Inference-friendly Graph Compression for Graph Neural Networks

Yangxin Fan
Case Western Reserve University
Cleveland, Ohio, USA
yxf451@case.edu

Haolai Che
Case Western Reserve University
Cleveland, Ohio, USA
hxc859@case.edu

Yinghui Wu
Case Western Reserve University
Cleveland, Ohio, USA
yxw1650@case.edu

ABSTRACT

Graph Neural Networks (GNNs) have demonstrated promising performance in graph analysis. Nevertheless, the inference process of GNNs remains costly, hindering their applications for large graphs. This paper proposes *inference-friendly graph compression* (IFGC), a graph compression scheme to accelerate GNNs inference. Given a graph G and a GNN M, an IFGC computes a small compressed graph G_c, to best preserve the inference results of M over G, such that the result can be directly inferred by accessing G_c with no or little decompression cost. (1) We characterize IFGC with a class of inference equivalence relation. The relation captures the node pairs in G that are not distinguishable for GNN inference. (2) We introduce three practical specifications of IFGC for representative GNNs: structural preserving compression (SPGC), which computes G_c that can be directly processed by GNN inference without decompression; (α, r)-compression, that allows for a configurable trade-off between compression ratio and inference quality, and anchored compression that preserves inference results for specific nodes of interest. For each scheme, we introduce compression and inference algorithms with guarantees of efficiency and quality of the inferred results. We conduct extensive experiments on diverse sets of large-scale graphs, which verifies the effectiveness and efficiency of our graph compression approaches.

PVLDB Reference Format:
Yangxin Fan, Haolai Che, and Yinghui Wu. Inference-friendly Graph Compression for Graph Neural Networks. PVLDB, 18(9): 3203 - 3215, 2025.
doi:10.14778/3746405.3746438

PVLDB Artifact Availability:
The source code, data, and/or other artifacts have been made available at https://github.com/Yangxin666/SPGC (Last accessed: June 12, 2025).

1 INTRODUCTION

Graph Neural Networks (GNNs) have shown promising performance in various analytical tasks such as node classification [30] and link prediction [54]. In general, a GNN $\mathcal{M}$ converts an input graph G (as a pair (X, A) of node feature matrix X and an adjacency matrix A), to a vector representation ("embeddings") $\mathcal{M}(G)$ via multiple layers. For each node, each layer applies a same "node update function" to uniformly update its embedding as a weighted aggregation of embeddings from its neighbors, subsequently transforming it towards an output embedding. The training of $\mathcal{M}$ is to optimize its model parameters ("weights") and obtain a proper update function to make it best fits a set of training data in a training graph. Given an input (test) graph G, the *inference* of $\mathcal{M}$ applies the node update function to generate output embeddings $\mathcal{M}(G)$ (a matrix of node embeddings). $\mathcal{M}(G)$ can be post-processed to task-specific output, such as class labels for node classification.

Despite their promising performances, GNNs incur expensive inference process when G is large [12, 52, 57]. The emerging need for large-scale testing, fine-tuning and benchmarking of graph learning models, require fast inferences of GNNs under various configurations. Consider the following scenarios.

(1) *Inference in large networks.* For a graph G with $|V|$ nodes and $|E|$ edges (where V and E refers to its node and edge set), with on average F features per node, an L-layered GNN M may typically take $O(L|E|dF^2 + L|V|F^2)$ time [12] (as summarized in Table 3). This can be prohibitively expensive for large real-world graphs.

(2) *Real-time Inference.* GNNs have been developed for traffic analysis [41], social recommendation [17], forecasting [2, 18, 29], computer vision [45], and edge devices [57]. Such scenarios require real-time response at milliseconds [57]. In such cases, even a linear time inference of "small" GNNs (when F and L are small constants) may still be infeasible for large graphs (when $|V|$ and $|E|$ are large).

(3) *Fine-tuning & Benchmarking.* Fine-tuning and testing pre-trained GNNs to adapt them for various domain-specific tasks is a routine process in GNN-based data analysis for *e.g.,* materials sciences, biomedicine, social science, and geosciences [24, 42, 47, 56]. Inference tests of large pool of "candidate" GNNs over domain-specific graph data (such as knowledge graphs) is a cornerstone in such context. Fast GNN inference can accelerate large-scale domain-specific testing and benchmarking.

Several approaches have been developed to accelerate GNN inference, by simplifying model architecture [40], (learning) to optimize inference process [46], or data sampling [13]. These methods typically works with specific GNN $\mathcal{M}$, requires prior knowledge of its internals (*e.g.,* model parameter values), and may incur new computation overhead each time a different GNN $\mathcal{M}$ is specified.

Compressing graphs for GNN **inference**. Unlike prior "model-specific" approaches, we propose a model-agnostic, "*once-for-all*" graph compression scheme to accelerate GNN inference, for a *set* of GNNs. Consider a set of GNNs $\mathbb{M}$ (a GNN "class") with the same form of inference, which apply the same "type" of the node update function but only differs in model weights (see Example 1). The inference of a GNN M over G can be characterized as an "inference query" [6, 22], which invokes the inference of $\mathcal{M}$ to compute $\mathcal{M}(G)$.

Proceedings of the VLDB Endowment, Vol. 18, No. 9 ISSN 2150-8097.
doi:10.14778/3746405.3746438

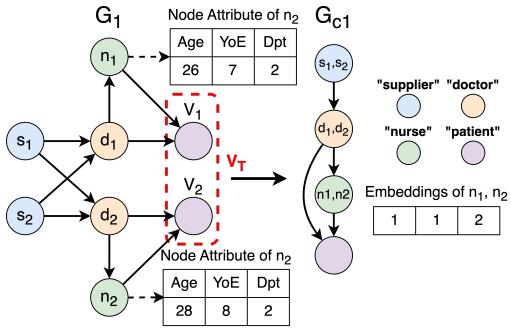

(a) Social role classification: a hospital network (G_1) can be "compressed" by merging nodes with equivalent social roles for testing a GCN-based classifier (adapted from [10]).

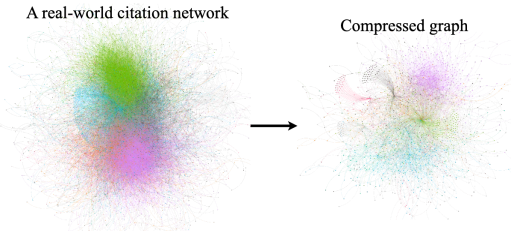

(b) Visualization of a fraction of real-world citation network [27] ($|G| = 1,335,586$) and its compressed counterpart ($|G_c| = 148,887$) for a GraphSAGE-based node classification. 88.6% of nodes and edges are compressed, reducing inference cost by 92.0%, achieving 12.5 times speed-up with up to 6.7% loss of accuracy.

Figure 1: Compression Scheme to scale node classification.

We advocate an *"inference-friendly" graph compression* scheme for GNN inference at scale. Given $\mathbb{M}$ and a large graph G,

- It uses a compression function $\mathcal{C}$ to compute a smaller counterpart G_c of G "once-for-all", for any GNN$\mathcal{M} \in \mathbb{M}$;
- For any inference query that requests $\mathcal{M}(G)$ for a specific GNN $\mathcal{M} \in \mathbb{M}$, it performs an inference directly over G_c instead of G to compute $\mathcal{M}(G_c)$, with a reduced time cost, such that $\mathcal{M}(G_c)$ (approximately) equals $\mathcal{M}(G)$.

Such a compression is desirable: (1) It readily reduces the cost for any single inference query that computes $\mathcal{M}(G)$; (2) An inference query often does not require the entire output $\mathcal{M}(G)$ but only a fraction $\mathcal{M}(G, V_T) \subseteq \mathcal{M}(G)$ of a specified test (node) set V_T of interests; (3) Multiple inference queries can be posed to request output from different GNNs in $\mathbb{M}$ in G. For any workload with inference queries that specify any GNN $\mathcal{M} \in \mathbb{M}$ and any V_T from G, one only need to compute G_c once, to reduce the total inference cost of the workload. These benefits applications in tests over large-scale graphs, real-time inference and benchmarking, as aforementioned.

While desirable, *is such a compression scheme doable?* We illustrate a case in the following example.

Example 1: Consider a 3-layer Vanilla GNN $\mathcal{M}$ [44] as a node classifier that assigns role labels {supplier, doctor, nurse, patient} in a social healthcare network G_1 (illustrated in Fig. 1). Each node in G_1 has attributes such as role, age group, department, etc. To infer the roles of test nodes $V_T = \{v_1, v_2\}$, an inference process of $\mathcal{M}$ starts by propagating a node feature matrix X with a node update function M_v. Via a 3-layer forward message passing, the embeddings of v_1 and v_2 are obtained, quantifying the likelihood of them being assigned to one of the labels. As the probability of "patient" is the highest for both, it infers both labels as "patient".

We take a closer look at the update function M_v at layer k:

$$X_v^k = \sigma(\Theta \cdot \sum_{u \in N(v)} X_u^{k-1})$$

where X_v^{k-1} (resp. X_v^k) is the embedding of a node v at the $(k-1)$-th (resp. k-th) layer; σ is an activation function, $N(v)$ refers to the neighbors of node v, and Θ refers to the learned weight matrix (same across all layers in the GNN).

If the input features of a pair of nodes x_1 and x_2 are the same (with x ranges over s, d, n and v), then the embedding of x_1 and x_2 will be the same during the inference computation, as long as the above the node update function is applied for any fixed model weights Θ and any fixed total number of layers. That is, x_1 and x_2 are "indistinguishable" for the inference of any 3-layered GNN $\mathcal{M}$ that adopts the above node update function M with the same aggregator AGG, regardless of how its Θ changes.

Note that feature equivalence does not necessarily mean that x_1 and x_2 have exactly the same attribute values as input. For example, while n_1 and n_2 refer to a 26 years old and a 28 years old nurse, respectively, their ages, years of experience (YoE), and department (Dpt) fall in the same group via (categorical or one-hot) encoding. Hence, they have the same input feature.

By "recursively" merging all such node pairs into a "group node" that are connected to neighbors that are also indistinguishable groups (*e.g.*, $[s] = \{s_1, s_2\}$, $[d] = \{d_1, d_2\}$, $[n] = \{n_1, n_2\}$, and $[v] = \{v_1, v_2\}$), a smaller graph G_c can be obtained. An inference directly over G_c can yield the same output for the test nodes v_1 and v_2 for $\mathcal{M}$ without decompression. To see this, one just need to "recover" their original value with a constant factor 2 (their original degrees) as auxiliary information at query time, for each layer:

$$X_{v_1}^k = X_{v_2}^k = 2 \times X_{[v]}^k = 2 \times \sigma(\Theta \cdot X_{[n]}^{k-1})$$

Such aggregated neighborhood information (*e.g.,* degrees, edge weights/attentions, or hyper-parameters) can be readily "remembered" at compression time, and be retrieved in constant time. This indicates an overall cheaper inference cost, and an exact query-time restore of the original embedding, for any node in G.

