

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

Qingyun Sun^{1*}, Ziyang Chen^{1*}, Beining Yang², Cheng Ji¹, Xingcheng Fu³, Sheng Zhou⁴,
Hao Peng¹, Jianxin Li[#], Philip S. Yu⁵

¹ Beihang University, ² University of Edinburgh, ³ Guangxi Normal University, ⁴ Zhejiang University, ⁵ University of Illinois, Chicago

Overview of GC-Bench

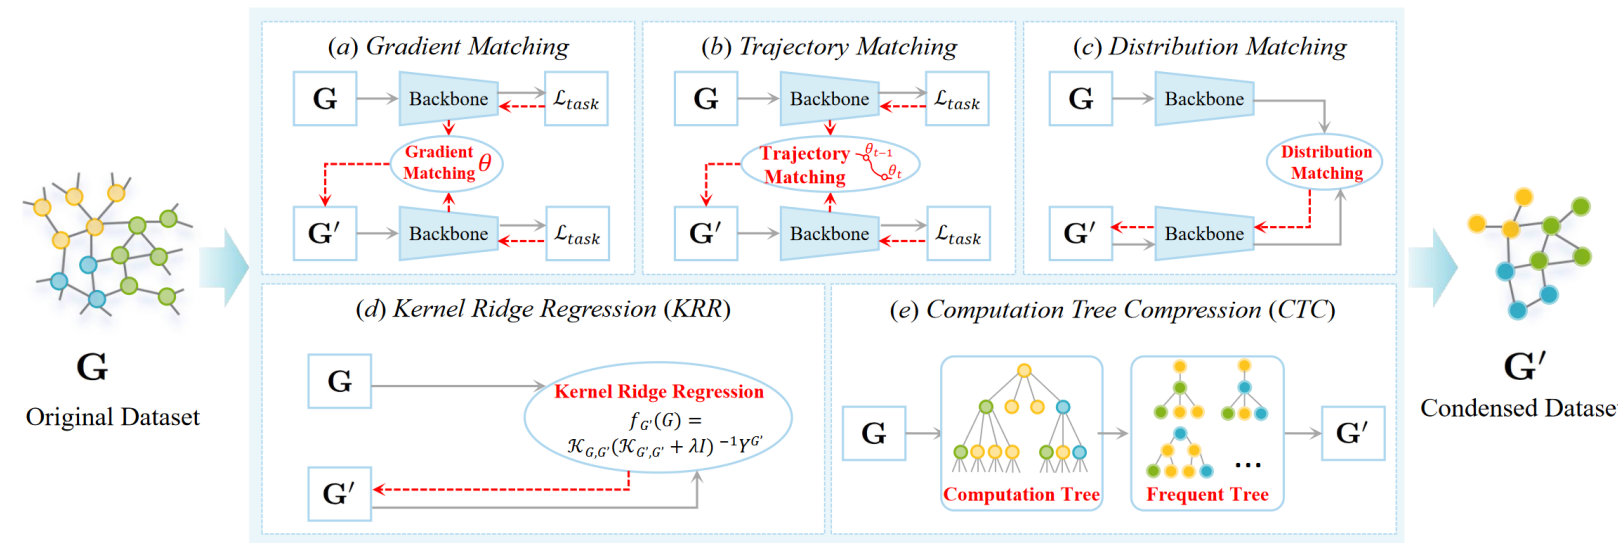


Figure 1: The GC methods can be broadly divided into two categories: The first category depends on the backbone model, refining the condensed graph by aligning it with the backbone’s *gradients* (a), *trajectories* (b), and output *distributions* (c) trained on both original and condensed graphs. The second category, independent of the backbone, optimizes the graph by matching its distribution with that of the original graph data (d) or by identifying frequently co-occurring computation trees (e).

Graph condensation (GC) works are in a promising direction and show tremendous performance by condensing the large-scale datasets’ valuable information. We introduce GC-Bench, an open and unified benchmark to systematically evaluate existing graph condensation methods focusing on the following aspects. **Effectiveness**: the progress in GC, and the impact of structure and initialization on GC; **Transferability**: the transferability of GC methods across backbone architectures and downstream tasks; **Efficiency**: the time and space efficiency of GC methods. The contributions of GC-Bench are as follows:

- **Comprehensive benchmark and Key findings.** GC-Bench systematically integrated 12 representative and competitive GC methods on both node-level and graph-level by unified condensation and evaluation, giving multi-dimensional analysis in terms of effectiveness, transferability, and efficiency
- **Open-sourced benchmark library and future directions.** GC-Bench is open-sourced and easy to extend to new methods and datasets, which can help identify directions for further exploration and facilitate future endeavors

