

---

# GC-Bench: An Open and Unified Benchmark for Graph Condensation

---

## Appendix

1	<b>Contents</b>	
2	<b>A Details of GC-Bench</b>	<b>2</b>
3	A.1 Datasets . . . . .	2
4	A.2 Algorithms . . . . .	2
5	A.3 Hyper-Parameter Setting . . . . .	4
6	A.4 Computation resources . . . . .	4
7	A.5 Discussion on Existing Benchmarks . . . . .	5
8	<b>B Settings and Additional Results</b>	<b>5</b>
9	B.1 Settings and Additional Results of Performance Comparison (RQ1) . . . . .	5
10	B.1.1 Comparison Setting . . . . .	5
11	B.1.2 Additional Results . . . . .	6
12	B.2 Settings and Additional Results of Structure in Graph Condensation (RQ2) . . . . .	7
13	B.2.1 Experimental Settings . . . . .	7
14	B.2.2 Additional Results . . . . .	8
15	B.3 Settings and Additional Results of Transferability in Different Tasks (RQ3) . . . . .	9
16	B.3.1 Link Prediction . . . . .	9
17	B.3.2 Node Clustering . . . . .	10
18	B.3.3 Anomaly Detection . . . . .	11
19	B.4 Settings and Additional Results of Transferability across Backbone Model Architectures (RQ4) . . . . .	11
20	B.4.1 Experimental Settings . . . . .	11
21	B.4.2 Additional Results . . . . .	14
22	B.5 Settings and Additional Results of Initialization Impacts (RQ5) . . . . .	14
23	B.5.1 Experimental Settings . . . . .	14
24	B.5.2 Additional Results . . . . .	15
25	B.6 Settings and Additional Results of Efficiency and Scalability (RQ6) . . . . .	15
26	B.6.1 Experimental Settings . . . . .	15
27	B.6.2 Additional Results . . . . .	15
28		
29	<b>C Reproducibility and Limitations</b>	<b>15</b>

## 30 A Details of GC-Bench

### 31 A.1 Datasets

32 The evaluation node-level datasets include 5 homogeneous datasets (3 transductive datasets, i.e.,  
33 *Cora*, *Citeseer* [14] and *ogbn-arxiv* [11], and 2 inductive datasets, i.e., *Flickr* [38] and *Reddit* [10])  
34 and 2 heterogeneous datasets (*ACM* [41] and *DBLP* [6]). The evaluation graph-level datasets include  
35 5 datasets (*NCII* [29], *DD* [3], *ogbg-molbace* [11], *ogbg-molhiv* [11], *ogbg-molbbbp* [11]).

36 We utilize the standard data splits provided by PyTorch Geometric [5] and the Open Graph Benchmark  
37 (OGB) [11] for our experiments. For datasets in TUDataset [25], we split the data into 10% for  
38 testing, 10% for validation, and 80% for training. For *ACM* and *DBLP* datasets, we follow the settings  
39 outlined in [23]. Dataset statistics are shown in Table A1.

Table A1: Dataset statistics. For heterogeneous datasets, the features are from the target nodes (papers in *ACM* and authors in *DBLP*).

	Dataset	#Nodes / #Avg. Nodes	#Edges / #Avg. Edges	#Classes	#Features / #Graphs
Node-level	<i>Cora</i>	2,708	5,429	7	1,433
	<i>Citeseer</i>	3,327	4,732	6	3,703
	<i>ogbn-arxiv</i>	169,343	1,166,243	40	128
	<i>Flickr</i>	89,250	899,756	7	500
	<i>Reddit</i>	232,965	57,307,946	210	602
	<i>ACM</i>	10,942	547,872	3	1,902
	<i>DBLP</i>	37,791	170,794	4	334
Graph-level	<i>ogbg-molhiv</i>	25.5	54.9	2	41,127
	<i>ogbg-molbace</i>	34.1	36.9	2	1,513
	<i>ogbg-molbbbp</i>	24.1	26.0	2	2,039
	<i>NCII</i>	29.8	32.3	2	4,110
	<i>DD</i>	284.3	715.7	2	1,178

### 40 A.2 Algorithms

41 We summarize the current GC algorithms in Table A2. We choose **12** representative ones for  
42 evaluation in this paper including *Random*, *K-Center* [27], *Herding* [33], *GCond* [13], *DosCond* [12],  
43 *SGDD* [37], *GCDM* [18], *DM* [22], *SFGC* [42], *GEOM* [40], *KiDD* [36], *Mirage* [9]. **We will**  
44 **continue to update and improve the benchmark to include more algorithms.** Here we introduce  
45 the GC algorithms in detail:

#### 46 • Traditional Core-set Methods

- 47 – **Random**: For node classification tasks, nodes are randomly selected to form a new subgraph.  
48 For graph classification, the graphs are randomly selected to create a new subset.
- 49 – **Herding** [33]: The nodes or graphs are selected samples that are closest to the cluster center.
- 50 – **K-Center** [27]: Nodes or graphs are chosen such that they have the minimal distance to the  
51 nearest cluster center, which is generated using the K-Means Clustering method.

#### 52 • Gradient Matching Methods

- 53 – **GCond** [13]: In GCond, the optimization of the synthetic dataset is framed as a bi-level  
54 problem. It adapts a gradient matching scheme to match the gradients of GNN parameters  
55 between the condensed and original graphs, while optimizing the model’s performance on  
56 the datasets. For generating the synthetic adjacency matrix, GCond employs a Multi-Layer  
57 Perceptron (MLP) to model the edges by using node features as input, maintaining the  
58 correlations between node features and graph structures.
- 59 – **DosCond** [12]: In DosCond, the gradient matching scheme only matches the network  
60 gradients for model initialization  $\theta_0$  while discarding the training trajectory of  $\theta_t$ , which  
61 accelerated the entire condensation process by only informing the direction to update the

Table A2: Summary of Graph Condensation (GC) algorithms. We also provide public access to the official algorithm implementations. “KRR” is short for Kernel Ridge Regression and “CTC” is short for computation tree compression. “GNN” is short for Graph, “GNTK” is short for graph neural tangent kernel, “SD” is short for spectral decomposition. “NC” is short for node classification, “LP” is short for link prediction, “AD” is short for anomaly detection, and “GC” is short for graph classification.

Taxonomy	Method	Initialization	Backbone Model	Downstream Task	Code	Venue
Traditional Methods	Random	—	—	—	—	—
	Herding [33]	—	—	—	<a href="#">link</a>	ICML, 2009
	K-Center [27]	—	—	—	<a href="#">link</a>	ICLR, 2018
Gradient Matching	GCond [13]	Random Sample	GNN	NC	<a href="#">link</a>	ICLR, 2021
	DosCond [12]	Random Sample	GNN	NC, GC	<a href="#">link</a>	SIGKDD, 2022
	MSGC [7]	Zero Matrix	GNN	NC	—	KBS, 2023
	SGDD [37]	Random Sample	GNN	NC, LP, AD	<a href="#">link</a>	NeurIPS, 2023
	GCARe [24]	—	GNN	NC	—	Appl. Sci. 2023
	CTRL [39]	K-Means	GNN	NC, GC	<a href="#">link</a>	arxiv, 2024
	GroC [17]	Random Sample	GNN	NC, GC	—	arxiv, 2023
	EXGC [4]	Random Sample	GNN	NC	<a href="#">link</a> <sup>1</sup>	WWW 2024
Distribution Matching	MCond [8]	Random Sample	GNN	NC	—	ICDE, 2024
	GCDM [18]	Random Sample	GNN	NC	—	arxiv, 2022
	DM [20, 22]	Random Sample	GNN	NC	—	ICDM, 2023
	GDEM [19]	Eigenbasis Approximation	SD	NC	<a href="#">link</a>	ICML, 2024
Trajectory Matching	FedGKD [26]	Random Noise	GNN	NC	—	arxiv, 2023
	SFGC [42]	K-Center	GNN	NC	<a href="#">link</a>	NeurIPS, 2023
	GEOM [40]	K-Center	GNN	NC	<a href="#">link</a>	ICML, 2024
KRR	GC-SNTK [31]	Random Noise	GNTK	NC	<a href="#">link</a>	WWW, 2024
	KIDD [36]	Random Sample	GNTK	GC	<a href="#">link</a>	SIGKDD, 2023
CTC	Mirage [9]	—	GNN	GC	<a href="#">link</a>	ICLR, 2024

<sup>1</sup> The code repository for EXGC is not fully developed.

62 synthetic dataset. DosCond also modeled the discrete graph structure as a probabilistic  
63 model and each element in the adjacency matrix follows a Bernoulli distribution.

- 64 – **MSGC** [7]: MSGC condenses a large-scale graph into multiple small-scale sparse graphs,  
65 leveraging neighborhood patterns as substructures to enable the construction of various  
66 connection schemes. This process enriches the diversity of embeddings generated by GNNs,  
67 enhances the representation power of GNNs on complex graphs.
- 68 – **SGDD** [37]: SGDD uses graphon approximation to ensure that the structural information  
69 of the original graph is retained in the synthetic, condensed graph. The condensed graph  
70 structure is optimized by minimizing the optimal transport (OT) distance between the original  
71 structure and the condensed structure.
- 72 – **GCARe** [24]: GCARe addresses biases in condensed graphs by regularizing the condensa-  
73 tion process, ensuring that the knowledge of different subgroups is distilled fairly into the  
74 resulting graphs.
- 75 – **CTRL** [39]: CTRL clusters each class of the original graph into sub-clusters and uses these  
76 as initial value for the synthetic graph. By considering both the direction and magnitude  
77 of gradients during gradient matching, it effectively minimizes matching errors during the  
78 condensation phase.
- 79 – **GroC** [17]: GroC uses an adversarial training (bi-level optimization) framework to explore  
80 the most impactful parameter spaces and employs a Shock Absorber operator to apply  
81 targeted adversarial perturbation.
- 82 – **EXGC** [4]: EXGC leverages Mean-Field variational approximation to address inefficiency  
83 in the current gradient matching schemes and uses the Gradient Information Bottleneck  
84 objective to tackle node redundancy.
- 85 – **MCond** [8]: MCond addresses the limitations of traditional condensed graphs in handling  
86 unseen data by learning a one-to-many node mapping from original nodes to synthetic nodes  
87 and uses an alternating optimization scheme to enhance the learning of synthetic graph and  
88 mapping matrix.