Better still, we only need to compute G_c "once-for-all", to reduce the unnecessary inference cost for any set of inference queries that specify a 3-layer GNN $\mathcal{M}$ (regardless of their weights Θ) that uses the same node update function M, and for any V_T in G. □

The above example verifies the possibility of a graph compression scheme by finding and merging node pairs that are indistinguishable for the inference process of GNNs. Our study verifies that real-world graphs are indeed highly compressible with such structures, and if compressed, well preserve inference output with no or small loss of accuracy, and meanwhile significantly reduce unnecessary inference computation (see Fig. 1 (b)).

Contributions. Our main contributions are as follows.

(1) We formally introduce *inference-friendly graph compression scheme* (IFGC), a as a general scheme to scale GNN inference to large graphs. We characterize IFGC with an *inference equivalence*

Methods	Category	LB	MA	IFGS	Compression Cost
Dspar [36]	S	×	✓	✓	$O(\frac{\lvert V\rvert log\lvert V\rvert}{\epsilon^2})$
AdaptiveGCN [33]	S	✓	×	×	n/a
NeuralSparse [55]	S	✓	✓	×	$O(q\lvert E\rvert)$
SCAL [28]	C	✓	✓	×	n/a
FGC [31]	C	✓	✓	×	$O(\lvert V\rvert^2\lvert V_c\rvert)$
SPGC (Ours)	C	×	✓	✓	$O(\lvert V\rvert + \lvert E\rvert)$
(α, r)-SPGC (Ours)	C	×	✓	✓	$O(\lvert V\rvert\lvert N_r\rvert + \lvert E\rvert)$
ASPGC (Ours)	C	×	✓	✓	$O(\lvert G_L\rvert)$

Table 1: Graph compression to accelerate GNN inference. S: Sparsification, C: Coarsening, LB: Learning-Based, MA: Model-Agnostic, IP: Inference-friendly w. guarantee, ϵ: a constant that controls approximation error, q: # of visits of the neighbors per node. N_r: the largest r-hop neighbor set for a node in G. G_L: the subgraph of G induced by L-hop of anchored node set V_A for L-layered GNNs. .

relation, which captures the nodes with embeddings that are indistinguishable for the inference process using the same type of node update function. We then introduce a sufficient condition for the existence of IFGC, which specifies G_c as the quotient graph of G induced by the inference equivalence relation.

(2) We specify IFGC for representative GNNs classes. We first introduce structural preserving compression (SPGC), that enforces node embedding equivalence and neighborhood connectivity. We show it computes a compressed graph G_c in $O(|E| \log |V|)$ time, which can be directly processed by the inference process to retrieve the original results *without* decompression. We further justify SPGC by showing that it can produce a unique, smallest G_c up to graph isomorphism among compressed graphs. We also show that SPGC preserves the discriminability of the GNNs.

(3) We further introduce two *configurable* variants of IFGC to allow flexible trade-off between compression ratio and the quality of inference output. (a) The (α, r)-SPGC groups nodes with similar features (determined by a threshold α), that also have similar counterparts within their r-hop neighbors. (b) The *anchored*-SPGC (ASPGC) adapts SPGC to an "anchored" node set of user's interests and preserve the inference results for such nodes only rather than the entire node set. For both variants, we introduce efficient compression and inference algorithms.

(4) We experimentally verify the effectiveness and efficiency of our graph compression schemes. We show that with cheap "once-for-all" compression, our compression methods can significantly reduce the inference cost of representative GNNs such as GCNs, GAT and GraphSAGE by 55%-85%, with little to no sacrifice of their accuracy.

Related work. Several approaches have been developed to accelerate GNN inference in large graphs [13, 20, 40, 57]. Closer to our approach is graph reduction, which simplifies graphs at a small sacrifice of model accuracy [25, 34]. There are three strategies.

Graph Sparsification. These methods (learn to) remove task-irrelevant edges from input graphs, such that the remaining part preserves the performance of GNNs. For example, AdaptiveGCN [33] learns an edge predictor to determine and remove task-irrelevant edges to accelerate GNN inference on CPU/GPU clusters. NeuralSparse [55] learns supervised DNNs to remove task-irrelevant edges. [13] proposed a framework to incorporates both model optimization and graph sparsification, which leverages lottery ticket hypothesisto identify subnetworks that can perform as well as the full network. Dspar [36] induces smaller subgraphs by removing edges that have similar "importance" (quantified by approximating a resistance measure as in circuits) to preserve graph spectrum.

Graph Coarsening. These methods group and amalgamate nodes into groups, without removing nodes. For example, SCAL [28] proposed the use of off-the-shelf coarsening methods LV[38] for scaling up GNN training and theoretically proved that coarsening can be considered a type of regularization and may improve the generalization as well as reduce the number of nodes by up to a factor of ten without causing a noticeable downgrade in classification accuracy. [31] introduced an optimization-based framework (FGC) that incorporates graph matrix and node features to jointly learn a coarsened graph while preserving desired properties such as spectral similarity [38]. GRAPE [50] is a GNN variant enhanced with sampled subgraph features from ego networks of automorphic equivalent nodes. It has a different goal of improving accuracy rather than reducing inference costs. [9] compresses graphs to accelerate GNN learning, using color refinement (with a case of bisimulation) to merge nodes within bounded radius. Node groups are iteratively refined based on a label encoding that concatenate node label and neighboring group colors. This is similar with SPGC. Nevertheless, no inference algorithm is provided. We show that our (α, r)-compression subsumes bisimulation compression.

Graph Sketch. These methods reduce the redundancy by generating a skeleton graph retaining essential structural information. For instance, Graph-Skeleton [11] constructs a compact, synthetic, and highly-informative graph for the target nodes classification by eliminating redundant information in the background nodes. While this approach may reduce the memory usage for a set of target nodes, it lacks flexibility and generalizability, as it must be tailored to specific nodes. In contrast, our SPGC performs one-time compression that applies universally to any set of test nodes $V_T \subset V$. NeutronSketch [35] focuses on eliminating the redundant information from the training portion of the graph, yet its compression is restricted to the training phase rather than accelerating inference, and does not ensure inference equivalence for the compressed graphs.

Our work differs from existing graph reduction approaches (summarized in Table. 1) in the following. (1) Our methods are model-agnostic and apply to any GNN that adopt the same inference process, without requiring model parameters, and incur no learning overhead. (2) We specify IFGC with variants that (approximately) preserve inferred results with invariant properties such as uniqueness and minimality, as well as fast compression and inference algorithms. These are not discussed in prior work. On the other hand, we remark that our scheme can be applied orthogonally: One can readily apply these approaches over compressed graphs from our method to further improve GNN training and inference. Our proposed SPGC can be potentially extended to the following learning settings: (1) transductive-learning: training on a compressed graph G_c while utilizing our proposed memoization structure $\mathcal{T}$ to enable accuracy-preserving inference; and (2) inductive-learning: applying semi-supervised or unsupervised learning techniques to refine the graph structure of G_c for a specific task.

2 GRAPHS AND GRAPH NEURAL NETWORKS

Graphs. A directed graph $G = (V, E)$ has a set of nodes V and a set of edges $E \subseteq V \times V$. Each node v carries a tuple $T(v)$ of attributes

and their values. The size of G, denoted as $|G|$, refers to the total number of its nodes and edges, *i.e.,* $|G| = |V| + |E|$.

Graph Neural Networks. A graph neural network (GNN) $\mathcal{M}$ is a mapping that takes as input a featurized representation $G = (X, A)$ to an output embedding matrix Z, *i.e.,* $\mathcal{M}(G) = Z$. Here X is a matrix of node features, and A is a (normalized) adjacency matrix of G. [1]

Inference. We take a query language perspective [6, 22] to characterize the inference process of GNNs. A GNN inference process is specified as a composition of *node update functions*.

Node update function. Given a GNN $\mathcal{M}$ with L layers, a node update function M_v uniformly computes the embedding of each node v at each layer k ($k \in [1, L]$), with a general recursive formula as

$$x_v^k = M_v^k(\Theta^k, \mathsf{AGG}(X_u^{k-1}, x_v^{k-1}, \forall u \in N^k(v)))$$

which is specified by (1) the learned model parameters Θ^k, (2) an aggregation function AGG (*e.g.,* $\sum$, CONCAT), and (3) the neighbors of v that participate in the inference computation at the k-th layer (denoted as $N^k(v)$). When k=1, X_v^0=$X_v \in X$, *i.e.,* the input features.

The *inference process* of a GNN $\mathcal{M}$ with L layers takes as input a graph $G = (X, A)$, and computes the embedding x_v^k for each node $v \in V$ at each layer $k \in [1, L]$, by recursively applying the node update function. A GNN $\mathcal{M}$ has a *fixed* inference process, if its node update function is specified by fixed input model parameters, layer number, and aggregator. It has a *deterministic* inference process, if $M(\cdot)$ always generates the same embedding for the same input.

We consider GNNs with fixed, deterministic inference processes. In practice, such GNNs are desired for consistent and robust performance. For simplicity, we assume that M_v specifies a proper set of neighbors that participate the inference process as $N(v) \subseteq \{u | (u, v) \in E \text{ or } (v, u) \in E\}$. This allows us to include GNNs that exploits neighborhood sampling (such as GraphSAGE), and directed message passing into discussion. In general, inferences of representative GNNs are in PTIME [12, 57] (see Table 3).

Classes of GNNs. We say a set of fixed, deterministic GNNs $\mathbb{M}$ belongs to a *class* of GNNs $\mathbb{M}^L$, if for every GNN $\mathcal{M} \in \mathbb{M}$, (1) $\mathcal{M}$ has L layers, and (2) $\mathcal{M}$ uses the same form of node update function M_v^k, for each node $v \in V$ and $k \in [1, L]$.

Table 3 summarizes several node update functions in their general forms for mainstream GNN classes. For example, Graph Convolution Networks (GCNs) [30] adopt a node update function as $X_v^k = \sigma(\Theta^k(\sum_{u \in \mathcal{N}(v)} \frac{1}{\sqrt{d_u d_v}} x_u^{k-1}))$. Here d_u or d_v denotes the degree of node u or v. $\sigma(.)$ is the non-linear activation function. A class of GNNs GCN^3 contains 3-layered GCNs that adopt such node update function. Note that two GNNs in the same class can have different Θ and output, given the same input.

3 INFERENCE-FRIENDLY COMPRESSION

Given a graph $G = (V, E)$, a *compressed graph* of G, denoted as $G_c = (V_c, E_c)$, is a graph where (1) each node $[v] \in V_c$ is a nonempty subset of V, and $V = \bigcup_{[v] \in V_c} [v]$; and (2) there is an edge $([v], [v']) \in E_c$, if there are at least one node $v \in [v]$ and $v' \in [v']$, such that $(v, v') \in E$.

[1] A feature vector X_v of a node v can be a word embedding or one-hot encoding [21] of $T(v)$. A is often normalized as $\hat{A} = A + I$, where I is the identity matrix.