Table 1: An overview of GC-Bench

Methods	
Traditional core-set methods	Random, Herding [36], K-Center [30]
Gradient matching	GCond [15], DosCond [14], SGDD [42]
Trajectory matching	SFGC [47], GEOM [45]
Distribution matching	GCDD [20], DM [24, 22]
Kernel Ridge Regression	KiDD [41]
Computation Tree Compression	Mirage [11]
Datasets	
Homogeneous datasets	Cora [16], Citeseer [16], ogbn-arxiv [13], Flickr [43], Reddit [12]
Heterogeneous datasets	ACM [46], DBLP [7]
Graph-level datasets	NCII [32], DD [4], ogbg-molbace [13], ogbg-molbbpp [13], ogbg-molhiv [13]
Downstream Tasks	
Node-level task	Node classification, Link prediction, Anomaly detection
Graph-level task	Graph classification
Evaluations	
Effectiveness	Performance under different condensation ratios, Impact of structural properties, Impact of initialization mechanism
Transferability	Different downstream tasks, Different backbone model architectures
Efficiency	Time and memory consumption



Paper: <https://arxiv.org/abs/2407.00615>

Code: <https://github.com/RingBDStack/GC-Bench>

Research Questions and Experimental Results

- RQ1: How much progress has been made by existing GC methods?
- **Key Takeaways 1:** Current node-level GC methods can achieve nearly lossless condensation performance. There is still a significant gap between graph-level GC and whole dataset training.
 - **Key Takeaways 2:** A large condensation ratio does not necessarily lead to better performance with current methods.

Table 2: Node classification accuracy (%) (mean±std) across datasets with varying condensation ratios r . \mathcal{H} denotes the homophily ratio [29]. The best results are shown in bold and the runner-ups are shown in underlined. Red color highlights entries that exceed the whole dataset performance.

Dataset	Ratio(r)	Traditional Core-set Methods					Distribution			Gradient			Trajectory			Whole Dataset
		Random	Herding	K-Center	GCDD	DM	DosCond	GCond	SGDD	SFGC	GEOM	GCDD	DM	GCDD		
Cora (H=0.81)	0.26%	31.6±1.2	48.6±1.4	48.6±1.4	39.0±0.4	35.6±3.9	78.7±1.3	79.8±0.8	78.6±2.7	78.8±1.2	50.1±1.2	80.8±0.3	80.8±0.3	80.8±0.3	80.8±0.3	
	0.2%	47.1±1.7	46.6±0.4	44.7±1.7	42.3±0.7	38.5±1.4	81.1±0.1	80.2±0.7	79.2±1.8	79.2±1.8	70.5±0.9	80.8±0.3	80.8±0.3	80.8±0.3	80.8±0.3	
	1.30%	62.3±1.0	69.9±0.8	62.0±1.3	63.4±0.1	63.2±0.5	81.2±0.4	80.2±2.2	81.6±0.9	79.4±0.7	78.5±1.9	80.8±0.3	80.8±0.3	80.8±0.3	80.8±0.3	
	2.60%	72.4±0.5	74.2±0.6	73.5±0.7	73.4±0.2	72.6±0.7	79.6±0.2	80.5±0.3	81.8±0.2	81.7±0.0	76.1±4.7	80.8±0.3	80.8±0.3	80.8±0.3	80.8±0.3	
	3.90%	74.6±0.6	76.0±0.4	77.4±0.3	76.4±0.3	77.3±0.1	78.9±0.1	79.8±0.4	82.1±0.1	81.8±0.0	78.3±2.0	80.8±0.3	80.8±0.3	80.8±0.3	80.8±0.3	
	5.20%	77.1±0.6	76.7±0.4	76.5±0.6	78.4±0.0	77.9±0.2	79.7±0.6	77.8±0.6	82.4±1.4	81.4±0.0	80.4±0.8	80.8±0.3	80.8±0.3	80.8±0.3	80.8±0.3	