89 • **Distribution Matching Methods**

- 90 – **GCDM** [18]: GCDM synthesizes small graphs with receptive fields that share a similar  
91 distribution to the original graph, achieved through a distribution matching loss quantified  
92 by maximum mean discrepancy (MMD).
- 93 – **DM** [20, 22]: DM can be regarded as a one-step variant of GCDM. In DM, the optimization  
94 is centered on the initial parameters. Notably, in [20] and [22], DM does not learn any  
95 structures for efficiency. However, for better comparison in our experiments, we continue to  
96 learn the adjacency matrix.
- 97 – **FedGKD** [26]: FedGKD trains models on condensed local graphs within each client to  
98 mitigate the potential leakage of the training set membership. FedGKD features a Federated  
99 Graph Neural Network framework that enhances client collaboration using a task feature  
100 extractor for graph data distillation and a task relator for globally-aware model aggregation.

#### 101 • **Trajectory Matching Methods**

- 102 – **SFGC** [42]: SFGC uses trajectory matching instead of a gradient matching scheme. It first  
103 trains a set of GNNs on original graphs to acquire and store an expert parameter distribution  
104 offline. The expert trajectory guides the optimization of the condensed graph-free data. The  
105 generated graphs are evaluated using closed-form solutions of GNNs under the graph neural  
106 tangent kernel (GNTK) ridge regression, avoiding iterative GNN training.
- 107 – **GEOM** [40]: GEOM makes the first attempt toward lossless graph condensation using  
108 curriculum-based trajectory matching. A homophily-based difficulty score is assigned to  
109 each node and the easy nodes are learned in the early stages while more difficult ones are  
110 learned in the later stages. On top of that, GEOM incorporated a Knowledge Embedding  
111 Extraction (KEE) loss into a matching loss.

#### 112 • **Kernel Ridge Regression Methods**

- 113 – **GC-SNTK** [31]: GC-SNTK introduces a Structure-based Neural Tangent Kernel(SNTK)  
114 to capture graph topology, replacing the inner GNNs training in traditional GC paradigm,  
115 avoiding multiple iterations.
- 116 – **KiDD** [36]: KiDD uses kernel ridge regression (KRR) with a graph neural tangent kernel  
117 (GNTK) for graph-level tasks. To enhance efficiency, KiDD introduces LiteGNTK, a  
118 simplified GNTK, and proposes KiDD-LR for faster low-rank approximation and KiDD-D  
119 for handling discrete graph topology using the Gumbel-Max reparameterization trick. We  
120 use KiDD-LR for experiments as it has generally demonstrated better performance compared  
121 to KiDD-D.

#### 122 • **Computation Tree Compression Methods**

- 123 – **Mirage** [9]: Mirage decomposes graphs in datasets into a collection of computation trees  
124 and then mines frequently co-occurring trees from this set. Mirage then uses aggregation  
125 functions (MEANPOOL, SUMPOOL, etc.) on the embeddings of the root node of each tree  
126 to approximate the graph embedding.

### 127 **A.3 Hyper-Parameter Setting**

128 For the implementation of various graph condensation methods, we adhere to the default parameters  
129 as specified by the authors in their respective original implementations. This approach ensures that  
130 our results are comparable to those reported in the foundational studies. For condensation ratios  
131 that were not explored in the original publications, we employ a grid search strategy to identify  
132 the optimal hyperparameters within the predefined search space. This includes experimenting with  
133 various combinations, such as differing learning rates for the feature optimizer and the adjacency  
134 matrix optimizer. The corresponding hyperparameter space are shown in Table A3.

### 135 **A.4 Computation resources**

136 All experiments were conducted on a high-performance GPU cluster to ensure a fair comparison.  
137 The cluster consists of 32 identical dell-GPU nodes, each featuring 256GB of memory, 2 Intel Xeon  
138 processors, and 4 NVIDIA Tesla V100 GPUs, with each GPU having 64 GB of GPU memory. If any  
139 experiment setting exceeds the GPU memory limit, it is reported as out-of-memory (OOM).

Table A3: Hyperparameter search space of different methods

Method	Hyperparameter	Values
General Settings	Learning Rate	0.1, 0.01, 0.001, 0.0001, 0.00001
	Epochs	300, 400, 500, 800, 1000, 2000, 3000, 4000, 5000
	Layers	2, 3
	Dropout Rate	0, 0.05, 0.1, 0.5, 0.6, 0.7, 0.8
	Weight Decay	0, 0.0005
	Hidden Units	128, 256
	Pooling	sum, mean
	Activation	LeakyReLU, ReLU, Sigmoid, Softmax
	Batch Size	(16,6000)
SGDD	mx_size	50, 100
	opt_scale	5, 10
GCond, DosCond, SGDD, GCDM, DM	outer loop	1, 2, 5, 10, 15, 20
GCond, SGDD, GCDM	inner loop	1, 5, 10, 15, 20
SFGC, GEOM	expert_epochs	50, 70, 100, 350, 600, 800, 1000, 1500, 1600, 1900
	start_epoch	10, 20, 50, 100, 200, 300
	teacher_epochs	800, 1000, 1200, 2400, 3000
GEOM	lam	0.6, 0.7, 0.75, 0.8, 0.85, 0.9, 0.95
	T	250, 500, 600, 800, 1000, 1200
	scheduler	linear, geom, root
KiDD	scale	uniform, degree
	rank	8, 16, 32
	orth_reg	0.01, 0.001, 0.0001

## 140 A.5 Discussion on Existing Benchmarks

141 To the best of our knowledge, the only concurrent work is GCondenser [21]. The comparison  
 142 of GCondser and our GC-Bench are list in Table A4. GCondenser [21] focus the node-level GC  
 143 methods for node classification on homogeneous graphs with limited evaluation dimensions in terms  
 144 of performance and time efficiency. Our GC-Bench analyzes more GC methods on a wider variety of  
 145 datasets (both homogeneous and heterogeneous) and tasks (node classification, graph classification),  
 146 encompassing both node-level and graph-level methods. In addition to performance and efficiency  
 147 analysis, we further explore the transferability across different tasks (link prediction, node clustering,  
 148 anomaly detection) and backbones (GNN models and the popular Graph Transformer). With GC-  
 149 Bench covering more in-depth investigation over a wider scope, we believe it will provide valuable  
 150 insights into existing works and future directions.

## 151 B Settings and Additional Results

152 In this section, we provide more details of the experimental settings and the additional results for the  
 153 proposed 6 research questions, respectively.

### 154 B.1 Settings and Additional Results of Performance Comparison (RQ1)

#### 155 B.1.1 Comparison Setting

156 **Node Classification Graph Dataset Setting.** We compared ten state-of-the-art GC methods. The  
 157 selection of the condensation ratio  $r$  is based on the labeling rates of different datasets. For datasets  
 158 like *Cora* and *Citeseer*, the labeling rates are less than 50%, we select  $r$  as a proportion of the labeling  
 159 rate, specifically at {5%, 10%, 25%, 50%, 75%, 100%}. For datasets like *ogbn-arxiv*, and inductive  
 160 datasets where all nodes in the training graphs are labeled, with a relatively higher labeling rate,  $r$   
 161 is chosen to be {5%, 10%, 25%, 50%, 75%, 100%}. Corresponding condensation rates are shown in  
 162 Table B2.

163 **Graph Classification Graph Dataset Setting.** We compared three state-of-the-art GC algorithms  
 164 on graph classification datasets: *DosCond* [12], *KiDD* [36], and *Mirage* [9]. *Mirage* [9] does not  
 165 condense datasets into unified graphs measurable by Graphs per Class(GPC) as *DosCond* [12] and

Table A4: Comparison of GCondenser and GC-Bench

Benchmark Coverage		GCondenser	GC-Bench	
Algorithms	Traditional Core-set Methods	Random, K-Center	Random, K-Center, Herding	
	Gradient Matching	GCond, DosCond, SGDD	GCond, DosCond, SGDD	
	Distribution Matching	GCDM, DM	GCDM, DM	
	Trajectory Matching	SFGC	SFGC, GEOM	
	KRR	—	KiDD	
	CTC	—	Mirage	
Datasets	Node-level Homogeneous	<i>Cora, Citeseer, ogbn-arxiv Flickr, Reddit, PubMed</i>	<i>Cora, Citeseer, ogbn-arxiv Flickr, Reddit</i>	
	Node-level Heterogeneous	—	<i>ACM, DBLP</i>	
	Graph-level	—	<i>NCII, DD, ogbg-molbase ogbg-molbbbp, ogbg-molhiv</i>	
Tasks	Node-level	node classification	node classification link prediction node clustering anomaly detection	
	Graph-level	—	graph classification	
Evaluation Dimensions	Perf.	Condensation Ratios	✓	
		Impact of Structure	structure v.s. structure-free	
		Impact of Initialization	✓	
	Trans.	Backbone Trans.	SGC and GCN transfer to SGC, GCN, GraphSAGE, APPNP, ChebyNet, MLP	SGC, GCN and Graph Transformer transfer to SGC, GCN, GraphSAGE, APPNP, ChebyNet, MLP, Graph Transformer
		Task Trans.	—	node classification link prediction node clustering anomaly detection
	Efficiency	Time	✓	✓
Space		—	✓	