Notation	Description
$G = (X, A)$	graph G, X: feature matrix, A: adjacency matrix
$\|G\|$	size of G; $\|G\| = \|V\| + \|E\|$
$\mathcal{M}$	a GNNs model
$\mathcal{M}(G)$ (resp. $\mathcal{M}(G, V_T)$)	output of $\mathcal{M}$ over G (resp. test set V_T)
$\mathcal{C}, \mathcal{P}$	compression, post-processing function
M_v	node update function
x_v^k	embedding of node v at layer k
G_c	compressed graph of G
α, r	similarity threshold, # hops
V_T, V_A	test node set, anchored nodes
$R^S, R^{(\alpha,r)}, R_L^A$	structural equivalence, (α, r) relation & anchored relation

Table 2: Summary of Notations.

Note that $|V_c| \leq |V|$ and $|E_c| \leq |E|$. Hence, $|G_c| \leq |G|$.

Inference-friendly Graph Compression. Given a set of GNNs $\mathbb{M}$ and a graph G, an *inference-friendly graph compression*, denoted as IFGC, is a pair $(\mathcal{C}, \mathcal{P})$ where

- $\mathcal{C}$ is a compression function that computes a compressed graph G_c of G ($G_c = \mathcal{C}(G)$);
- $\mathcal{P}$ is a function that restore the auxiliary information of nodes in G_c to their counterparts in G; and moreover,
- $\mathcal{M}(G) = \mathcal{M}(\mathcal{P}(G_c))$, for *any* GNN $\mathcal{M} \in \mathbb{M}$.

An IFGC aims to generate a compressed graph G_c with a smaller size, such that an inference query that requests output $\mathcal{M}(G, V_T)$ for any $V_T \subseteq V$ can be computed by a faster inference process of $\mathcal{M}$ over G_c only, even with a query-time overhead incurred by $\mathcal{P}$.

A Sufficient Condition. We next introduce a sufficient condition for the existence of IFGC. To this end, we start with a notion of *inference-equivalent* relation.

Inference equivalence. Given a class of GNN $\mathbb{M}^L$ and a graph G, a pair of nodes (v, v') in G are *inference equivalent w.r.t.* $\mathbb{M}^L$, denoted as $v \sim_M^L v'$, if for any $M \in \mathbb{M}^L$, $X_v^k = X_{v'}^k$ for any $k \in [1, L]$.

One can readily infer that for any two nodes $v \sim_M^L v'$, $M(v, G) = M(v', G)$. That is, inference equivalence of nodes ensure that they are all "indistinguishable" for the inference of any GNN $\mathcal{M} \in \mathbb{M}^L$.

Denote the binary relation (v, v') induced by inference equivalence as R_M^L, *i.e.,* $(v, v') \in R_M^L$ if and only if $v \sim_M^L$. We say R_M^L is *nontrivial* if there is at least one pair $(v, v') \in R_M^L$, where $v \neq v'$. We can readily verify the following result.

Lemma 1: *Given $\mathbb{M}$ and G, the binary relation R_M^L is an equivalence relation,* i.e., *it is reflexive, symmetric, and transitive.* □

The *equivalent class* of v under an equivalence relation R_M^L, denoted as $[v]$, refers to the set $\{v' | (v, v') \in R_M^L\}$. The equivalent classes induced by the inference equivalence relation R_M^L forms a node partition V_R of V. The *quotient graph* induced by R_M^L is a graph G_R with nodes V_R and edges E_R, where each node in V_R is a distinct equivalent class induced by R_M^L, and there is an edge $([v], [v']) \in E_R$ if and only if there exists a node $v \in [v]$ and $v' \in [v']$, such that $(v, v') \in E$.

Lemma 2: *Given a class of* GNNs $\mathbb{M}^L$ *and a graph G, a graph compression scheme $(\mathcal{C}, \mathcal{P})$ is an* IFGC w.r.t. $\mathbb{M}^L$ *and G, if for any*

GNNs Classes	Node Update Function (general form)	Training Cost	Inference Cost
Vanilla [44]	$X_v^k = \sigma(\Theta \cdot \text{AGG}(X_u^{k-1}, \forall u \in \mathcal{N}(v)))$	$O(L\|E\| + L\|V\|)$	$O(L\|E\| + L\|V\|)$
GCN [12, 30]	$X_v^k = \sigma(\Theta^k(\sum_{u\in\mathcal{N}(v)} \frac{1}{\sqrt{d_u d_v}} x_u^{k-1}))$	$O(L\|E\|F + L\|V\|F^2)$	$O(L\|E\|F + L\|V\|F^2)$
GAT [8, 48]	$X_v^k = \sigma(\sum_{u\in\mathcal{N}(v)} \alpha_{uv}\Theta^k X_u^{k-1}))$	$O(L\|E\|dF^2 + L\|V\|F^2)$	$O(L\|E\|dF^2 + L\|V\|F^2)$
GraphSAGE [12, 23]	$X_v^k = \sigma(\Theta^k \cdot (X_v^{k-1}\|\|\text{AGG}(X_u^{k-1}, \forall u \in \mathcal{N}(v))))$	$O(L\|V\|dF + L\|V\|F^2)$	$O(L\|V\|dF + L\|V\|F^2)$
GIN [8, 51]	$X_v^k = \sigma(\text{MLP}((1+\gamma)x_v^{k-1} + \sum_{u\in\mathcal{N}(v)} x_u^{k-1})$	$O(L\|E\|F + L\|V\|F^2)$	$O(L\|E\|F + L\|V\|F^2)$

Table 3: Comparison of Representative GNNs with node update functions, training cost, and inference cost. σ: an activation function *e.g.,* ReLU or LeakyReLU. AGG: aggregation function; can be *e.g.,* sum ($\sum$), average (Avg), or concatenation (||). L, $|E|$, $|V|$, F, and d denote the number of layers, edges, nodes, features per node, and maximum node degree of G, respectively.

$\mathcal{M} \in \mathbb{M}^L$, *(1)* $C(G)$ *computes a quotient graph* G_c *induced by a nontrivial inference equivalent relation* R_M^L w.r.t. $\mathbb{M}^L$ *and* G*, and (2)* $\mathcal{P}$ *is a function that restores* X_v^k *with* $X_{[v]}^k$ *by a scaling factor derived from auxiliary information of* v*, for each layer* $k \in [1, L]$. □

Proof sketch: Let R_M^L be a non-empty inference equivalent relation *w.r.t.* $\mathbb{M}^L$ and G, and G_c be the quotient graph induced by $\mathbb{M}^L$. (1) Given **Lemma 1**, R_M^L is a nontrivial equivalence relation. Hence there exists at least one equivalent class $[v]$ with size larger than one, *i.e.,* $|G_c| < |G|$. As function $\mathcal{P}$ does not introduce new node or edge to G_c, we have $|C(G))| = |G_c| < |\mathcal{P}(G_c)| < |G|$. (2) To see $\mathcal{M}(G) = \mathcal{M}(\mathcal{P}(C(G)))$, *i.e.,* G_c preserves inference result, it suffices to show that for every node $v \in G$, $\mathcal{M}(G, \{v\}) = M(\mathcal{P}(C(G)), \{[v]\})$. This is ensured by (a) the fixed deterministic inference process that applies the same node update function M_v, and (b) $\mathcal{P}$ restores X_v^k with only $X_{[v]}^k$ and a scaling factor, for any layer $k \in [1, L]$. We list examples of $\mathcal{P}$ and scaling factors for mainstream GNNs in Table 4. Hence $(C, \mathcal{P})$ is an IFGC. □

We next introduce practical IFGC for representative GNN classes, with efficient compression (implementing C) and inference (involving $\mathcal{P}$) algorithms. We summarize notations in Table 2.

4 STRUCTURAL-PRESERVING COMPRESSION

We introduce a first IFGC for GNN inference. We specify R_M^L as an extended version of *structural equivalence*. The latter has origins in role equivalence in social science [37], and simulation equivalence of Kripke structures in model checking [4, 15]. By enforcing equivalence on embeddings and neighborhood connectivity, it ensures an IFGC to accelerate GNN inference *without* decompression.

4.1 Compression Scheme

Structural equivalence. Given a graph G=(X, A), a *structural equivalence* relation, denoted as R^S, is a non-empty binary relation such that for any node pair (v, v') in G, $(v, v') \in R^S$, if and only if:

- $X_v^0 = X_{v'}^0$, *i.e.,* v and v' have the same input features;
- for any neighbor u of v ($u \in N(v)$), there exists a neighbor u' of v' (u' in $N(v')$), such that $(u, u') \in R^S$; and
- for any neighbor u'' of v' in $N(v')$, there exists a neighbor u''' of v in $N(v)$, such that $(u'', u''') \in R^S$.

Structural-preserving Compression. Given a GNN class $\mathbb{M}^L$ and graph G, a *structural-preserving compression*, denoted as SPGC *w.r.t.* $\mathbb{M}^L$, is a pair $(C, \mathcal{P})$ where (1) C computes G_c as the quotient graph of R^S, where R^S is the non-empty, maximum structural equivalence relation in G, and (2) $\mathcal{P}$ is a function that restores node embeddings with matching scaling factors for $\mathbb{M}^L$.

Example 2: Consider the graphs G_2 and G_3 in Fig. 2, and their compressed counterpart obtained by SPGC, G_{c_2} and G_{c_3}, respectively.

(1) G_2 has 10 nodes and 10 edges. The nodes having the same labels $'a', 'c', 'd'$ also have the same input features, respectively. For example, $X_{a_1}^0 = X_{a_2}^0$, and $X_{d_1}^0 = X_{d_2}^0 = X_{d_3}^0$. For nodes labeled with $'b'$, $X_{b_1}^0 = X_{b_3}^0 \neq X_{b_2}^0$. One can verify that $R^S = \{(b_1, b_3), (c_1, c_2), (d_1, d_3)\}$. A compressed graph G_{c_2} is illustrated with 7 nodes and 7 edges.

Observe that despite d_1, d_2 and d_3 have the same input features, $d_2 \not\sim_M^L d_1$, and $d_2 \not\sim_M^L d_3$ for GNNs with $L \geq 1$. Indeed, d_2 has a neighbor b_2 that has no counterpart in the neighbors of d_1 or d_3 that share the same embedding; hence the output embedding of d_2 may be different from either d_1 or d_3, and should be separated from equivalent class $\{d_1, d_3\}$ in G_c.

(2) G_3 is a cycle in the form of $\{c_n, b_n, a_n, \ldots, c_1, b_1, a_1\}$, where $X_{a_i}^0 = X_{a_j}^0, X_{b_i}^0 = X_{b_j}^0$, and $X_{c_i}^0 = X_{c_j}^0$, for any $i, j \in [1, n]$. We can verify that $R^S = \bigcup_{i,j\in[1,n]}\{(a_i, a_j), (b_i, b_j), (c_i, c_j)\}$. A smallest compressed graph G_{c_3} is illustrated with only three nodes: $A = \{a_i\}$, $B = \{b_i\}$, $C = \{c_i\}$, $\forall i \in [1, n]$, regardless of how large n is. □

The result below tells us that any two nodes that are structural equivalent are "indistinguishable" for GNN inference process.