Table 3: Graph classification performance on GIN (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 C/Is	Ratio(r)	Traditional Core-set methods			Gradient		KRR		CTC		Whole Dataset
			Random	Herding	K-Center	DosCond	KiDD	Mirage	Mirage			
NCII Acc. (%)	1	0.06%	50.90±2.10	51.90±1.60	51.90±1.60	49.20±1.10	61.40±0.50	50.80±2.20	51.30±1.10	51.30±1.10	80.1±1.8	
	5	0.24%	52.10±1.00	60.50±2.40	47.00±1.10	51.10±0.80	63.20±0.20	51.30±1.10	51.70±1.40	51.70±1.40	80.1±1.8	
	10	0.49%	55.60±1.90	61.30±1.50	49.40±1.80	50.30±1.30	64.20±0.10	51.70±1.40	52.10±2.20	52.10±2.20	80.1±1.8	
	20	0.97%	58.70±1.40	60.90±1.90	55.20±1.60	50.30±1.30	60.90±0.70	52.10±2.20	52.10±2.20	52.10±2.20	80.1±1.8	
	50	2.43%	61.10±1.20	59.00±1.50	62.70±1.90	50.30±1.30	65.40±0.60	52.40±2.20	52.40±2.20	52.40±2.20	80.1±1.8	
DD Acc. (%)	1	0.21%	49.70±1.10	58.80±1.10	58.80±1.10	46.30±2.20	71.30±1.10	74.00±2.10	74.00±2.10	74.00±2.10	70.1±2.2	
	5	1.06%	40.80±4.30	58.70±5.80	51.30±5.30	57.50±5.00	70.90±1.10	-	-	-	70.1±2.2	
	10	2.12%	63.10±5.20	64.10±5.80	53.40±3.10	46.30±8.50	71.50±0.50	-	-	-	70.1±2.2	
	20	4.25%	56.40±4.30	67.00±2.60	58.50±5.70	40.70±6.00	71.20±0.00	-	-	-	70.1±2.2	
	50	10.62%	58.90±6.30	68.40±4.00	62.30±2.50	44.00±6.70	71.80±1.00	-	-	-	70.1±2.2	

RQ2: How do the potential flaws of the structure affect the GC performance?

- **Key Takeaways 3:** Existing GC methods primarily address single graph data. There is significant room for improvement in preserving complex structural properties.

Table B4: Homophily ratio comparison of different condensed datasets

Dataset	Ratio(r)	GCDD	DM	DosCond	GCond	SGDD	
Cora	0.81	1.30% 0.11 5.20%	0.76 0.74 0.21	0.88 0.16 0.15	0.20 0.64 0.62	0.19 0.19 0.15	
	Citeseer	0.74	0.90% 1.80% 3.60%	0.16 0.08 1.00	0.75 0.30 0.34	0.19 0.36 0.15	0.14 0.19 0.15
		Flickr	0.24	0.05% 0.50% 1.00%	0.28 0.29 0.36	0.29 0.22 0.18	0.25 0.08 0.06



- RQ3: Can the condensed graphs be transferred to different types of tasks?
- **Key Takeaways 4:** All condensed datasets struggle to perform well outside the context of the specific tasks for which they were condensed.