166 *KiDD* [36] do. Therefore, we measure the condensed dataset size by storing its elements in .pt format,  
167 similar to *DosCond* [12] and *KiDD* [36]. We select the *Mirage*-condensed dataset size closest to  
168 *DosCond*'s as the corresponding GPC. *KiDD* [36] generally occupies more disk space than *DosCond*  
169 under the same GPC. The size of *Mirage* datasets is determined by two parameters: the number of  
170 GNN layers ( $L$ ) and the frequency threshold  $\Theta$ . We fix  $L = 2$ , consistent with the 2-layer model used  
171 for validation, and employ a grid search strategy to identify the threshold combination that yields a  
172 dataset size closest to the targeted GPC. The corresponding disk space, GPC, and threshold choices  
173 are presented in Table B1. Note that for small thresholds, the MP Tree search algorithms used in  
174 *Mirage* [9] may reach recursive limits. Consequently, in *DD* and *ogbg-molbase*, certain GPCs lack  
175 corresponding threshold values.

176 **Heterogeneous Graph Dataset Setting.** Due to the absence of condensation methods specifically  
177 for heterogeneous graphs, we convert heterogeneous datasets into homogeneous graphs for conden-  
178 sation, focusing on target nodes. We uniformly summed the adjacency matrices corresponding to  
179 various meta-paths as in [23], and applied zero-padding to match the maximum feature dimension as  
180 well as one-hot encoding for nodes without features. Specifically, in *GEOM* [40], when calculating  
181 heterophily, all nodes without labels (non-target nodes) are assigned the same distinct label, ensuring  
182 a consistent heterophily calculation.

### 183 B.1.2 Additional Results

184 The graph classification performance on GCN is shown in Table B3. *DosCond* [12] with GCN  
185 demonstrates significant advantages in 12 out of 25 cases, while *KiDD* [36] underperforms in most  
186 scenarios. Notably, *DosCond* [12] and *Mirage* [9] even outperform the results of the whole dataset on

Table B1: Comparison of Disk Size and Graph per Class (GPC) for condensed datasets between *Mirage* and *DosCond*.

Dataset	Graph/ Cls	Mirage Disk Size (Bytes)	DosCond Disk Size (Bytes)	Class 0 Threshold	Class 1 Threshold
<i>NCII</i> [29]	1	14,455	18,425	451	441
	5	81,622	82,745	351	381
	10	142,228	162,301	301	291
	20	195,609	324,035	251	231
	50	995,277	806,403	201	171
<i>DD</i> [3]	1	38,352	855,077	15	9
	5	—	4,265,957	—	—
	10	—	8,529,583	—	—
	20	—	17,056,751	—	—
	50	—	42,638,383	—	—
<i>ogbg-molbace</i> [11]	1	13,836	14143	120	90
	5	60,047	60,927	230	80
	10	106,077	119,497	120	80
	20	232,191	236,489	140	70
	50	—	587,337	—	—
<i>ogbg-molbbbp</i> [11]	1	8,817	8,831	29	198
	5	34,699	34,175	49	109
	10	66,433	65,929	30	90
	20	104,091	129,289	20	80
	50	324,425	319,369	17	87
<i>ogbg-molhiv</i> [11]	1	9,606	9,717	8,000	250
	5	54,669	38,837	1,760	170
	10	74,524	75,263	1,680	130
	20	148,028	148,095	1,420	110
	50	330,498	366,463	800	110

Table B2: Different condensation ratios of transductive datasets. For heterogeneous datasets, the number of nodes in the original graph is the sum of all types of nodes.

Ratio ( $r$ )	<i>Cora</i>	<i>Citeseer</i>	<i>ACM</i>	<i>DBLP</i>
5%	0.26%	0.18%	0.003%	0.002%
10%	0.52%	0.36%	0.007%	0.004%
25%	1.30%	0.90%	0.013%	0.007%
50%	2.60%	1.80%	0.033%	0.019%
75%	3.90%	2.70%	0.066%	0.037%
100%	5.20%	3.60%	0.332%	0.186%

187 *ogbg-molbace*. For *Mirage* [9], due to the algorithm’s recursive depth under low threshold parameters,  
 188 we have only one result corresponding to GPC 1 on *DD*. However, this single result already surpasses  
 189 all datasets condensed by *KiDD* [36] and the dataset with GPC 1 condensed by *DosCond*.

## 190 B.2 Settings and Additional Results of Structure in Graph Condensation (RQ2)

### 191 B.2.1 Experimental Settings

192 The homophily ratio we use is the edge homophily ratio, which represents the fraction of edges that  
 193 connect nodes with the same labels. It can be calculated as:

$$\mathcal{H}(G) = \frac{1}{|\mathcal{E}|} \sum_{(j,k) \in \mathcal{E}} \mathbf{1}(y_j = y_k), \quad i \in \mathcal{V}, \quad (\text{A.1})$$

194 where  $\mathcal{V}$  is the node set,  $\mathcal{E}$  is the edge set,  $|\mathcal{E}|$  is the number of edges in the graph,  $y_i$  is the label of  
 195 node  $i$  and  $\mathbf{1}(\cdot)$  is the indicator function. A graph is typically considered to be highly homophilous  
 196 when  $\mathcal{H}$  is large (typically,  $0.5 \leq \mathcal{H} \leq 1$ ), such as *Cora* and *Reddit*. Conversely, a graph with a low  
 197 edge homophily ratio is considered to be heterophilous, such as *Flickr*.

198 We also calculate the homophily ratio of condensed datasets. Since the condensed datasets have  
 199 weighted edges, we first sparsify the graph by removing all edges with weights less than 0.05,

Table B3: **Graph classification performance on GCN** (mean±std) across datasets with varying condensation ratios  $r$ . The best results are shown in **bold** and the runner-ups are shown in underlined. **Red** color highlights entries that exceed the whole dataset values.

Dataset	Graph /Cls	Ratio( $r$ )	Traditional Core-set methods			Gradient	KRR	CTC	Whole Dataset
			Random	Herding	K-Center	DosCond	KIDD	Mirage	
<i>NCII</i> Acc. (%)	1	0.06%	53.30±0.6	55.20±2.6	55.20±2.6	<b>57.30</b> ±0.9	49.30±1.1	49.10±0.9	71.1±0.8
	5	0.24%	55.00±1.4	56.50±0.9	53.20±0.6	<b>58.40</b> ±1.4	56.10±1.0	49.60±2.2	
	10	0.49%	<u>58.10</u> ±2.2	<b>58.60</b> ±0.8	57.00±2.6	57.80±1.6	57.50±1.1	48.60±0.1	
	20	0.97%	54.40±0.8	59.10±1.1	<b>60.10</b> ±1.3	<b>60.10</b> ±3.2	56.40±0.6	48.70±0.0	
	50	2.43%	56.80±1.1	58.70±1.1	<b>64.40</b> ±0.9	58.20±2.8	<u>59.90</u> ±0.6	48.60±0.1	
<i>DD</i> Acc. (%)	1	0.21%	59.70±1.5	66.90±2.8	66.90±2.8	68.30±6.6	58.60±2.4	<b>71.20</b> ±6.6	78.4±1.7
	5	1.06%	61.90±1.1	<u>66.20</u> ±2.5	62.00±1.7	<b>73.10</b> ±2.2	58.60±1.1	-	
	10	2.12%	63.70±2.8	<u>68.00</u> ±3.6	62.50±2.3	<b>71.30</b> ±8.3	61.60±3.8	-	
	20	4.25%	64.70±5.3	<u>69.70</u> ±0.8	63.10±1.9	<b>73.00</b> ±5.8	62.60±1.4	-	
	50	10.62%	66.60±2.1	68.50±1.4	<u>68.90</u> ±1.8	<b>74.20</b> ±3.6	59.30±0.0	-	
<i>ogbg-molbace</i> ROC-AUC	1	0.17%	0.510±.083	0.515±.040	0.517±.044	0.658±.064	0.568±.047	<b>0.733</b> ±.012	0.711±.019
	5	0.83%	0.612±.036	0.653±.043	0.508±.087	0.691±.06	0.356±.022	<b>0.760</b> ±.002	
	10	1.65%	0.620±.054	0.658±.046	0.646±.047	0.702±.045	0.542±.027	<b>0.759</b> ±.002	
	20	3.31%	0.642±.053	0.631±.051	0.575±.03	0.659±.049	0.526±.014	<b>0.761</b> ±.003	
	50	8.26%	<u>0.677</u> ±.015	0.629±.053	0.576±.087	<b>0.714</b> ±.032	0.446±.042	-	
<i>ogbg-molbbbp</i> ROC-AUC	1	0.12%	0.534±.041	0.560±.017	0.560±.017	<b>0.600</b> ±.023	0.504±.042	<b>0.600</b> ±.002	0.646±.013
	5	0.61%	0.561±.014	0.574±.022	<u>0.585</u> ±.005	0.579±.056	0.561±.004	<b>0.609</b> ±.061	
	10	1.23%	0.566±.011	<u>0.590</u> ±.024	<b>0.598</b> ±.025	0.556±.063	0.550±.005	0.517±.028	
	20	2.45%	0.593±.023	0.568±.019	0.545±.009	0.590±.057	0.594±.022	<b>0.626</b> ±.032	
	50	6.13%	0.587±.007	0.579±.022	<b>0.621</b> ±.011	0.598±.024	<u>0.603</u> ±.01	0.602±.018	
<i>ogbg-molhiv</i> ROC-AUC	1	0.01%	<u>0.733</u> ±.008	0.727±.012	0.727±.012	<b>0.734</b> ±.002	0.725±.007	0.728±.012	0.750±.010
	5	0.03%	0.729±.006	0.720±.018	<b>0.739</b> ±.01	0.736±.008	0.738±.003	0.717±.003	
	10	0.06%	0.724±.011	0.726±.014	0.723±.012	<b>0.736</b> ±.007	0.731±.008	<b>0.735</b> ±.028	
	20	0.12%	0.723±.015	<u>0.726</u> ±.015	0.724±.01	<b>0.733</b> ±.007	0.703±.097	0.710±.016	
	50	0.30%	0.712±.014	<u>0.723</u> ±.019	0.721±.012	<b>0.731</b> ±.011	0.723±.011	0.718±.022	

\**Mirage* cannot directly generate graphs with the required ratio. Thus, we search the parameter space and aligned the generated graph to match *DosCond*'s disk usage as substitution (see Appendix B.1).