Theorem 3: *Given a class of* GNNs $\mathbb{M}^L$ *and graph* G*, the relation* R^S *over* G *is an inference equivalence relation* w.r.t. $\mathbb{M}^L$. □

Proof sketch: First, R^S is an equivalence relation. It then suffices to show that for any pair $(v, v') \in R^S$, $v \sim_M^L v'$. We perform an induction on the number of layers k for GNNs. Consider a "matching" relation h between a pair $(v, v') \in R^S$, such that $h(v) = v'$. At any layer, for any node v and every neighbor $u \in N(v)$, there exists a "match" $h(v)$ and a "match" $h(u) \in N(h(v))$ with the same (intermediate) embedding. This ensures the equivalence of aggregated embedding computed by the node update function at v and $h(v)$, and vice versa. Hence for any pair $(v, v') \in R^S$, $v \sim_M^L v'$. R^S is thus an inference-friendly relation by definition. □

Following Lemma 2 and Theorem 3, SPGC is an IFGC.

Example 3: We also compare SPGC with several other possible graph compression scheme. We illustrate three more compressed graphs, $G_{c2}^1, G_{c2}^2, G_{c2}^3$ of G_2, following Exact Compression [9], Bisimulation [15], and Automorphism [50], respectively.

(1) Exact Compression [9] applies an iterative color refinement process[2] starting with groups that contains nodes agreeing on embeddings and color encoding (labels). It then iteratively split the groups, where each node is updated by concatenating its color

[2]The original method is used to simplify GNNs learning problems [9]; we make a comparison by applying color refinement for graph compression alone.

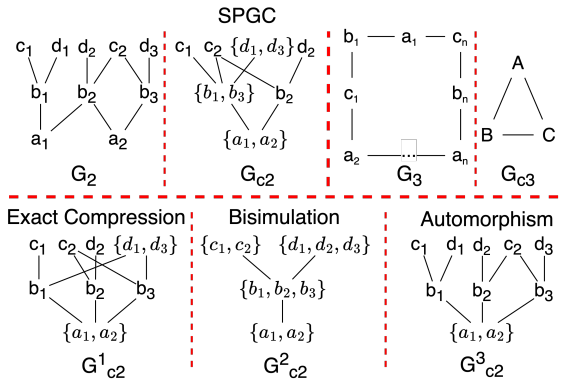

Figure 2: Compressing graphs with SPGC and a comparison with exact compression [9] , bisimulation [15], and automorphism (adapted from [50]). Exact compression uses node features for initial coloring and $d = 2$ for color refinement.

encoding with those from their neighbors, and refine groups. It finally derives $G^1_{c_2}$ with 8 nodes and 9 edges after two rounds, where only node pairs d_1, d_3 and a_1, a_2 share the same node representation. The concatenation is more sensitive to the impact of *e.g.,* degrees, preventing more possible merge, hence less can be compressed.

(2) Bisimulation [15] ignores embedding equivalence and groups nodes only with topology-level equivalence. In this case, b_2 can be merged with b_1 and b_3 due to bisimulation in connectivity. This leads to smaller compressed graph $G^2_{c_2}$ with 4 nodes and 3 edges compared with G_{c_2}, yet at the cost of inaccurate inference at *e.g.,* nodes c_1, c_2, b_2, and d_2 due to "overly" compressed structure.

(3) Automorphism [50] partitions nodes into same automorphism equivalence sets, which poses strong topological equivalence on graph isomorphism in their neighbors. By enumerating the automorphism groups of G_2 and considering the embedding similarities, only a_1 and a_2 can be merged as shown in G^3_{c2} while other node pairs like c_1, d_1 or b_1, b_3 cannot be merged due to different embeddings or different connection patterns of the neighbors. This can be an overkill for reducing unnecessary inference computation. In addition, computing automorphism remains NP-hard, which indicates more expensive compression cost; while SPGC computes maximum structural equivalence in PTIME (see Section 4.3). □

4.2 Properties and Guarantees

We next justify SPGC by showing a *minimality and uniqueness* property. We show that an SPGC can generate a smallest G_c, which is "unique" up to graph isomorphism. That is, if there is another smallest compressed graph G'_c generated by an SPGC, then G^*_c and G'_c are isomorphic.

Lemma 4: *Given a* GNN *class* $\mathbb{M}^L$ *and* G*, there is an* SPGC *that computes a smallest* G^*_c *which is unique up to graph isomorphism.* □

Proof sketch: We show the minimality property by a constructive proof as follows: (1) given $\mathbb{M}^L$ and G, there exists a unique, largest inference equivalence relation R^{*S}; (2) we construct an SPGC that computes G^*_c as the quotient graph of R^{*S}. The uniqueness of G^*_c can be shown by a contradiction: if there exists another smallest G'_c that is not graph isomorphic to G^*_c, then either G'_c is not smallest in sizes, or R^{*S} is not the (largest) inference equivalence relation, *i.e.,* there is a pair (v, v'), such that either $v \sim^R_M v'$ but are not in $[v]$, or $v \nsim^R_M v'$, but are included in $[v]$. Either leads to contradiction. □

Algorithm 1 : SPGC

Input: Graph G, node feature matrix X, a class of GNNs $\mathbb{M}^L$ with node update function M_v;
Output: A compressed graph G_c with memoization structure $\mathcal{T}$;
1: set $R^S := \emptyset$; set $EC := \{V\}$; set $\mathcal{T} := \emptyset$; graph $G_c := \emptyset$;
2: $R^S := \text{DPP}(G)$;
3: $R^S := R^S \setminus \{(v, v') \mid X^0_v \neq X'^0_v\}$;
4: $EC := V/R^S$; /* induce partition EC from refined R^S*/
5: $(G_c, \mathcal{T}) := \text{CompressG}(\mathcal{T}, G_c, EC, G, M)$;
6: **return** G_c and $\mathcal{T}$;

Figure 3: Algorithm SPGC

We next justify SPGC by showing that it properly preserves the discriminative set of GNNs, which has been used as one way to characterize the expressiveness power of GNNs as queries [6, 22].

Discriminative set of GNNs [22]. Given a set of graphs $\mathcal{G}$, the *discriminative set* of a GNN $\mathcal{M}$, denoted as $\mathcal{G}_\mathcal{M}$, refers to the maximum set of pairs $\{(G, G')\}$, where $G, G' \in \mathcal{G}$, such that $M(G) = M(G')$. In the case of equivariant GNNs [43], the strongest discriminativeness can be achieved, for which the set contains all pairs (G, G') such that G and G' are isomorphic [5]. In other words, these GNNs can "solve" graph isomorphic problem: one can issue a Boolean inference query to test if an input pair of graphs are isomorphic.

Given a set of graphs $\mathcal{G}$, denote the set of corresponding compressed graphs generated by SPGC as $\mathcal{G}_c$, *i.e.,* $\mathcal{G}_c = \{G_c | G_c = C(G); G \in \mathcal{G}\}$. We have the following result.

Lemma 5: *Given* $\mathbb{M}^L$ *and a set of graphs* $\mathcal{G}$*, an* SPGC *can compute a compressed set* $\mathcal{G}_c$*, such that for every* GNN $\mathcal{M} \in \mathbb{M}^L$*, and any pair* $(G, G') \in \mathcal{G}_\mathcal{M}$*, there exists a pair* $(G_c, G'_c) \in \mathcal{G}_{c_\mathcal{M}}$. □

This result tells us that SPGC "preserves" the discriminativeness of GNNs. Moreover, it suggests a practical compression scheme for large-scale graph classification. One can apply SPGC to compress $\mathcal{G}$ to a smaller counterpart $\mathcal{G}_c$. As the discriminativeness set is preserved over $\mathcal{G}_c$ for every GNN $M \in \mathcal{M}^L$, SPGC reduces the overall classification overhead, via a post-processing $\mathcal{P}$ that readily groups $\mathcal{G}$ by corresponding label groups over $\mathcal{G}_c$.

Due to limited space, we present the detailed proofs in [1].

4.3 Compression Algorithm

We next present a compression algorithm (function C) in SPGC.

General idea. The algorithm, simply denoted as SPGC, follows Lemma 4 to construct the smallest G^*_c induced by the maximum structure equivalence relation R^*_S. To ensure efficient inferences that only refer to G_c without decompression, it (1) uses a *memoization* structure $\mathcal{T}$ to cache the neighborhood statistics specified by node update function M_v, and (2) *rewrites* M_v to an equivalent counterpart $M_{[v]}$ (see Table 4 for examples), such that the inference can directly process on each $[v]$ in G_c, and "looks up" $\mathcal{T}$ at runtime, to obtain the embeddings for all the nodes in $[v]$, in a single batch.

Compression Algorithm. The SPGC algorithm, as illustrated in Fig. 3, takes as input a featurized input $G = (X, A)$ and a GNN class

Algorithm 2 Procedure CompressG($\mathcal{T}$, G_c, EC, G, M)

1: **for** $[v] \in EC$ **do**
2: $V_c = V_c \cup \{[v]\}$;
3: initialize $[v]_T$; /* with row pointers as $v \in [v]$*/
4: **for** edge $(u, v) \in E$ **do**
5: $E_c = E_c \cup \{([u], [v])\} \mid u \in [u], v \in [v]$;
6: **if** $M.\phi$ is topology sensitive **then**
7: $[v]_T(v, [u])$ += $\frac{1}{\sqrt{deg(u)}}$;
8: **else if** $M.\phi$ is weight sensitive **then**
9: $[v]_T(v, [u])$ += $\alpha_{v,u}$;
10: **else**
11: $[v]_T(v, [u])$ += 1;
12: $\mathcal{T} = \bigcup_{[v] \in V_c} [v]_T$;
13: **return** G_c and $\mathcal{T}$;

Figure 4: Procedure CompressG

$\mathbb{M}^L$ with node update function M_v. (1) It first extends Dovier-Piazza-Policriti (DPP) algorithm [15] to compute the maximum structural equivalence relation R_S^*, by enforcing embedding equivalence as an additional equivalence constraint (lines 2-4). This induces a set of equivalence classes EC (a node partition). It then invokes a procedure CompressG to construct G_c as the quotient graph of R_S^*, as well as the memoization structure $\mathcal{T}$ (line 5).

Procedure CompressG. CompressG is a light-weighted compression approach enabling an easy-to-implement and feasible way to derive compressed graph G_c from large-scale G while computing the memoization structure $\mathcal{T}$. Given the induced equivalence classes EC and an encoding of the node update function M_v, procedure CompressG (illustrated in Fig. 4), generates the compressed graph G_c and memoization structure $\mathcal{T}$. For each equivalent class $[v]$ in EC, CompressG initializes a node in G_c, (lines 1-3). For each edge $(u, v) \in E$, CompressG adds an edge $([u], [v])$ (lines 4-5).