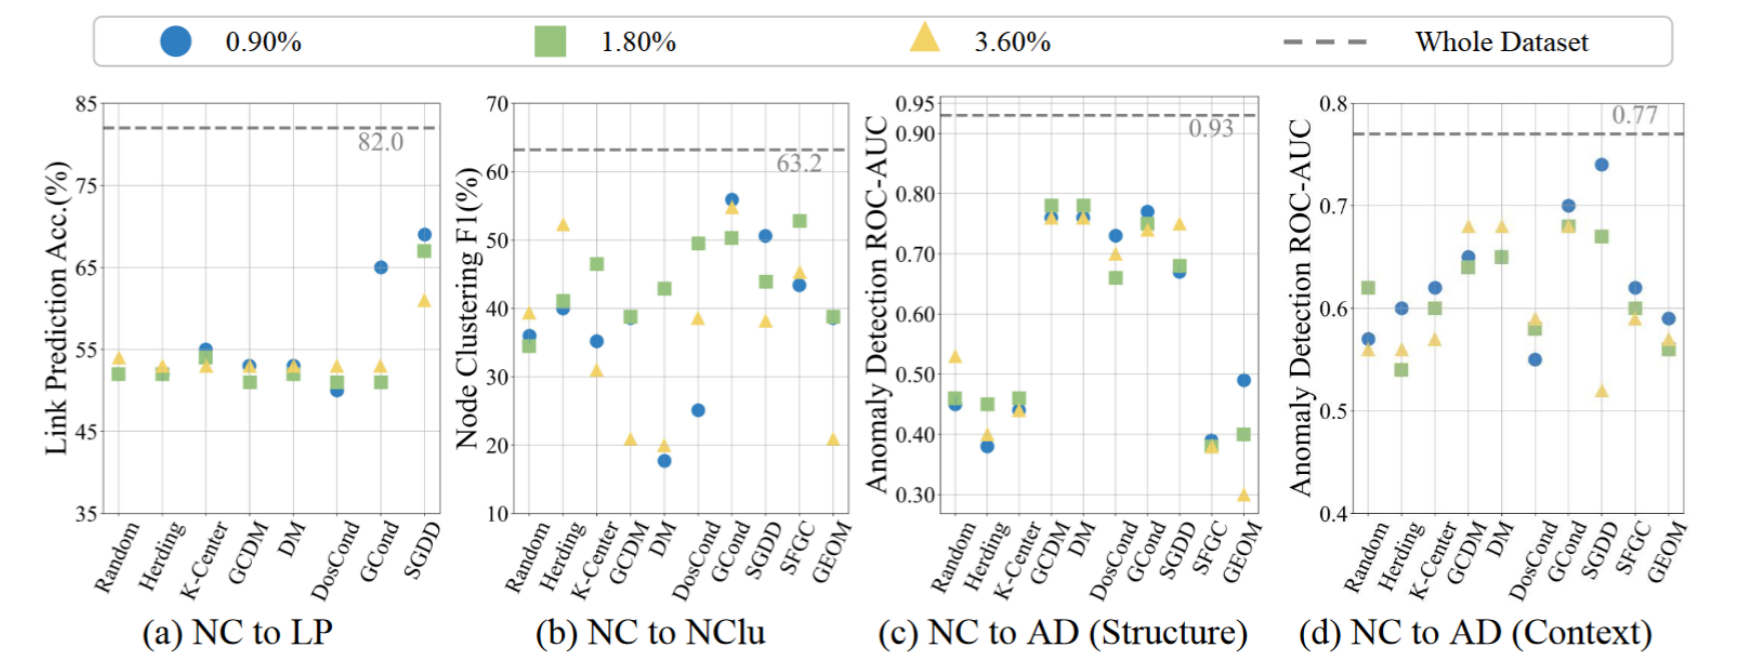


Figure 2: Cross-task performance on Citeseer. For all downstream tasks, the models are trained solely using data of graphs condensed by node classification. For anomaly detection (c, d), structural and contextual anomalies [3] are injected into both the condensed graph and the original graph.

- RQ4: How does the backbone used for condensation affect the performance?

- **Key Takeaways 5:** Current GC methods exhibit significant performance variability when transferred to different backbone architectures. Involving the entire training process potentially may lead to encoding backbone-specific details in the condensed datasets.
- **Key Takeaways 6:** Despite their strong performance in general graph learning tasks, transformers surprisingly yield suboptimal results in graph condensation.

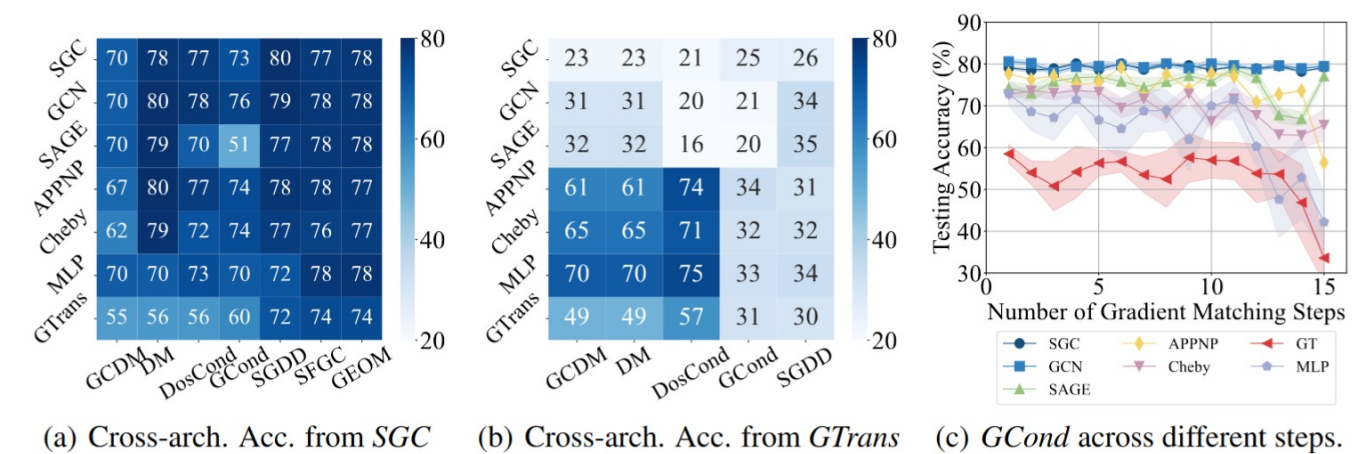


Figure 3: Cross-architecture performance. Using SGC and Graph Transformer (GTrans) to condense Cora with a 2.6% ratio, we then test the accuracy on various downstream architectures (a, b). Furthermore, we evaluate the influence of gradient matching steps on GCond (c).