200 then calculate the homophily ratio by adjusting the fraction to a weighted fraction, which can be  
201 represented as:

$$\mathcal{H}(G) = \frac{\sum_{(j,k) \in \mathcal{E}} w_{jk} \mathbf{1}(y_j = y_k)}{\sum_{(j,k) \in \mathcal{E}} w_{jk}}, \quad i \in \mathcal{V}, \quad (\text{A.2})$$

202 where  $w_{jk}$  is the weight of the edge between nodes  $j$  and  $k$ .

## 203 B.2.2 Additional Results

204 The results of homophily ratios of condensed datasets are shown in Table B4. It appears that  
205 condensed datasets often struggle to preserve the homophily properties of the original datasets. For  
206 instance, in the case of the heterophilous dataset *Flickr*, an increase in the homophily rate is observed  
207 under most methods and ratios.

Table B4: Homophily ratio comparison of different condensed datasets

	Whole Dataset	Ratio ( $r$ )	GCDM	DM	DosCond	GCond	SGDD
<i>Cora</i>	0.81	1.30%	0.76 ↓	0.88 ↑	0.20 ↓	0.64 ↓	0.19 ↓
		2.60%	0.11 ↓	0.74 ↓	0.16 ↓	0.55 ↓	0.19 ↓
		5.20%	1.00 ↑	0.21 ↓	0.15 ↓	0.62 ↓	0.15 ↓
<i>Citeseer</i>	0.74	0.90%	0.16 ↓	0.75 ↑	0.19 ↓	0.57 ↓	0.14 ↓
		1.80%	0.08 ↓	0.30 ↓	0.20 ↓	0.36 ↓	0.19 ↓
		3.60%	1.00 ↑	0.34 ↓	0.15 ↓	0.22 ↓	0.15 ↓
<i>Flickr</i>	0.24	0.05%	0.28 ↑	0.29 ↑	0.25 ↑	0.28 ↑	0.32 ↑
		0.50%	0.29 ↑	0.22 ↓	0.08 ↓	0.28 ↑	0.30 ↑
		1.00%	0.36 ↑	0.18 ↓	0.06 ↓	0.28 ↑	0.26 ↑

208 We visualize the condensed datasets using force-directed graph visualization, as shown in Figure B.1,  
209 Figure B.2, and Figure B.3. Since *SFGC* [42] and *GEOM* [40] synthesize edge-free datasets, we



210 do not visualize the datasets they condensed. As shown in the visualization, graphs condensed  
 211 by different methods exhibit distinct structural characteristics. For example, distribution matching  
 212 methods often result in less pronounced community structures compared to other methods.

213 We also visualize the node degree distribution of the original graph and the condensed graphs in  
 214 Figure B.4. Note that the graphs condensed by *GCDM* [18] and *DM* [22] are dense and each edge has  
 215 an extremely small weight under most situations, the degree of each node is also small. We observe  
 216 that the degree distributions of most condensed datasets deviate significantly from the original graph.  
 217 Among them, *SGDD* [37] demonstrates a relatively similar degree distribution to that of the original  
 218 graph.

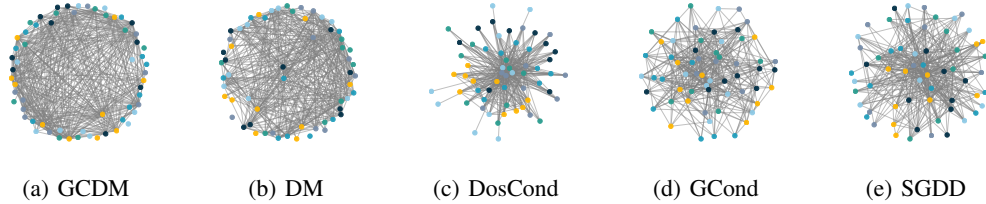


Figure B.1: Visualization of the Condensed *Citeseer* (1.80%) Dataset. Only the top 20% of edges ranked by weight are visualized.

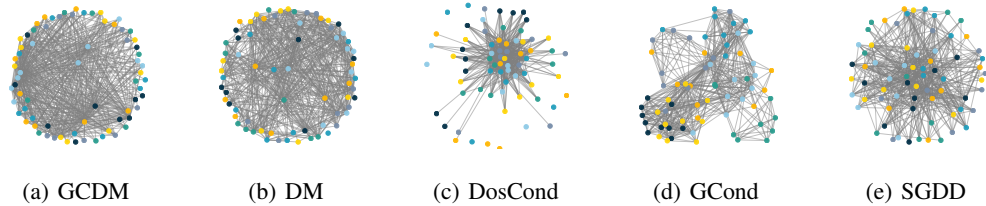


Figure B.2: Visualization of the Condensed *Cora* (2.60%) Dataset. Only the top 20% of edges ranked by weight are visualized.

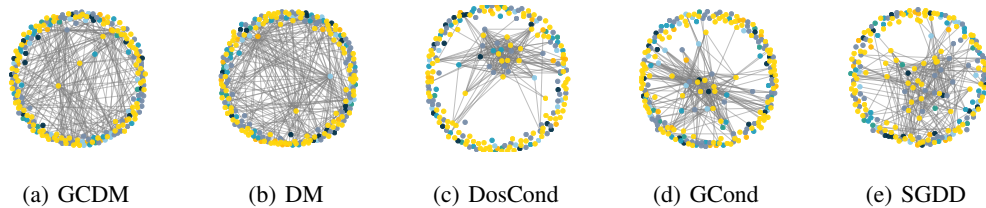


Figure B.3: Visualization of the Condensed *Flickr* (0.50%) Dataset. Only the top 1% of edges ranked by weight are visualized.

### 219 B.3 Settings and Additional Results of Transferability in Different Tasks (RQ3)

#### 220 B.3.1 Link Prediction

221 For the link prediction task, we utilize a graph autoencoder (GAE)[15] based on Graph Convolutional  
 222 Networks (GCN[14]). The GAE consists of a two-layer GCN encoder that creates node embeddings.  
 223 During training, we enhance the dataset by randomly adding negative links and use a decoder to  
 224 perform binary classification on edges. During evaluation, we test the model using the test set of the  
 225 original graph. Since trajectory matching methods do not generate any edges, we do not use them  
 226 for link prediction tasks. The results of condensed datasets on the link prediction task are shown in

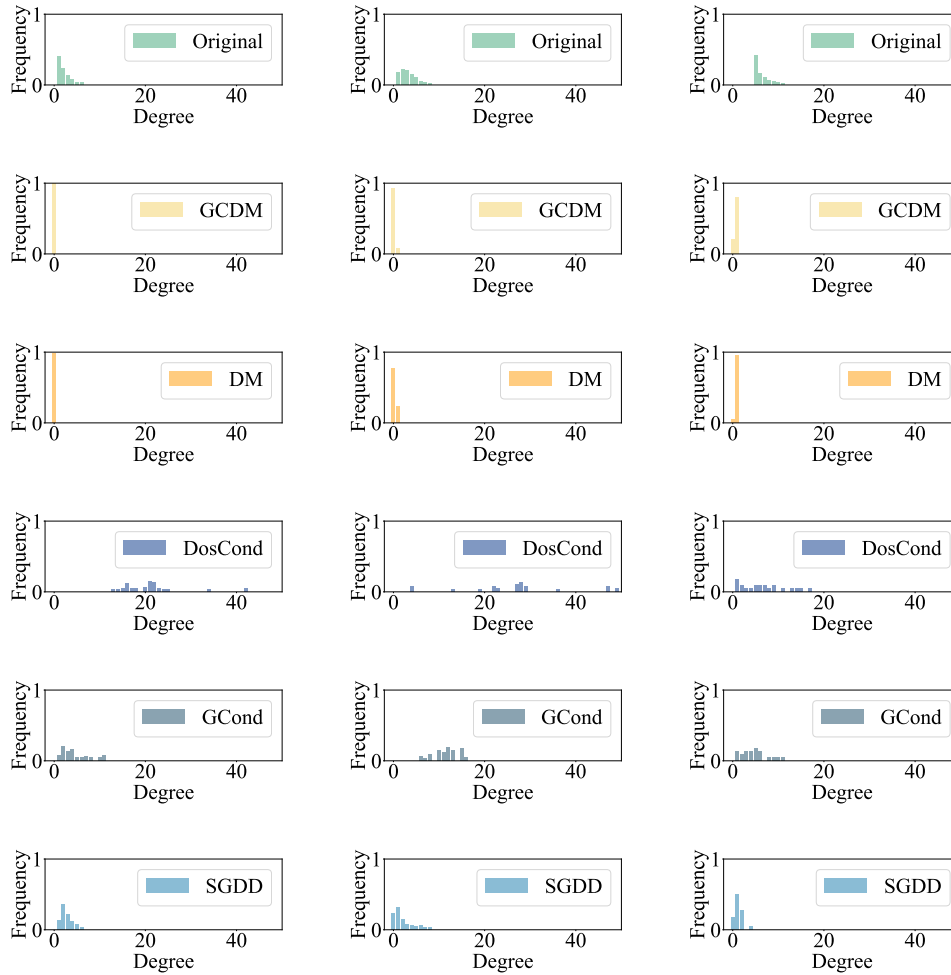


Figure B.4: Degree distribution in the condensed graphs for *Citeseer* (1.80%), *Cora* (2.60%), and *Flickr* (0.05%). The first, second, and third columns correspond to *Citeseer*, *Cora*, and *Flickr*, respectively.