Compression time memoization. CompressG dynamically maintains a memoization structure $\mathcal{T}$ that is shared by all GNNs in $\mathcal{M}^L$, to cache useful auxiliary neighborhood information used by the node update function for efficient inference (see Section 4.4). For each node $[v] \in V_c$, it assigns $[v]_{\mathcal{T}}$, a compact table, such that for every $v \in [v]$, and every neighbor $[u] \in N([v])$, an entry $[v]_{\mathcal{T}}(v, [u])$ records an aggregation of auxiliary neighborhood information (*e.g.,* sum of node degree, edge weights) of $N(v) \subseteq N([v])$.

When processing an edge $(u, v) \in E$, it follows a case analysis of $\mathbb{M}^L$ with node update function M_v. For example, (1) "topology sensitive" means the degrees of neighbors of v is required, as seen in GCNs; (2) "weight sensitive" means additional edge weights, such as edge attentions in GATs. For GATs, the edge weights from pre-trained M are optional such that if given the information of all edge weights, the inference accuracy can be preserved based on the formula shown in Table. 4.4. Such information can be readily obtained by tagging the input GNN class $\mathbb{M}^L$ or encoded as rules. It then updates the entry $[v]_{\mathcal{T}}(v, [u])$ accordingly (lines 6-12).

Example 4: Consider a GNNs class GIN, and graph G_4 shown in Fig. 5. (1) SPGC invokes the DPP algorithm to compute R^S. It next refines R^S based on feature embeddings and returns the induced

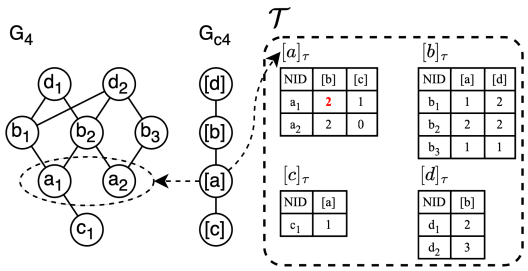

Figure 5: Run-time generation of Memoization structure $\mathcal{T}$.

partition $EC = \{[a], [b], [c], [d]\}$, where nodes with same labels are merged, *e.g.,* $[a] = \{a_1, a_2\}$. (2) We illustrate how CompressG dynamically updates the memoization structure $\mathcal{T}$ by considering the processing of two edges (b_1, a_1) and (b_2, a_1). For $[a] \in EC$, it first initializes $[a]_{\mathcal{T}}$ as an empty table. It next iterates over edges in E. As the node update function of GIN does not require exact degree or additional edge weight (not topology or weight sensitive), the entry $[a]_T(a_1, [b])$ is updated to 1, to "memoize" that there is one neighbor of a_1 in $N([b])$ that will contribute to a "unit" value to the embedding computation of a_1, via edge (b_1, a_1). Similarly, when it reaches edge (b_2, a_1), $[a]_T(a_1, [b])$ is updated to 2. Following this processes, all entries in $\mathcal{T}$ will be updated to memoize neighbors' information while compressing the graph. □

Correctness and cost. SPGC correctly computes G_c^* as ensured by (1) the correctness of DPP algorithm and (2) the follow-up refinement by enforcing embedding equivalence. For time cost, it takes SPGC $O(|V| + |E|)$ time to initialize R^S with DPP algorithm. The refinement of R^S and EC takes $O(|V|)$ time (lines 3-4). Procedure CompressG processes each equivalent class in EC ($|EC| \leq |V|$) and each edge in G once, hence in $O(|V| + |E|)$ time to construct G_c^* and update $\mathcal{T}$. The total cost is thus in $O(|V| + |E|)$ time.

4.4 Inference Process

Inference algorithm. We outline an algorithm that directly obtains $M(G)$ by referring to G_c^* only, without decompression. Our strategy rewrites the node update function M_v to an equivalent counterpart $M_{[v]}$, that takes as input $[v]$ and the corresponding tuple $[v]_{\mathcal{T}}(v)$ in $\mathcal{T}$, to "scale" the embedding computation with the memorized edge weights. The algorithm performs inference directly in G_c^* with $M_{[v]}$, and simply "scale up" the results at $[v]$ for each node $v \in [v]$ with a scaling factor. The scaling factor can be directly looked up from the table $[v]_{\mathcal{T}}(v)$ (function $\mathcal{P}$). Table 4 illustrates the scaling factors for mainstream GNN classes.

Example 5: Continuing with Example. 4, an inference at node a_1 looks up, in constant time, the values from entries $[a]_T(a_1, [b])$ and $[a]_T(a_1, [c])$ which are 2 and 1 separately (as shown in Fig. 5). It next assigns the values as scaling factors (Table. 4) to restore messages and the embedding of node a_1 as in original G_4. □

Inference cost. As SPGC requires no decompression on neighborhood structures of nodes, an inference query can be directly applied to G_c without incurring additional overhead. The overall inference cost is in $O(L|E_c|F + L|V_c|F^2)$. We remark that this result is derived

GNNs	Node Update Function M_v	equivalent rewriting $M_{[v]}$; scaling factors are marked in red	notes
Vanilla [44]	$X_v^k = \sigma(\Theta \cdot \text{AGG}(X_u^{k-1}, \forall u \in \mathcal{N}(v)))$	$X_v^k = \sigma(\Theta \cdot \text{AGG}(\mathbf{[v]_T(v,[u])} X_{[u]}^{k-1}, \forall [u] \in \mathcal{N}([v])))$	AGG: $\sum$ or AVG; for AVG, need to multiple by RF_v
GCN [30]	$X_v^k = \sigma(\Theta^k(\sum_{u\in\mathcal{N}(v)} \frac{1}{\sqrt{deg_u deg_v}} x_u^{k-1}))$	$X_v^k = \sigma(\Theta^k(\sum_{[u]\in\mathcal{N}([v])} \frac{1}{\sqrt{\text{deg}_v}} \mathbf{[v]_T(v,[u])} x_{[u]}^{k-1}))$	deg_v: degree of node v in G, topology sensitive
GAT [48]	$X_v^k = \sigma(\sum_{u\in\mathcal{N}(v)} \alpha_{uv}\Theta^k X_u^{k-1}))$	$X_v^k = \sigma(\sum_{[u]\in\mathcal{N}([v])} \mathbf{[v]_T(v,[u])}\Theta^k X_{[u]}^{k-1}))$	weight sensitive
GraphSAGE [23]	$X_v^k = \sigma(\Theta^k \cdot (X_v^{k-1} \Vert \text{AGG}(X_u^{k-1}, \forall u \in \mathcal{N}(v))))$	$X_v^k = \sigma(\Theta^k \cdot (X_{[v]}^{k-1} \Vert \text{AGG}(\mathbf{RF_v \times [v]_T(v,[u])} X_{[u]}^{k-1}, \forall [u] \in \mathcal{N}([v]))))$	$\Vert$: concatenation; AGG: AVG
GIN [51]	$X_v^k = \sigma(\text{MLP}((1+\gamma)x_v^{k-1} + \sum_{u\in\mathcal{N}(v)} x_u^{k-1}))$	$X_v^k = \sigma(\text{MLP}((1+\gamma)x_{[v]}^{k-1} + \sum_{[u]\in\mathcal{N}([v])} \mathbf{[v]_T(v,[u])} x_{[u]}^{k-1}))$	

Table 4: Rewriting of node update functions for mainstream GNN classes (scaling factors highlighted in red).

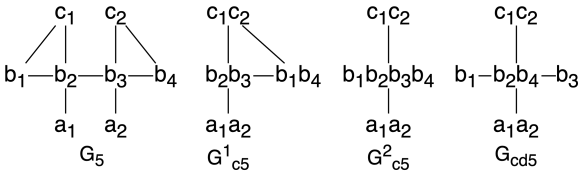

Figure 6: Compressing G_5 with $(0.5, 2)$-SPGC.

by scaling down a common upper bound of inference costs for mainstream GNNs in Table 3. For other and more complex GNNs variants, the inference costs can be derived similarly by scaling down from their counterparts over G.

5 CONFIGURABLE GRAPH COMPRESSION

While SPGC generates G_c that can be directly queried by inference queries without decompression, it enforces node embedding equivalence, which may be an overkill for nodes with similar embeddings and can be processed in a single batch with tolerable difference in query outputs. Users may also want to *configure* the compression schemes to balance among accuracy and speed up, or to contextualize the compression with inference queries that specifies a set of test nodes $V_T \subseteq V$ of interests, such that $M(G, V_T) = M(G_c, V_T)$.

In response, we next introduce two variants of SPGC: (α, r)-SPGC (Section 5.1), and anchored SPGC (Section 5.2), respectively.

5.1 Compression with Structural and Embedding Similarity

We start with a relation called (α, r)-relation, which approximates R^S by lifting its equivalence constraints.

(α, r)-relation. Given graph G, a configuration (xsim, α, r) is a triple where $\text{xsim}(\cdot)$ is a *feature similarity function* that computes a similarity score of two node embeddings, α is a similarity threshold ($\alpha \in [0, 1]$), and r an integer. Let $N_r(v)$ be the nodes within r-hop neighbors of v. A binary relation $R^{(\alpha,r)} \subseteq V \times V$ is an (α, r)-relation if for any node pair $(v, v') \in R^{(\alpha,r)}$,

- $\text{xsim}(X_v^0, X_{v'}^0) \geq \alpha$;
- for any node $u \in N(v)$, there exists a node $u' \in N_r(v'))$, such that $(u, u') \in R^{(\alpha,r)}$; and
- for any node $u'' \in N(v')$, there exists a node $u''' \in N_r(v)$, such that $(u'', u''') \in R^{(\alpha,r)}$.

Note that $R^{(1,1)}$ is an R^S, as $\alpha = 1$ ensures embedding equivalence, and $r = 1$ preserves indistinguishable neigbhors for node update functions in GNN inference. On the other hand, (α, r)-relation is no longer an equivalence relation, as it "relaxes" structural equivalence by lifting both embedding equality, and the strict neighborhood-wise equivalence, in trade for better compression ratio. We next clarified the relationship between r and L. Based on the definition of SPGC, the embedding similarity between two vertices is dependent on r instead of L. In other words, SPGC is model-agnostic as long as $r \leq L$. In practice, since L is small due to over-smoothing issue of larger L, r is usually a small integer like 1 or 2.

Based on the relation $R^{(\alpha,r)}$, we introduce a variant of SPGC.

(α, r)-SPGC. Given a graph G, and a configuration α and r *w.r.t.* an embedding similarity measure and a threshold, an (α, r)-SPGC is a graph compression scheme if C computes a graph G_c induced by the relation $R^{(\alpha,r)}$. Specifically,

- for any node pair $(v, v') \in R^{(\alpha,r)}$, $v \in [v]$, $v' \in [v]$; and
- there is an edge between $([u], [v])$ if $(u, v) \in E$.

Lemma 6: *Given a* GNN *class $\mathcal{M}^L$ and a graph G, an (α, r)-*SPGC *incurs compression cost in $O(|V||N_r| + |E|)$ time ($|N_r|$ refers to the largest size of r-hop neighbors of a node in G), and an inference cost in $O(L|E|F + L|V_c|F^2)$ time.* □

As a constructive proof, we next introduce algorithms that implements (α, r)-SPGC with the above guarantees.