- RQ5: How does the initialization mechanism affect the performance of GC?

- **Key Takeaways 7:** Different datasets have their preferred initialization methods for optimal performance even for the same GC method.
- **Key Takeaways 8:** The initialization mechanism primarily affects the convergence speed with little impact on the final performance. The smaller the condensed graph, the greater the influence of different initialization strategies on the convergence speed.

- RQ6: How efficient are these GC methods in terms of time and space?

- **Key Takeaways 9:** GC methods that rely on backbones and full-scale data training have large time and space consumption.

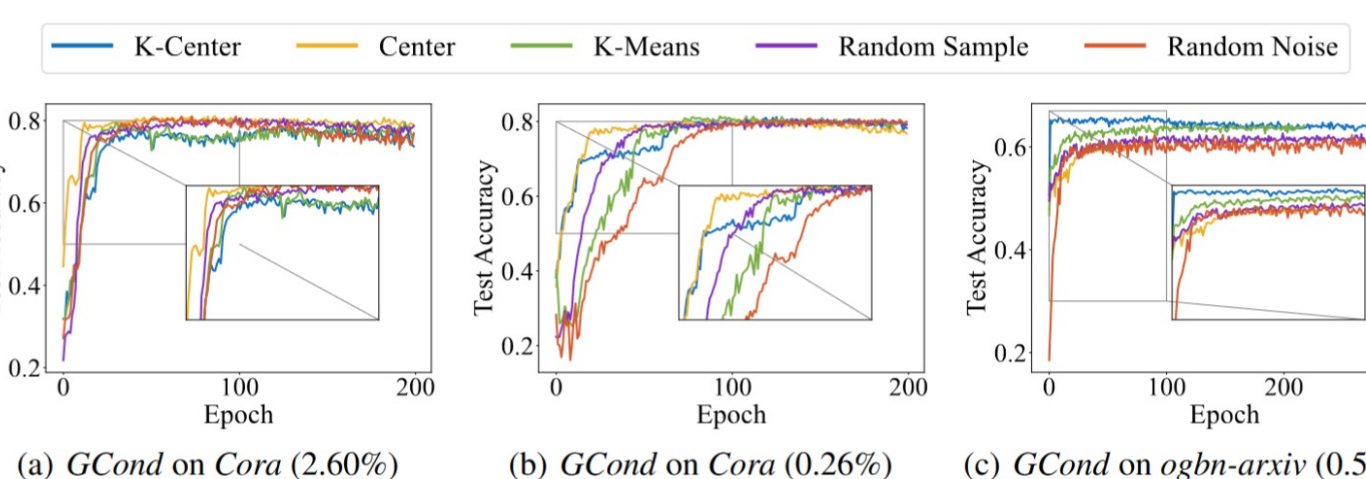


Figure 4: The impact of initialization under different condensation ratios (a, b) and the impact across different datasets Cora (a, b) and ogbn-arxiv (c).

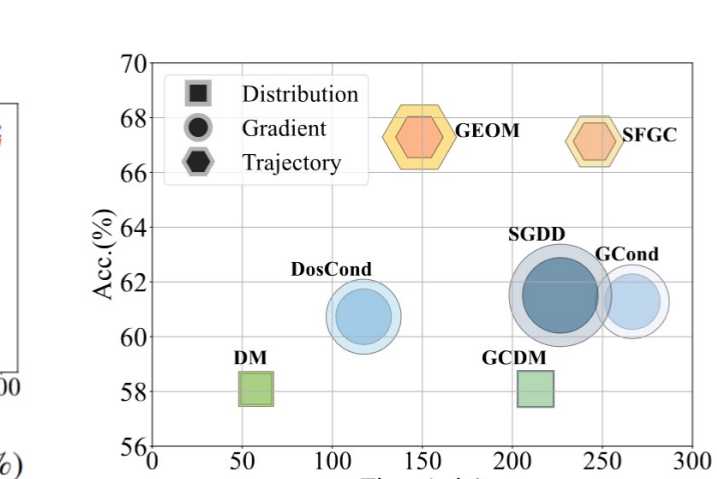


Figure 5: Time and memory consumption of different methods on ogbn-arxiv (0.50%).