227 Table B5. We observe that most condensed datasets underperform in link prediction tasks, especially  
 228 on *ogbn-arxiv* and *Flickr*. Most methods' condensed datasets consistently underperform compared to  
 229 traditional core-set methods, indicating room for improvement.

### 230 B.3.2 Node Clustering

231 For the node clustering tasks on condensed datasets, we utilize DAEGC [30] to train on synthetic  
 232 datasets condensed using the node classification task. We then test the trained model on the original  
 233 large-scale datasets and include the results of other methods on the original graph for comprehensive  
 234 comparison. Due to the performance degradation of GAT with large neighborhood sizes, we use  
 235 GCN as the encoder. Performance metrics include Accuracy (Acc.), Normalized Mutual Information  
 236 (NMI), F-score, and Adjusted Rand Index (ARI).

237 To fully leverage the condensed datasets, we include the results of node clustering with perturbing. In  
 238 this experiment, the GCN encoder is first trained on the synthetic datasets with a node classification  
 239 task, which incorporates the synthetic labels' information. Using the pre-trained GCN as an encoder,  
 240 we then perform node clustering on the synthetic datasets and the original graph. Results of node  
 241 clustering tasks, both without and with perturbing are shown in Table B6 and Table B7 respectively.  
 242 We observe that most condensed datasets perform worse in the node clustering task compared to the  
 243 original dataset. However, when additional information from the pretraining of the node classification

Table B5: **Link Prediction Accuracy (%)** of different condensed datasets. The best results are shown in **bold**.

Dataset	Ratio ( $r$ )	Random	Herding	K-Center	GCDM	DM	DosCond	GCond	SGDD	Whole Dataset
Citeseer	0.90%	0.52	0.52	0.55	0.53	0.53	0.50	0.65	<b>0.69</b>	0.82
	1.80%	0.52	0.52	0.54	0.51	0.52	0.51	0.51	<b>0.67</b>	
	3.60%	0.54	0.53	0.53	0.53	0.53	0.53	0.53	<b>0.61</b>	
Cora	1.30%	0.58	0.54	0.58	<b>0.72</b>	0.71	0.67	0.61	0.51	0.78
	2.60%	0.55	0.55	0.56	0.69	0.67	0.58	<b>0.77</b>	0.62	
	5.20%	0.57	0.56	0.58	0.70	<b>0.71</b>	0.59	0.65	0.56	
ogbn-arxiv	0.05%	<b>0.76</b>	0.68	0.67	0.66	0.68	0.63	0.60	0.70	0.75
	0.20%	0.72	0.72	<b>0.73</b>	0.72	0.72	0.69	0.71	0.51	
	0.50%	<b>0.74</b>	0.73	<b>0.74</b>	0.71	0.73	0.72	0.72	0.70	
Flickr	0.05%	0.55	0.54	0.53	<b>0.60</b>	0.53	0.52	0.54	0.51	0.75
	0.20%	0.63	0.63	0.63	0.63	0.51	0.53	0.57	<b>0.70</b>	
	0.50%	<b>0.70</b>	0.68	<b>0.70</b>	0.56	0.65	0.62	0.67	0.61	

244 task on condensed dataset is utilized, the results of node clustering significantly improve. Notably,  
 245 some datasets in Table B6 exhibit identical results with the Adjusted Rand Index (ARI) being 0 or  
 246 even negative. This occurs because the clustering results do not match the number of classes in  
 247 the labels, requiring manual splitting of clusters in such scenarios. An ARI of 0 indicates that the  
 248 clustering result is as good as random, while a negative ARI suggests it is worse than random.

### 249 B.3.3 Anomaly Detection

250 For the anomaly detection tasks, we generate two types of anomalies, *Contextual Anomalies* and  
 251 *Structural Anomalies*, following the method described in [2]. We set the anomaly rate to 0.05; if the  
 252 condensed dataset is too small, we inject one contextual anomaly and two structural anomalies.

253 **Contextual Anomalies:** Each outlier is generated by randomly selecting a node and substituting its  
 254 attributes with those from another node with the maximum Euclidean distance in attribute space.

255 **Structural Anomalies:** Outliers are generated by randomly selecting a small group of nodes and  
 256 making them fully connected, forming a clique. The nodes in this clique are then regarded as structural  
 257 outliers. This process is repeated iteratively until a predefined number of cliques are generated.

258 We conduct anomaly detection by training a DOMINANT model [2], which features a shared graph  
 259 convolutional encoder, a structure reconstruction decoder, and an attribute reconstruction decoder.  
 260 Initially, we inject predefined anomalies into the test set of the original graph and use it for evaluation  
 261 across different condensed datasets derived from this graph. The model is then trained on these  
 262 condensed datasets, which were injected with specific types of anomalies before training. The  
 263 DOMINANT model measures reconstruction errors as anomaly scores for both the graph structure  
 264 and node attributes, combining these scores to detect anomalies. The results are evaluated using the  
 265 ROC-AUC metric, as shown in Table B8 and B9.

## 266 B.4 Settings and Additional Results of Transferability across Backbone Model Architectures (RQ4)

### 268 B.4.1 Experimental Settings

269 For transferability evaluation, we use different models as backbones to test the condensation methods.  
 270 For distribution matching methods, two backbone models with shared parameters are used to generate  
 271 embeddings that are matched. For trajectory matching methods, two backbone models are used  
 272 to generate expert trajectories and student trajectories, respectively, to match the corresponding  
 273 parameters. For gradient matching methods, two backbone models with shared parameters are  
 274 used to generate gradients for real and synthetic data. Models are selected using grid-searched  
 275 hyperparameters. The details of the backbone architecture are as follows:

- 276 • **MLP:** MLP is a simple neural network consisting of fully connected layers. The MLP we use  
 277 is structured similarly to a GCN but without the adjacency matrix input, effectively functioning

Table B6: **Node Clustering without Pretraining Results** on *Cora* and *Citeseer* with varying condensation ratios ( $r$ ). The best results are highlighted in **bold**, the runner-ups are underlined, and the best results of condensed datasets are shaded in gray .

Methods	Ratio ( $r$ )	<i>Citeseer</i>				Ratio( $r$ )	<i>Cora</i>			
		Acc.	NMI	ARI	F1		Acc.	NMI	ARI	F1
K-means		54.4	31.2	28.5	41.3		50.0	31.7	<u>37.6</u>	23.9
DAEGC [30]	Full	<b>67.2</b>	<b>39.7</b>	<b>41.0</b>	<b>63.6</b>	Full	<b>70.4</b>	<b>52.8</b>	<b>68.2</b>	<u>49.6</u>
Random	0.90%	40.6	19.1	17.5	36.0	1.30%	36.6	13.5	9.0	34.3
	1.80%	38.3	14.8	13.6	34.5	2.60%	33.5	13.9	7.1	33.4
	3.60%	41.8	18.1	16.9	39.4	5.20%	30.2	0.4	0.0	6.8
Herding	0.90%	41.9	16.9	15.3	40.0	1.30%	37.4	18.2	11.7	35.0
	1.80%	44.9	18.7	16.0	41.1	2.60%	36.6	16.4	11.9	34.0
	3.60%	58.1	27.8	29.2	52.3	5.20%	26.7	13.7	2.9	20.6
K-Center	0.90%	37.9	13.4	11.1	35.2	1.30%	34.3	13.5	7.8	32.4
	1.80%	50.0	23.5	22.9	46.5	2.60%	42.5	22.3	15.0	42.3
	3.60%	31.9	14.0	10.2	31.0	5.20%	30.2	0.4	0.0	6.8
GCDM	0.90%	41.4	16.9	16.2	38.6	1.30%	30.2	0.4	0.0	6.8
	1.80%	44.1	18.1	18.1	38.8	2.60%	30.2	0.4	0.0	6.8
	3.60%	22.8	1.8	1.2	20.9	5.20%	30.2	0.4	0.0	6.8
DM	0.90%	23.5	2.1	1.1	17.7	1.30%	30.2	0.4	0.0	6.8
	1.80%	45.3	19.1	17.7	42.9	2.60%	29.2	2.0	0.0	9.5
	3.60%	25.9	4.5	3.5	20.0	5.20%	30.2	0.4	0.0	6.8
DosCond	0.90%	28.6	10.2	6.3	25.1	1.30%	30.2	0.4	0.0	6.8
	1.80%	57.1	31.4	26.2	49.5	2.60%	30.2	0.4	0.0	6.8
	3.60%	44.3	20.6	17.0	38.6	5.20%	29.6	16.2	7.7	23.4
GCond	0.90%	61.8	34.0	34.7	55.9	1.30%	46.6	36.7	27.3	41.2
	1.80%	59.6	33.0	32.6	50.3	2.60%	49.9	39.3	27.9	44.3
	3.60%	57.8	32.0	30.2	54.8	5.20%	44.6	<u>40.9</u>	25.1	37.3
SGDD	0.90%	56.5	27.3	26.8	50.6	1.30%	30.2	0.4	0.0	6.8
	1.80%	45.4	24.0	20.0	43.9	2.60%	30.2	0.4	0.0	6.8
	3.60%	42.5	23.6	20.8	38.2	5.20%	33.2	17.9	8.8	25.5
SFGC	0.90%	46.7	19.9	18.8	43.4	1.30%	42.1	23.5	17.7	39.2
	1.80%	56.8	27.4	27.6	52.8	2.60%	54.4	31.8	26.4	<b>50.2</b>
	3.60%	47.7	19.0	16.9	45.3	5.20%	30.1	0.4	-0.1	6.8
GEOM	0.90%	41.4	16.9	16.2	38.6	1.30%	40.7	16.9	11.6	37.3
	1.80%	44.1	18.1	18.1	38.8	2.60%	30.8	12.9	9.3	29.2
	3.60%	22.8	1.8	1.2	20.9	5.20%	35.6	16.0	11.5	33.6