Compression Algorithm. We describe the compression algorithm (α, r)-SPGC. It follows the same principle to compute (α, r)-relation and induce a compressed graph. The difference is that rather than inducing equivalence class and quotient graph from G, (1) it first induces a graph G_r by linking nodes to their r-hop neighbors, (2) it then computes an R^S relation by invoking DPP algorithm, and refines it by a re-grouping of nodes determined by similarity function $\text{xsim}(\cdot)$ with α as similarity threshold. (3) It generates $[v]$ to include all the pairs (v, v') in $R^{(\alpha,r)}$, and accordingly the edges. It updates the memoization structure $\mathcal{T}$ following the edges in G_r, similarly as in SPGC. Here $\mathcal{T}$ caches the statistics from the r-hop neighbors of each node v in the original graph G.

Example 6: Consider graph G_5 shown in Fig. 6. A $(0.5, 2)$-SPGC initializes a $(1, 1)$-relation, and refines it to $R^{(0.5,2)} = \{(a_1, a_2), (c_1, c_2), (b_1, b_2), (b_1, b_3), (b_1, b_4), (b_2, b_3), (b_2, b_4), (b_3, b_4)\}$. This yields a compressed graph $G_{c_5}^2$ with only 3 nodes and 2 edges. We also illustrate $G_{c_5}^1$ compressed from G_5 (induced by an R^S as a $(1, 1)$-relation). Due to strictly enforced embedding equivalence, b_1 and b_2 cannot be merged, and similarly for b_3 and b_4. This yields $G_{c_5}^1$ with more nodes and edges. □

Compression Cost. It takes $O(|V| \cdot |N_r(v)|)$ time to derive G_r for $v \in V$. It then takes $O(|V| + |E_r|)$ to compute and refine $R^{(\alpha,r)}$, for G_r with edge set E_r. Here $|E_r| \leq |V| \cdot |N_r|$, where N_r refers to the largest r-hop neighbor set for a node in G. Procedure CompressG

constructs G_c in $O(|V| + |E_r|)$ time, and generates memoization structure $\mathcal{T}$ in $O(|E|)$ time. The total cost is thus in $O(|V||N_r|+|E|)$.

As (α, r)-SPGC is specified by $R^{(\alpha,r)}$ that approximates an inference-friendly relation, it is no longer an IFGC, hence a direct inference over the G_c from it may not preserve the original output. To mitigate accuracy loss, the inference specifies a procedure $\mathcal{P}$ to perform run-time decompression with small overhead.

Inference process with decompression. The inference algorithm directly processes each node $[v]$ in G_r as in SPGC. For (α, r)-SPGC. the difference is that it ad-hocly invokes a decompression procedure decompG (the decompG algorithm and its example illustrated in Figure. 18 are shown in the Appendix [1]) to reconstruct the neighbors of v in G_r, and performs an inference using original 1-hop neighbors of v to obtain an embedding X_v as close as its original counterpart in G_r. To minimize decompression cost, when compressing the graph, algorithm (α, r)-SPGC (shown in Appendix B) incorporates *Re-Pair*, a reference encoding method [14, 32], to derive AL_c and rules for later fast decompression from a compact encoding structure. Specifically, within each *EC*, procedure decompG sorts the processing order of nodes by their degrees in G_r. In other words, procedure decompG prioritizes the decompression of nodes having the most shared r-hop neighbors with others in G_r, to (1) maximally reduce redundant computation in decompression process, and (2)"maximize" the likelihood for more accurate inference computation. For example, a partially decompressed graph G_{cd_5} that resolves 1-hop neighbors of b_2 (in trade for more accurate embedding) is illustrated in Fig. 6 (please refer to Example. 8 and Figure. 18 in Appendix for details [1]). The decompressed neighbors are kept in G_{cd} until the inference terminates.

As the decompression restores at most $|E|$ edges, the overall inference process takes $O(L|E|F + L|V_c|F^2)$ time, including the decompression overhead. We present the details of decompression algorithm in [1]. The above analysis completes the proof of Lemma 6.

5.2 Anchored Graph Compression

We next introduce our second variant of SPGC, notably, *anchored* SPGC, which permits a decompression-free, inference friendly compression, *relative* to a specific set of nodes of interests.

We present our main result below.

Theorem 7: *Given $\mathcal{M}^L$ and G with a set of targeted nodes V_A, there exists an* IFGC *that computes a compressed graph in $O(|G_L|)$ time to preserve the inference output for every node in V_A at an inference cost in $O(L|E_c|F + L|V_c|F^2)$ time. Here $|G_L|$ refers to the subgraph of G induced by L-hop of V_A, and $|V_c|$ and $|E_c|$ are bounded by $|G_L|$.* □

As a constructive proof, we introduce a notion of anchored relation, and construct such an IFGC. Given a graph G with a set of designated targeted nodes V_A, and a class of GNNs $\mathcal{M}^L$, a graph compression scheme $(\mathcal{C}, \mathcal{P})$ is an IFGC relative to V_A, if (1) $|\mathcal{P}(\mathcal{C}(G))| < |G|$; and (2) for any GNN $M \in \mathcal{M}^L$, and any $v \in V_A$, $M(G, \{v\}) = M(\mathcal{P}(\mathcal{C}(G), \{v\})$.

Anchored relation. Given graph G, an integer L, and a designated *anchor set* $V_A \subseteq V$, we define the L-hop neighbors of V_A, denoted as $N_L(V_A)$, as $\bigcup_{v \in V_A} N_L(v)$, where $N_L(v)$ refers to the set of nodes within L-hop of v in G. An *anchored relation* R_L^A *w.r.t.* V_A refers to the structural equivalence relation defined over the subgraph G_L of G induced by $N_L(V_A)$.

One may verify that (1) R_L^A is an equivalence relation over V, and (2) $R^S = R_L^A$ if $V_A = V$, and L is larger than the diameter of G.

Anchored SPGC. Given graph G and an anchor set V_A from G, an anchored SPGC, denoted as ASPGC, is a graph compression where $\mathcal{C}$ computes a compressed graph that is the quotient graph of R_L^A. The computation of compressed graph G_c using ASPGC follows its counterpart in SPGC. The difference is that it induces a subgraph G_L of G with the L-hop neighbors of all the nodes in V_A. It then invokes the compression algorithm of SPGC to derive R_L^A and applies the compression algorithm of SPGC on G_L to compute the G_c and memoization structure $\mathcal{T}$. The inference process over G_c, similarly, follows its SPGC counterpart over G_c, which consistently leverages $\mathcal{T}$ to efficiently recover the auxiliary information of neighborhoods. Hence, it preserves the inference results of the nodes in V_A for GNNs classes with up to L layers. In practice, one may set V_A simply as a set of test nodes V_T to adapt ASPGC for specific inference queries. We present details and an example in [1].

Analysis. The correctness of ASPGC follows from the data locality of L-layered GNN inference when V_A is specified, which involves the subgraph L-hops induced of $G^L(V_A)$. ASPGC next follows SPGC to compute G_c from G_L. For compression cost, it takes $O(|N^L(V_A)|+|E|)$ time to induce G_L, and $O(|V_L|+|E_L|)$ time to construct G_c from G_L. (3) The inference cost is in $O(L|E_c|F + L|V_c|F^2)$, where both $|E_c|$ and $|V_c|$ are bounded by $|G_L|$.

Given the above analysis, Theorem 7 follows.

6 EXPERIMENTAL STUDY

Using both real-world graph datasets and large synthetic graphs, we conducted four sets of experiments, to understand (1) effectiveness of our compression methods, in terms of compression ratio, and the trade-off between inference cost and accuracy loss; (2) their efficiency, in terms of the compression cost and inference cost, (3) impact of critical factors, and (4) an ablation study to evaluate the effectiveness of memoization, and decompression overhead.

6.1 Experimental Settings

Datasets. We employ four real-world graph benchmark datasets (summarized in Table. 5): (1) **Cora** [39], a citation network where nodes represent documents, and edges are citations among the documents. Node features are described by a 0/1-valued word vector indicating the absence/presence of the corresponding word from the dictionary; (2) **Arxiv** [27], a citation network with nodes representing arXiv papers and edges denoting one paper cites another one. The features of each node includes a 128-dimensional feature vector obtained by averaging the embeddings of words in its title and abstract; (3) **Yelp** [53] is prepared from the raw json data of businesses, comprising a dense network of user-business interactions. Nodes represent users and edges are generated if two users are friends. Node features are the word embeddings derived by reviews of products; (4) ***Products*** [27], a product co-purchase network, where nodes represent products sold in Amazon and edges denote products purchased together; and (5) **ogbn-papers100M**, a billion-scale citation graph of 111 million papers and 1.6 billion

Dataset	|V|	|E|	# node types	# attributes
Cora	2,708	5,429	7	1,433
Arxiv	169K	1.2M	40	128
Yelp	717K	7.9M	100	300
Products	2.4M	61.9M	47	100
WS-MAG	2M	1.192B	153	768
ogbn-Papers100M	111M	1.6B	40	128

Table 5: Summary of Datasets.

edges. Its graph structure and node features are constructed in the same way as **Arxiv**. The size of test nodes $|V_T| = 5\% * |V|$ across all datasets. For ASPGC, $|V_A| = |V_T|$ as $V_A = V_T$ by default.

Besides real-world benchmark datasets, we also generated a large synthetic dataset **WS-MAG**, by extending a core of real MAG240M network [26] (a citation network) with a small world generator [49]. **WS-MAG** has a fixed number of nodes ($|V|$ = 2 million), while the number of edges $|E|$ increases linearly from 4 million to 1.192 billion.

Graph Neural Networks. We have pre-trained three classes of representative GNNs: GCNs[3] [30], GATs[3] [48], and GraphSAGE[3] [23], for each dataset. For a fair comparison, (1) for all the datasets, we consider node classification, and (2) all compression methods are applied for the same set of GNNs.

Compression Methods. We compare SPGC and its variants, (α, r)-SPGC and ASPGC, with two state-of-the-art compression methods: (1) DSpar [36], a graph sparsification method that performs edge down-sampling to preserve graph spectral information, and (2) FGC [31], a latest learning-based graph coarsening approach that learns a coarsened graph matrix and feature matrix to preserve desired graph properties, such as homophily. We are aware of other learning-based approaches, yet they are model-specific and require the learned model parameters. Our work is orthogonal to these methods, and is not directly comparable. When conducing the experiments using DSpar and FGC, we set the longest waiting limit = 5 hours which means if the compression graph cannot be generated after 5 hours. We simply exclude them since the compression cost is already comparable to training a new GNNs model.

