674 total runtime (including hyperparameter searching) of baseline graph data augmentation methods, our
675 proposed AutoGDA, and community-specific augmentations with other searching methods. All base-
676 line methods require additional hyperparameter search on the magnitudes of augmentation operations,
677 while the policy network in AutoGDA is able to search for community customized augmentations
678 strategies automatically. Therefore, our proposed AutoGDA is capable of out-performing baseline
679 augmentation methods without requiring much additional runtime. Moreover, we also evaluate when
680 substituting the policy network in AutoGDA with other hyperparameter searching methods: Random
681 Search [83] and Bayesian Search (implemented with Optuna [84]). We observe that although they
682 can achieve similar performance as AutoGDA, they require much more time than AutoGDA to find
683 the optimal augmentation strategy due to the huge searching space.

684 **Extension of AutoGDA on heterogeneous graphs.** Table 7 shows the entity classification perfor-
685 mances of AdaEdge [11], DropEdge [18], and AutoGDA using RGCN [41] as backbone. We note
686 that more advanced graph data augmentation baselines FLAG [12], GAUGM, and GAUGO [13] are not
687 compatible with heterogeneous graphs. Moreover, as the datasets do not have raw node features, we
688 only search for the DROPEdge operation in AutoGDA on these datasets (i.e., $\mathcal{A} = \{\text{DROPEdge}\}$).
689 We observe that our proposed AutoGDA achieves the best performance on both datasets, showing
690 that AutoGDA can be easily generalized to heterogeneous graphs.

691 **Sensitivity of AutoGDA .** Figure 5a shows the sensitivity analysis of our proposed AutoGDA on
692 the choices of community detection methods and the number of communities. Figure 5b shows the

³<https://aws.amazon.com/ec2/instance-types/p4/>

⁴<https://aws.amazon.com/ec2/instance-types/g4/>

⁵<https://www.dgl.ai/>

⁶<https://github.com/zhao-tong/GAug>

⁷<https://github.com/DLR-RM/stable-baselines3>

⁸<https://github.com/openai/gym>

693 modularity measurement for the two community detection methods (Louvain [25] and METIS [85])
694 at different number of communities. We observe that AutoGDA is generally good when $8 \leq N_c \leq 10$
695 for the PubMed dataset, which is also where the modularity curve converge. Thus, a good number of
696 communities for AutoGDA is generally easy to find with the help of the modularity measurement.