278 as a standard multi-layer perceptron (MLP). The MLP we adopted consists of 2 layers with 256  
279 hidden units in each layer.

- 280 • **GCN [14]:** GCN is the most common architecture for evaluating condensed datasets in main-  
281 stream GC methods. GCN defines a localized, first-order approximation of spectral graph  
282 convolutions, effectively aggregating and combining features from a node’s local neighborhood,  
283 leveraging the graph’s adjacency matrix to update node representations through multiple layers.  
284 We adhere to the setting in previous work [13] and use 2 graph convolutional layers for node  
285 classification, each followed by ReLu activation and batch normalization depending on the config-  
286 uration. For graph classification, we use a 3-layer GCN with a sum pooling function. The hidden  
287 unit size is set to 256.
- 288 • **SGC [34]:** SGC is the standardized model used for condensation in previous works. It can be  
289 regarded as a simplified version of GCN, which ignores the nonlinear activation function but still  
290 keeps two Graph Convolution layers, thereby preserving similar graph filtering behaviors. In the  
291 experiments, we use 2-layer SGC with no bias.
- 292 • **Cheby [1]:** Cheby utilizes Chebyshev polynomials to approximate the graph convolution oper-  
293 ations, which retains the essential graph filtering properties of GCN while reducing the com-

Table B7: **Node Clustering with Pretraining Results** on *Cora* and *Citeseer* with varying condensation ratios ( $r$ ). The best results are highlighted in **bold** and the runner-ups are underlined.

Methods	Ratio ( $r$ )	<i>Citeseer</i>				Ratio ( $r$ )	<i>Cora</i>			
		Acc.	NMI	ARI	F1		Acc.	NMI	ARI	F1
Random	0.90%	27.3	5.5	4.7	24.6	1.30%	41.7	15.8	13.5	37.3
	1.80%	32.7	9.7	7.8	31.4	2.60%	36.5	14.6	9.1	35.4
	3.60%	44.6	16.0	14.1	43.0	5.20%	44.4	23.5	14.9	45.7
Herding	0.90%	36.7	12.8	11.1	34.4	1.30%	40.7	18.3	12.9	40.0
	1.80%	36.8	13.1	10.2	36.2	2.60%	36.1	14.6	8.7	34.9
	3.60%	39.4	16.9	14.1	38.1	5.20%	35.0	16.6	10.9	32.0
K-Center	0.90%	33.7	9.7	8.3	29.5	1.30%	41.8	19.3	14.5	39.2
	1.80%	37.6	15.6	13.9	34.9	2.60%	38.5	20.8	14.8	38.3
	3.60%	41.7	17.1	14.3	40.5	5.20%	38.5	17.4	10.9	36.3
GCDM	0.90%	31.1	9.6	6.6	27.3	1.30%	21.3	3.7	1.7	20.1
	1.80%	33.1	11.9	11.1	30.4	2.60%	27.0	10.9	5.7	26.7
	3.60%	39.7	18.0	15.2	34.4	5.20%	30.0	12.4	7.0	29.6
DM	0.90%	36.5	15.7	12.9	30.0	1.30%	27.3	9.3	4.5	25.7
	1.80%	37.1	10.6	8.6	31.4	2.60%	20.8	3.3	0.9	19.0
	3.60%	29.2	6.0	4.0	23.6	5.20%	23.5	4.8	1.6	16.3
DosCond	0.90%	<b>62.7</b>	<b>35.9</b>	<b>35.1</b>	<b>60.6</b>	1.30%	60.2	42.5	29.4	61.2
	1.80%	45.2	17.9	15.4	40.8	2.60%	44.5	30.1	16.6	46.5
	3.60%	<u>58.6</u>	29.6	28.5	<u>55.8</u>	5.20%	25.4	9.8	5.0	25.0
GCond	0.90%	44.0	22.5	18.7	40.3	1.30%	67.4	45.1	40.4	65.8
	1.80%	58.5	<u>30.9</u>	<u>29.6</u>	54.9	2.60%	<u>63.7</u>	44.5	36.2	61.8
	3.60%	52.0	<u>26.8</u>	<u>22.5</u>	46.6	5.20%	60.9	<u>47.1</u>	37.1	56.0
SGDD	0.90%	46.7	23.5	19.1	42.3	1.30%	65.1	44.6	37.1	64.6
	1.80%	55.4	28.0	25.8	50.9	2.60%	35.7	19.2	11.7	34.8
	3.60%	40.5	18.3	14.3	34.8	5.20%	<b>74.8</b>	<b>51.9</b>	<b>53.1</b>	<b>72.8</b>
SFGC	0.90%	34.2	9.8	8.4	32.2	1.30%	41.2	21.2	13.9	40.2
	1.80%	47.1	21.7	20.6	43.5	2.60%	38.7	20.7	13.5	36.2
	3.60%	48.5	23.3	21.5	44.8	5.20%	37.3	21.1	14.4	34.1
GEOM	0.90%	32.7	10.5	8.6	31.7	1.30%	39.1	20.1	11.4	40.0
	1.80%	48.2	23.6	22.7	45.2	2.60%	32.2	14.5	8.9	29.4
	3.60%	54.2	25.7	24.9	52.1	5.20%	38.1	22.0	12.7	34.7

Table B8: **Structural Anomaly Detection results (ROC-AUC)** on *Cora* and *Citeseer* with varying condensation ratios. The best results are shown in **bold** and the runner-ups are shown in underline.

Dataset	Ratio ( $r$ )	Random	Herding	K-Center	GCDM	DM	DosCond	GCond	SGDD	SFGC	GEOM
<i>Citeseer</i>	0.90%	0.44	0.38	0.44	<u>0.76</u>	<u>0.76</u>	0.73	<b>0.77</b>	0.67	0.62	0.59
	1.80%	0.46	0.45	0.46	<b>0.78</b>	<b>0.78</b>	0.66	<u>0.75</u>	0.68	0.60	0.56
	3.60%	0.44	0.40	0.44	<b>0.76</b>	<b>0.76</b>	0.70	0.74	<u>0.75</u>	0.59	0.57
<i>Cora</i>	1.30%	0.56	0.59	0.62	<u>0.80</u>	<u>0.80</u>	0.79	<b>0.81</b>	0.75	0.54	0.51
	2.60%	0.50	0.65	0.67	<u>0.80</u>	<u>0.80</u>	<b>0.82</b>	0.79	<u>0.81</u>	0.53	0.53
	5.20%	0.65	0.55	0.67	<b>0.82</b>	<b>0.82</b>	<b>0.82</b>	<u>0.81</u>	0.71	0.54	0.55

294 putational complexity. We use a 2-layer Cheby with 256 hidden units and ReLU activation  
 295 function.

- 296 • **GraphSAGE [10]**: GraphSAGE is a spatial-based graph neural network that directly samples  
 297 and aggregates features from a node’s local neighborhood. In the experiments, We use a two-layer  
 298 architecture and a hidden dimension size of 256 while using a mean aggregator.
- 299 • **APPNP [16]**: APPNP leverages personalized PageRank to propagate information throughout  
 300 the graph. This method decouples the neural network used for prediction from the propagation  
 301 mechanism, enabling the use of personalized PageRank for message passing. In the experiments,  
 302 we use a 2-layer model implemented with ReLU activation and sparse dropout to condense and  
 303 evaluate.

Table B9: **Contextual Anomaly Detection results (ROC-AUC)** on *Cora* and *Citeseer* with varying condensation ratios. The best results are shown in **bold** and the runner-ups are shown in underline.