Evaluation Metrics. Given graph G, a class of GNNs $\mathcal{M}^L$, a graph compression scheme $(C, \mathcal{P})$ that computes a compressed graph G_c, and its matching inference process over G_c, we use the following metrics. (1) For efficiency, we evaluate (a) the time cost of compression, and (b) the speed-up of inference, which is defined by $\frac{T_{MG}}{T_{MC}}$, where T_{MG} refers to the inference time cost over G, and T_{MC} represents its counterpart over G_c. (2) For effectiveness, we report (a) a normalized compression ratio, which is defined as ncr = $1 - \frac{|G_c|}{|G|}$; Intuitively, it quantifies the fraction of G that is "reduced": the larger, the better; and (b) the model performance quantified by accuracy and $F1$-score over G_c. In particular for Yelp, over which the benchmark task is a multi-class node classification, we report micro $F1$-score. We report the average performance of 200 inference tests, for each GNN model over each dataset.

Environment. SPGC and its variants are developed in Python with PyTorch Geometric [19] and BisPy [3] libraries. All tests are conducted on 4 Intel(R) Xeon(R) Silver 4216 CPU @ 2.10GHz, 128 GB Memeory, 16 cores, and 1 32GB NVIDIA V100 GPU. Our source code, datasets, and a full version of the paper are made available[3].

[3]https://github.com/Yangxin666/SPGC

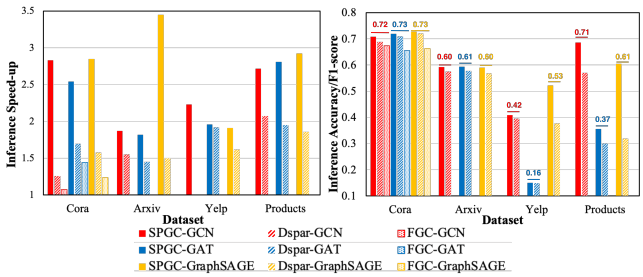

Figure 7: Comparison of $(0.5, 1)$-SPGC with the Baselines in Inference Speed-up and Inference Accuracy/F1-score.

6.2 ExperimentalResults

Exp-1: Effectiveness: Accuracy vs. Speed-up. Fig. 7 compares $(0.5, 1)$-SPGC with DSpar and FGC, in terms of inference speed-up in left figure and inference accuracy/F1-score in the right figure, over all four real datasets. Note that FGC can only generate results for **Cora** since its compression cannot be completed within 5-hours for all other datasets. Here a test annotated as "compression method-GNN" refers to the setting that a GNN inference is applied on a compressed graph generated by the method. (1) $(0.5, 1)$-SPGC outperforms DSpar and FGC across all four datasets for all the GNN classes on inference speed-up. It can improve the inference efficiency better over larger graphs. For example, for ***Arxiv***, $(0.5, 1)$-SPGC achieves a speed-up of 3.4 for inference with GraphSAGE, while DSpar achieves a speed-up to 1.5. (2) Consistently, we found that SPGC achieves higher ncr than DSpar and FGC. For example, for **Cora**, SPGC achieves ncr up to 74.5% while DSpar and FGC achieves 46.2% and 30.7% respectively. (3) SPGC outperforms DSpar and FGC in terms of its ability to preserve inference results for GNNs. We observe that i) $(0.5, 1)$-SPGC achieves inference accuracy comparable to direct inference on the original graph G across all four datasets (as illustrated in Fig. 7); and ii) $(0.5, 1)$-SPGC retains highest inference accuracy/F1-scores across all datasets and models. For example, for **Cora**, $(0.5, 1)$-SPGC achieves 0.71 accuracy which outperforms 0.68 and 0.67 achieved by DSpar and FGC.

Exp-2: Effectiveness: Impact of Factors. We first investigate the impact of number of layers L, which evaluates whether the quality of SPGC-based compression is affected by the complexity of GNNs classes. Then we evaluate the performance of configurable compression (α, r)-SPGC, in terms of the impact of α and r.

Varying Number of Layers. We varied the number of layers of GNNs from 2 to 4 over **Arxiv** and report its impact on inference speed-up (resp. accuracy) in Fig. 8(a) (resp. 8(b)). (1) $(0.5, 1)$-SPGC consistently outperforms all the baselines in both inference speed-up and accuracy, due to that it preserves the inference results with small compressed graphs. (2) In general, the speed-up of inference achieved by $(0.5, 1)$-SPGC is not sensitive to the number of layers. This verifies our theoretical analysis that it preserves inference results with unique smallest compressed graphs, which is independent of model complexity. (3) While the accuracy of GNNs drops as the number of layers become larger, in all cases, $(0.5, 1)$-SPGC preserves the accuracy with smallest "gap" compared with other methods, for the same class of GNNs.

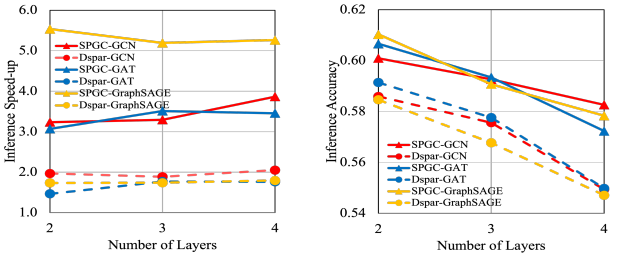

(a) Num. Layers v.s. Speed-up (b) Num. Layers v.s. Accuracy

Figure 8: Varying Num. Layers in GNNs (Arxiv).

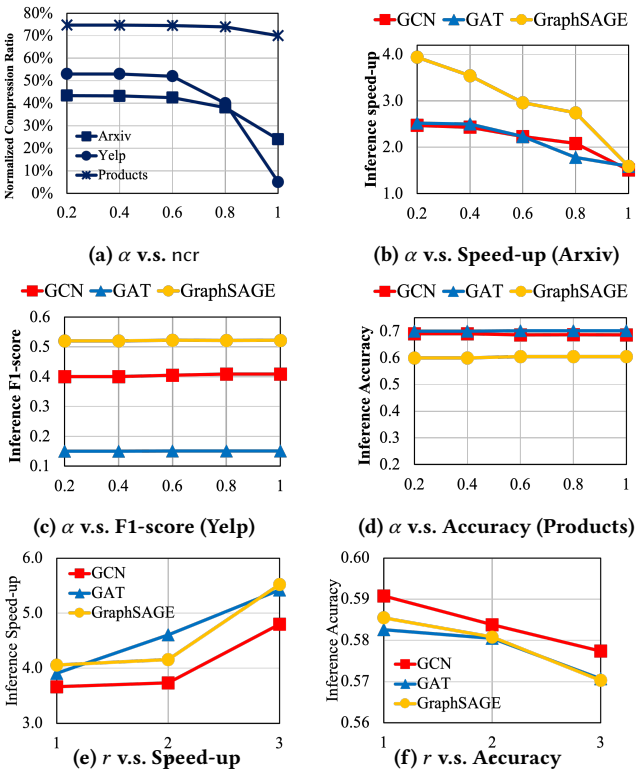

(a) α v.s. ncr (b) α v.s. Speed-up (Arxiv)

(c) α v.s. F1-score (Yelp) (d) α v.s. Accuracy (Products)

(e) r v.s. Speed-up (f) r v.s. Accuracy

Figure 9: Varying α and r in (α, r)-SPGC

Varying α. Fixing $r = 1$, we varied α from 0.2 to 1, and report the results in Figs. 9(a) to 9(d). It tells us the followings.

(1) As α is increased from 0.2 to 1, ncr drops as illustrated in Fig. 9(a). As expected, larger α makes it harder for (α, r)-SPGC to merge nodes that are less close in their representations, leaving more nodes in compressed graphs, hence worsening compression ratio. (2) Consistently for larger α and for all the three GNNs classes, it is harder for (α, r)-SPGC to improve their inference efficiency due to larger compressed structure G_c as shown in Fig. 9(b). We observe consistent results for Yelp and Products (not shown; see [1]).

(3) As α increases, the F1-score (resp. accuracy) of the inference results over **Yelp** (resp. **Products**) remains insensitive, as shown in Figs.9(c) (resp. 9(d)). Our observation over **Arxiv** remains consistent, and we omitted it due to limited space. This indicates that (α, r)-SPGC does not lose much on the quality of the inference while significantly improved inference efficiency.

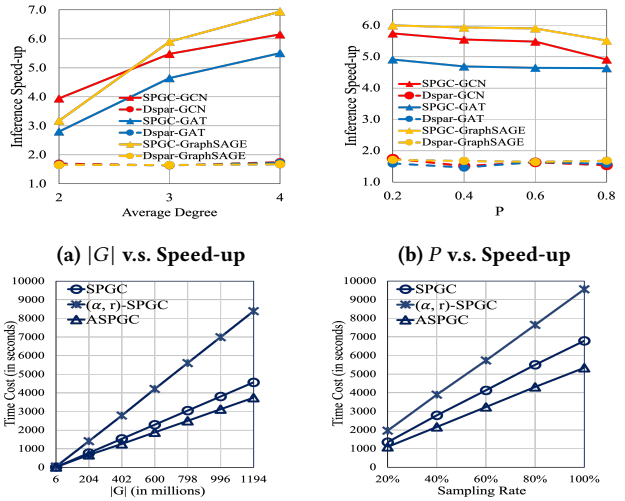

(a) $|G|$ v.s. Speed-up (b) P v.s. Speed-up

(c) Scalability: WS-MAG (d) Scalability: ogbn-papers100M

Figure 10: Efficiency and Scalability at Billion scale.

Varying r in SPGC. Fixing $\alpha = 0.25$, we vary r from 1 to 3 over **Arxiv** and report the result in Fig. 9. (1) As r is varied from 1 to 3, the inference speed-up achieved by $(0.25, r)$-SPGC for all GNNs classes notably increased. Indeed, larger r allows (α, r)-SPGC to find and merge more node pairs with equivalent embeddings, which may not be direct neighbors of another pair in the $(\alpha, 1)$-relation. (2) As r increases, the inference accuracy for all GNNs classes slightly drops, and all within a small range of 0.02. This demonstrates that (α, r)-SPGC is capable of preserving inference accuracy while increasing r in trading for larger speed up.

Exp-3: Efficiency & Scalability. We next evaluate the compression and inference costs of SPGC and its variants.

Average degree and "small-world" effect v.s. Inference Speed-up. We simulate **WS-MAG** based on (K, P) Watts-Strogatz algorithm [49] with fixed $|V| = 2M$ from **MAG240M** dataset. K and P represent average degree and re-wiring probability respectively. As P goes up, the less "small-world" (*i.e.,* more random) the graphs become.

Fixing $P = 0.6$, we vary the average degree from 2 to 4. Fig. 10(a) shows that the inference speed-up for different GNNs types. (1) As the average degree of the graph increases from 2 to 4, inference speed-up achieved by SPGC increases approximately linearly from 3.0×-4.0× to 5.0×-7.0×. (2) Inference speed-up achieved by DSpar remains relatively stable (1.6×-1.8×) and smaller than SPGC. As the average degree goes up, SPGC may takes the advantage of higher density that makes nodes more likely to be merged, resulting in a smaller G_c and greater speed-up.