Dataset	Ratio ( $r$ )	Random	Herding	K-Center	GCDM	DM	DosCond	GCond	SGDD	SFGC	GEOM
<i>Citeseer</i>	0.90%	0.62	0.60	0.62	0.65	0.65	0.55	0.70	<b>0.74</b>	0.62	0.59
	1.80%	0.60	0.54	0.60	0.64	0.65	0.58	<b>0.68</b>	<u>0.67</u>	0.60	0.56
	3.60%	0.57	0.56	0.57	<b>0.68</b>	<b>0.68</b>	<u>0.59</u>	<b>0.68</b>	<u>0.52</u>	<u>0.59</u>	0.57
<i>Cora</i>	1.30%	0.52	0.48	<u>0.53</u>	0.52	0.52	0.45	<b>0.54</b>	0.41	<b>0.54</b>	0.51
	2.60%	0.50	0.45	<u>0.54</u>	0.54	0.54	0.56	0.55	<b>0.57</b>	0.53	0.53
	5.20%	0.56	0.58	<u>0.59</u>	0.55	0.55	<u>0.55</u>	0.57	<b>0.62</b>	0.54	0.55

- 304 • **GIN [35]:** GIN aggregates features by linearly combining the node features with those of their  
305 neighbors, achieving classification power as strong as the Weisfeiler-Lehman graph isomorphism  
306 test. We specifically applied a 3-layer GIN with a mean pooling function to compress and evaluate  
307 graph classification datasets. For the datasets *DD* and *NCII*, we use negative log-likelihood loss  
308 function for training and softmax activation in the final layer. For *ogbg-molhiv*, *ogbg-molbbbp*  
309 and *ogbg-molbace*, we use binary cross-entropy with logits and sigmoid activation in the final  
310 layer.
- 311 • **Graph Transformer [28]:** The Graph Transformer leverages the self-attention mechanism of the  
312 Transformer to capture long-range dependencies between nodes in a graph. It employs multi-head  
313 self-attention to dynamically weigh the importance of different nodes, effectively modeling  
314 complex relationships within the graph. We use a two-layer model with layer normalization and  
315 gated residual connections, following the settings outlined in [28].

## 316 B.4.2 Additional Results

317 Table B10 shows the node classification accuracy of datasets condensed by traditional core-set  
318 methods, which is backbone-free, evaluated across different backbone architectures on *Cora*.

Table B10: **Node Classification Accuracy (%)** of core-set datasets across different backbone architectures on *Cora* (2.6%).

Methods	SGC	GCN	GraphSage	APPNP	Cheby	GTrans.	MLP
Full Dataset	80.8	80.8	80.8	80.3	78.8	69.6	81.0
Herding	74.8	74.0	74.1	73.3	69.6	65.4	74.1
K-Center	72.5	72.4	71.8	71.5	63.0	64.3	72.2
Random	71.7	72.4	71.6	71.3	65.3	62.7	71.6

## 319 B.5 Settings and Additional Results of Initialization Impacts (RQ5)

### 320 B.5.1 Experimental Settings

321 The details of evaluated initialization mechanism are as follows:

- 322 • **Random Sample.** We randomly select features from nodes in the original graph that correspond  
323 to the same label, using these features to initialize the synthetic nodes.
- 324 • **Random Noise.** Consistent with prevalent dataset condensation methods, we initialize node  
325 features by sampling from a Gaussian distribution.
- 326 • **Center.** This method involves extracting features from nodes within the same label, applying the  
327 K-Means clustering algorithm to these features while treating the graph as a singular cluster and  
328 utilizing the centroid of this cluster as the initialization point for all synthetic nodes bearing the  
329 same label.
- 330 • **K-Center.** Similar to the Center initialization method, but employ the K-Means Clustering  
331 method on original nodes by dividing each class of the original graph nodes into  $n$  clusters,  
332 where  $n$  is the number of synthetic nodes per class. We select the center of these clusters as the  
333 initialization of synthetic nodes in this class.

334 • **K-Means.** Similar to the K-Center initialization method, but instead of using the centroids of  
 335 clusters to initialize the synthetic dataset, randomly select one node from each cluster to serve as  
 336 the initial state for the synthetic node.

### 337 B.5.2 Additional Results

338 The performance of different initialization mechanism on *Cora* (2.6%) and *Cora* (0.26%) are shown  
 339 in Table B11 and Table B12, respectively. It is evident that distribution matching methods are highly  
 340 sensitive to the choice of initialization, especially when the dataset is condensed to a smaller scale.  
 341 Additionally, trajectory matching methods perform poorly with random noise initialization and often  
 fail to converge.

Table B11: Performance comparison of different initialization on various methods for *Cora* (2.60%). The best results are shown in **bold**.

Methods	Random Noise	Random Sample	Center	K-Center	K-Means
GCDM	34.5	73.3	77.4	<b>78.7</b>	75.9
DM	34.5	73.7	77.7	<b>78.1</b>	75.9
DosCond	78.8	81.9	81.8	<b>82.5</b>	81.8
GCond	74.8	75.1	<b>76.3</b>	76.2	75.1
SGDD	81.7	81.8	82.6	<b>82.7</b>	82.5
SFGC	52.5	80.7	79.7	81.5	<b>81.8</b>
GEOM	-	77.9	48.3	78.8	<b>78.9</b>

Table B12: Performance comparison of different initialization on various methods for *Cora* (0.26%). The best results are shown in **bold**.

Methods	Random Noise	Random Sample	Center	K-Center	K-Means
GCDM	32.3	37.8	<b>78.7</b>	<b>78.7</b>	34.3
DM	32.2	38.4	<b>77.9</b>	<b>77.9</b>	34.2
DosCond	78.7	<b>82.4</b>	80.5	82.0	81.9
GCond	80.2	<b>81.6</b>	80.1	81.2	80.7
SGDD	82.2	82.2	<b>82.7</b>	<b>82.7</b>	81.5
SFGC	79.7	79.7	<b>79.8</b>	<b>79.8</b>	72.0
GEOM	-	49.6	51.3	51.3	<b>65.0</b>

342

## 343 B.6 Settings and Additional Results of Efficiency and Scalability (RQ6)

### 344 B.6.1 Experimental Settings

345 For a fair comparison, all the experiments are conducted on a single NVIDIA A100 GPU. Then we  
 346 report the overall condensation time (min) when achieving the best validation performance, the peak  
 347 CPU memory usage (MB) and the peak GPU memory usage (MB).

### 348 B.6.2 Additional Results

349 The detailed time and space consumption of the node-level GC methods on *ogbn-arxiv* (0.50%) and  
 350 graph-level GC methods on *ogbg-molhiv* (1 Graph/Cls) are shown in Table B13 and Table B14  
 351 respectively. For node-level methods, although trajectory matching methods (*SFGC* [42], *GEOM* [40])  
 352 may consume less time and memory due to their offline matching mechanism, the expert trajectories  
 353 generated before matching can occupy up to 764 GB of space as shown in Table B15, significantly  
 354 impacting storage requirements. Among all the graph-level GC methods, *Mirage* [9] stands out by  
 355 not relying on any GPU resources for calculation and can condense data extremely quickly, taking  
 only 1% of the time required by other methods.

Table B13: Time and memory consumption of different methods on *ogbn-arxiv* (0.50%).

Consumption	GCDM	DM	DosCond	GCond	SGDD	SFGC	GEOM
Time (min)	212.90	57.70	117.38	266.57	226.62	245.65	148.37
Acc. (%)	58.09	58.09	60.73	61.28	61.51	67.13	67.29
CPU Memory (MB)	2720.88	2708.70	5372.60	5270.70	5426.30	3075.30	3335.10
GPU Memory (MB)	2719.74	2552.63	3850.24	3850.24	8326.35	4050.12	5308.42

356

## 357 C Reproducibility and Limitations

358 **Accessibility and license.** All the datasets, algorithm implementations, and experimental settings  
 359 are publicly available in our open project (<https://github.com/RingBDStack/GC-Bench>). Our  
 360 package (codes and datasets) is licensed under the MIT License. This license permits users to freely  
 361 use, copy, modify, merge, publish, distribute, sublicense, and sell copies of the software, provided

Table B14: Time and memory consumption of different methods on *ogbg-molhiv* (1 Graph/Cls).

Consumption	DosCond	KiDD	Mirage
Time (min)	218.11	202.38	2.91
Acc. (%)	67.41	66.44	71.09
CPU Memory (MB)	2666.29	3660.79	752.22
GPU Memory (MB)	1005.98	6776.42	0.00

Table B15: Expert trajectory size (GB) for trajectory matching methods.

<i>Citeseer</i>	<i>Cora</i>	<i>ogbn-arxiv</i>	
129	152	15	
<i>Flickr</i>	<i>Reddit</i>	<i>ACM</i>	<i>DBLP</i>
21	42	312	764

362 that the original copyright notice and permission notice are included in all copies or or substantial  
 363 portions of the software. The MIT License is widely accepted for its simplicity and permissive terms,  
 364 ensuring ease of use and contribution to the codes and datasets. We bear all responsibility in case of  
 365 violation of rights, *etc.*, and confirmation of the data license.

366 **Datasets.** *Cora*, *Citeseer*, *Flickr*, *Reddit* and *DBLP* are publicly available online<sup>1</sup> with the MIT  
 367 license. *ogbn-arxiv*, *ogbg-molbase*, *ogbg-molbbbp* and *ogbg-molhiv* are released by OGB [11]  
 368 with the MIT license. *ACM* [41] is the subset hosted in [32] with the MIT license. *NCII* [29] and  
 369 *DD* [3] are available in TU Datasets [25] with the MIT license. All the datasets are consented to by  
 370 the authors for academic usage. All the datasets do not contain personally identifiable information or  
 371 offensive content.

372 **Limitations.** GC-Bench has some limitations that we aim to address in future work. Our current  
 373 benchmark is limited to a specific set of graph types and graph tasks and might not reflect the full  
 374 potential and versatility of GC methods. We hope to implement more GC algorithms for various tasks  
 375 (e.g. subgraph classification, community detection) on more types of graphs (e.g., dynamic graph,  
 376 directed graph). Besides, due to resource constraints and the availability of implementations, we  
 377 could not include some of the latest GC algorithms in our benchmark. We will continuously update  
 378 our repository to keep track of the latest advances in the field. We are also open to any suggestions  
 379 and contributions that will improve the usability and effectiveness of our benchmark, ensuring it  
 380 remains a valuable resource for the IGL research community.

## 381 References

- 382 [1] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs  
 383 with fast localized spectral filtering. *arXiv: Learning, arXiv: Learning*, 2016.
- 384 [2] Kaize Ding, Jundong Li, Rohit Bhanushali, and Huan Liu. *Deep Anomaly Detection on Attributed Networks*,  
 385 page 594–602. May 2019.
- 386 [3] Paul D Dobson and Andrew J Doig. Distinguishing enzyme structures from non-enzymes without  
 387 alignments. *Journal of molecular biology*, 2003.
- 388 [4] Junfeng Fang, Xinglin Li, Yongduo Sui, Yuan Gao, Guibin Zhang, Kun Wang, Xiang Wang, and Xiangnan  
 389 He. Exgc: Bridging efficiency and explainability in graph condensation. *arXiv preprint arXiv:2402.05962*,  
 390 2024.
- 391 [5] Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with pytorch geometric. *arXiv*  
 392 *preprint arXiv:1903.02428*, 2019.
- 393 [6] Tao-yang Fu, Wang-Chien Lee, and Zhen Lei. Hin2vec: Explore meta-paths in heterogeneous information  
 394 networks for representation learning. In *CIKM*, pages 1797–1806, 2017.
- 395 [7] Jian Gao and Jianshe Wu. Multiple sparse graphs condensation. *Knowledge-Based Systems*, 278:110904,  
 396 2023.
- 397 [8] Xinyi Gao, Tong Chen, Yilong Zang, Wentao Zhang, Quoc Viet Hung Nguyen, Kai Zheng, and Hongzhi  
 398 Yin. Graph condensation for inductive node representation learning. In *ICDE*, 2024.
- 399 [9] Mridul Gupta, Sahil Manchanda, Sayan Ranu, and Hariprasad Kodamana. Mirage: Model-agnostic graph  
 400 distillation for graph classification. In *ICLR*, 2024.
- 401 [10] William L. Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs.  
 402 In *NeurIPS*, 2017.

<sup>1</sup>[https://github.com/pyg-team/pytorch\\_geometric](https://github.com/pyg-team/pytorch_geometric)



- 403 [11] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and  
404 Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs. In *NeurIPS*, 2020.
- 405 [12] Wei Jin, Xianfeng Tang, Haoming Jiang, Zheng Li, Danqing Zhang, Jiliang Tang, and Bing Yin. Condensing  
406 graphs via one-step gradient matching. In *SIGKDD*, 2022.
- 407 [13] Wei Jin, Lingxiao Zhao, Shi-Chang Zhang, Yozen Liu, Jiliang Tang, and Neil Shah. Graph condensation  
408 for graph neural networks. In *ICLR*, 2021.
- 409 [14] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks.  
410 *arXiv preprint arXiv:1609.02907*, 2016.
- 411 [15] Thomas N Kipf and Max Welling. Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308*,  
412 2016.
- 413 [16] Johannes Klicpera, Aleksandar Bojchevski, and Stephan Günnemann. Predict then propagate: Graph  
414 neural networks meet personalized pagerank. In *ICLR*, 2018.
- 415 [17] Xinglin Li, Kun Wang, Hanhui Deng, Yuxuan Liang, and Di Wu. Attend who is weak: Enhancing graph  
416 condensation via cross-free adversarial training. *arXiv preprint arXiv:2311.15772*, 2023.
- 417 [18] Mengyang Liu, Shanchuan Li, Xinshi Chen, and Le Song. Graph condensation via receptive field  
418 distribution matching. *arXiv preprint arXiv:2206.13697*, 2022.
- 419 [19] Yang Liu, Deyu Bo, and Chuan Shi. Graph condensation via eigenbasis matching. In *ICML*, 2024.
- 420 [20] Yilun Liu, Ruihong Qiu, and Zi Huang. Cat: Balanced continual graph learning with graph condensation.  
421 In *ICDM*, pages 1157–1162. IEEE, 2023.
- 422 [21] Yilun Liu, Ruihong Qiu, and Zi Huang. Gcondenser: Benchmarking graph condensation. *arXiv preprint*  
423 *arXiv:2405.14246*, 2024.
- 424 [22] Yilun Liu, Ruihong Qiu, Yanran Tang, Hongzhi Yin, and Zi Huang. Puma: Efficient continual graph  
425 learning with graph condensation. *arXiv preprint arXiv:2312.14439*, 2023.
- 426 [23] Qingsong Lv, Ming Ding, Qiang Liu, Yuxiang Chen, Wenzheng Feng, Siming He, Chang Zhou, Jianguo  
427 Jiang, Yuxiao Dong, and Jie Tang. Are we really making much progress?: Revisiting, benchmarking and  
428 refining heterogeneous graph neural networks. In *SIGKDD*, 2021.
- 429 [24] Runze Mao, Wenqi Fan, and Qing Li. Gcare: Mitigating subgroup unfairness in graph condensation  
430 through adversarial regularization. *Applied Sciences*, 13(16):9166, 2023.
- 431 [25] ChristopherJ. Morris, NilsM. Kriege, Franka Bause, Kristian Kersting, Petra Mutzel, and Marion Neumann.  
432 Tudataset: A collection of benchmark datasets for learning with graphs. *arXiv: Learning, arXiv: Learning*,  
433 Jul 2020.
- 434 [26] Qiyang Pan, Ruofan Wu, Tengfei Liu, Tianyi Zhang, Yifei Zhu, and Weiqiang Wang. Fedgkd: Unleashing  
435 the power of collaboration in federated graph neural networks. *arXiv preprint arXiv:2309.09517*, 2023.
- 436 [27] Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach.  
437 In *ICLR*, 2018.
- 438 [28] Yunsheng Shi, Zhengjie Huang, Shikun Feng, Hui Zhong, Wenjin Wang, and Yu Sun. Masked label  
439 prediction: Unified message passing model for semi-supervised classification, 2021.
- 440 [29] Nikil Wale, Ian A Watson, and George Karypis. Comparison of descriptor spaces for chemical compound  
441 retrieval and classification. *Knowledge and Information Systems*, 2008.
- 442 [30] Chun Wang, Shirui Pan, Ruiqi Hu, Guodong Long, Jing Jiang, and Chengqi Zhang. Attributed graph  
443 clustering: A deep attentional embedding approach. In *IJCAI*, 2019.
- 444 [31] Lin Wang, Wenqi Fan, Jiatong Li, Yao Ma, and Qing Li. Fast graph condensation with structure-based  
445 neural tangent kernel. In *The Web Conference*, 2024.
- 446 [32] Xiao Wang, Houye Ji, Chuan Shi, Bai Wang, Yanfang Ye, Peng Cui, and Philip S Yu. Heterogeneous graph  
447 attention network. In *The Web Conference*, 2019.
- 448 [33] Max Welling. Herding dynamical weights to learn. In *ICML*, pages 1121–1128, 2009.

- 449 [34] Felix Wu, Tianyi Zhang, AmauriH. Souza, Christopher Fifty, Tao Yu, and KilianQ. Weinberger. Simplifying  
450 graph convolutional networks. *arXiv: Learning, arXiv: Learning*, 2019.
- 451 [35] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks?  
452 In *ICLR*, 2018.
- 453 [36] Zhe Xu, Yuzhong Chen, Menghai Pan, Huiyuan Chen, Mahashweta Das, Hao Yang, and Hanghang Tong.  
454 Kernel ridge regression-based graph dataset distillation. In *SIGKDD*, pages 2850–2861, 2023.
- 455 [37] Beining Yang, Kai Wang, Qingyun Sun, Cheng Ji, Xingcheng Fu, Hao Tang, Yang You, and Jianxin Li.  
456 Does graph distillation see like vision dataset counterpart? In *NeurIPS*, 2023.
- 457 [38] Hanqing Zeng, Hongkuan Zhou, Ajitesh Srivastava, Rajgopal Kannan, and Viktor K. Prasanna. Graphsaint:  
458 Graph sampling based inductive learning method. In *ICLR*, 2020.
- 459 [39] Tianle Zhang, Yuchen Zhang, Kun Wang, Kai Wang, Beining Yang, Kaipeng Zhang, Wenqi Shao, Ping  
460 Liu, Joey Tianyi Zhou, and Yang You. Two trades is not baffled: Condense graph via crafting rational  
461 gradient matching. *arXiv preprint arXiv:2402.04924*, 2024.
- 462 [40] Yuchen Zhang, Tianle Zhang, Kai Wang, Ziyao Guo, Yuxuan Liang, Xavier Bresson, Wei Jin, and Yang  
463 You. Navigating complexity: Toward lossless graph condensation via expanding window matching. *CoRR*,  
464 abs/2402.05011, 2024.
- 465 [41] Jianan Zhao, Xiao Wang, Chuan Shi, Zekuan Liu, and Yanfang Ye. Network schema preserving heteroge-  
466 neous information network embedding. In *IJCAI*, 2020.
- 467 [42] Xin Zheng, Miao Zhang, Chunyang Chen, Quoc Viet Hung Nguyen, Xingquan Zhu, and Shirui Pan.  
468 Structure-free graph condensation: From large-scale graphs to condensed graph-free data. In *NeurIPS*,  
469 2023.