Next, fixing $|G| = 8M$, we increase P from 0.2 to 0.8. We have the following observations. (1) As P increases, the inference speed-up achieved by SPGC on three GNNs exhibits a slight drop, albeit still much larger than the speed-up by DSpar. (2) Inference speed-ups achieved by DSpar on three GNNs stay flat without even breaking 2× speed-up. This indicates that SPGC may perform better to compress "small-world" graphs, which are common in real-world networks such as social networks [7, 49].

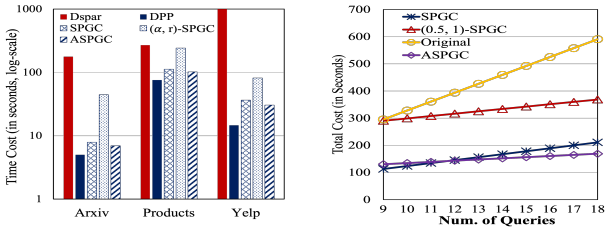

(a) Compression Cost **(b) Total Cost**

Figure 11: Time Cost Analysis (left: compression cost comparison; right: total cost for varying query loads).

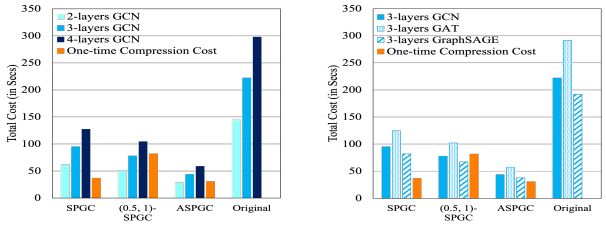

(a) Varying GNN Layer Number **(b) Varying GNN class**

Figure 12: Total Cost: Varying GNN Layers and Class.

Scalability Tests. We evaluated the scalability of SPGC, (α, r)-SPGC and ASPGC on large-scale graphs up to the billion-scale using **WS-MAG** (see Section 6.1) and **ogbn-papers100M** [27].

We varied the size of **WS-MAG** by enlarging its sizes with a generator that uses Watts-Strogatz algorithm [49]. We varied the graph size $|G|$ from 6 million to 1.194 billion in a sequence. As shown in Fig. 10(c), We observe the following. (1) SPGC, ASPGC, and (α, r)-SPGC scale well with $|G|$. This is consistent with our time cost analysis on the one-time compression in Table. 1. (2) (α, r)-SPGC incurs a relatively higher and more sensitive compression cost due to the overhead of computing and refining the (α, r)-relation. The overall cost remains to grow linearly with graph size, while its additional overhead yields improved compression ratios, leading to higher inference speed-ups on the compressed graphs G_c.

We also conducted a scalability test using a real-world billion-scale **ogbn-papers100M** (as illustrated in Fig. 10(d)). We generate subgraphs from 20% to 100% of the size of **ogbn-papers100M** using a subgraph sampling approach [16] and evaluate the compression costs of SPGC and its variants. As the sampling rate increases, we observe that SPGC, ASPGC, and (α, r)-SPGC all scale well. (α, r)-SPGC incurs relatively higher and more sensitive compression costs, consistent with its performances on **WS-MAG**.

Compression Cost Comparison. We compare the one-time compression cost induced by the SPGC or its variants with the cost of DSpar as shown in Fig. 11(a). Note that the compression cost of SPGC and its variants include the time cost of DPP. We have the following discoveries. (1) SPGC, (α, r)-SPGC and ASPGC outperforms DSpar on all the real-world datasets. In particular, SPGC are 95.52% and 58.45% times faster than DSpar over **Arxiv** and **Products**, respectively. For **Yelp**, DSpar does not run to completion after 1,000 seconds. (2) For SPGC and its variants, ASPGC outperforms SPGC and (α, r)-SPGC by 12.71% and 68.34% on average. This is because ASPGC exploits data locality of GNN inference and perform compression up to certain hops of the anchored nodes.

Total cost: Varying query workloads. We compare the total costs of SPGC, (α, r)-SPGC, ASPGC to the total cost induced by the original G. For SPGC and its variants, the total cost is the sum of "one-time" compression cost and the inference cost for all inference queries; for "Original", it refers to the inference time on the original graph. Varying the number of inference queries from 9 to 18 (each with fixed number of test node $|V_T| = 0.05|V|$), Fig. 11(b) reports the total cost of a 3-layers GCN on **Products**. (1) All SPGC methods are able to reduce the original inference cost, and all can improve the inference efficiency better with larger amount of queries. (2) Among SPGC variants, for queries with given anchored nodes, ASPGC performs best in improving the inference efficiency. We also observe that (α, r)-SPGC is more sensitive to different datasets, compared with SPGC and ASPGC. Indeed, SPGC enforces rigidly embedding equivalence, while ASPGC benefits most from data locality from anchored node set of fixed size, hence both are less sensitive to datasets. (α, r)-SPGC on the other hand is most adaptive for tunable trade-off between inference speed up and compression ratio, depending on the heterogeneity of node embeddings.

Total cost: varying GNN complexity. We next investigate the impact of GNN model complexity, in terms of number of layers, and types. Using GCNs, we varying the number of layers L from 2 to 4 (We denote GCN_2 as the 2-layers GCN; similarly for others), and a query workload of 3 different inference queries with the same size $|V_T|$. Fig. 12(a) tells us that as the number of layers increase, the inference cost grows as expected. ASPGC outperforms SPGC and $(0.5, 1)$-SPGC consistently for all layers in compression efficiency.

Fixing query workload size as 3, and the number of layers as 3 for all GNN types, we report the total cost over GCN, GAT and GraphSAGE. As shown in Fig. 12(b), all three types of GNNs consistently and significantly benefit from SPGC, with an improvement of inference time by a factor of 2.21, 2.36 and 2.11, respectively; similar for (α, r)-SPGC and ASPGC. The one-time compression costs are consistent with our prior observations in Fig. 12(a). In general, GAT may benefit most from inference-friendly compression. A possible reason is that the memoization effectively cached its additional edge weights as part of the auxiliary information, which are a source of overhead for inference over G.

7 CONCLUSION

We have proposed IFGC, a graph compression scheme to generate compressed graphs that can be directly queried by GNN inference to obtain the original output. We have introduced three practical specifications of IFGC, SPGC, for inference without decompression, and configurable (α, r)-compression, to enable trade-off between compression ratio and inference cost, and anchored SPGC, that preserves inference results for targeted nodes. Our theoretical analysis and experimental study have verified that IFGC can significantly accelerate the inference on real-world graphs with a small loss on inference accuracy. A future topic is to extend our compression scheme to accelerate GNN training. Another topic is to develop a parallel compression scheme for large graphs.

ACKNOWLEDGMENTS

Che and Wu are supported by NSF OAC-2104007. Fan and Wu are supported by DE-NA0004104.

REFERENCES

[1] 2025. Full version. http://github.com/Yangxin666/SPGC/blob/main/SPGC_full.pdf

[2] Arman Ahmed, Sajan K Sadanandan, Shikhar Pandey, Sagnik Basumallik, Anurag K Srivastava, and Yinghui Wu. 2022. Event Analysis in Transmission Systems Using Spatial Temporal Graph Encoder Decoder (STGED). *IEEE Transactions on Power Systems* (2022).

[3] Francesco Andreuzzi. 2021. BisPy: Bisimulation in Python. *Journal of Open Source Software* (2021).

[4] Adnan Aziz, Vigyan Singhal, Felice Balarin, Robert K Brayton, and Alberto L Sangiovanni-Vincentelli. 1994. Equivalences for fair kripke structures. In *Automata, Languages and Programming: 21st International Colloquium, ICALP 94 Jerusalem, Israel, July 11–14, 1994 Proceedings 21*. 364–375.

[5] Waïss Azizian and Marc Lelarge. 2021. Expressive Power of Invariant and Equivariant Graph Neural Networks. In *ICLR*.

[6] Pablo Barceló, Egor V Kostylev, Mikaël Monet, Jorge Pérez, Juan L Reutter, and Juan-Pablo Silva. 2020. The expressive power of graph neural networks as a query language. *SIGMOD Record* (2020).

[7] Danielle S Bassett and Edward T Bullmore. 2017. Small-world brain networks revisited. *The Neuroscientist* (2017).

[8] Maciej Besta and Torsten Hoefler. 2024. Parallel and distributed graph neural networks: An in-depth concurrency analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2024).

[9] Jeroen Bollen, Jasper Steegmans, Jan Van den Bussche, and Stijn Vansummeren. 2023. Learning graph neural networks using exact compression. In *Proceedings of the 6th Joint Workshop on Graph Data Management Experiences & Systems (GRADES) and Network Data Analytics (NDA)*. 1–9.

[10] Stephen P Borgatti and Martin G Everett. 1992. Notions of position in social network analysis. *Sociological methodology* (1992).

[11] Linfeng Cao, Haoran Deng, Yang Yang, Chunping Wang, and Lei Chen. 2024. Graph-Skeleton:¯ 1% Nodes are Sufficient to Represent Billion-Scale Graph. In *Proceedings of the ACM on Web Conference 2024*. 570–581.

[12] Ming Chen, Zhewei Wei, Bolin Ding, Yaliang Li, Ye Yuan, Xiaoyong Du, and Ji-Rong Wen. 2020. Scalable graph neural networks via bidirectional propagation. *NeurIPS* (2020).

[13] Tianlong Chen, Yongduo Sui, Xuxi Chen, Aston Zhang, and Zhangyang Wang. 2021. A unified lottery ticket hypothesis for graph neural networks. In *ICML*.

[14] Francisco Claude and Gonzalo Navarro. 2007. A fast and compact Web graph representation. In *International Symposium on String Processing and Information Retrieval*. 118–129.

[15] Agostino Dovier, Carla Piazza, and Alberto Policriti. 2001. A fast bisimulation algorithm. In *Proceedings of the 13th International Conference on Computer Aided Verification*.

[16] Pierre Dupont, Jérôme Callut, Grégoire Dooms, Jean-Noël Monette, Yves Deville, and B Sainte. 2006. Relevant subgraph extraction from random walks in a graph. *Universite Catholique de Louvain, UCL/INGI, Number RR* 7 (2006).

[17] Wenqi Fan, Yao Ma, Qing Li, Yuan He, Eric Zhao, Jiliang Tang, and Dawei Yin. 2019. Graph neural networks for social recommendation. In *WWW*.

[18] Yangxin Fan, Xuanji Yu, Raymond Wieser, David Meakin, Avishai Shaton, Jean-Nicolas Jaubert, Robert Flottemesch, Michael Howell, Jennifer Braid, et al. 2023. Spatio-Temporal Denoising Graph Autoencoders with Data Augmentation for Photovoltaic Data Imputation. *SIGMOD* (2023).

[19] Matthias Fey and Jan E. Lenssen. 2019. Fast Graph Representation Learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*.

[20] Xinyi Gao, Wentao Zhang, Yingxia Shao, Quoc Viet Hung Nguyen, Bin Cui, and Hongzhi Yin. 2022. Efficient Graph Neural Network Inference at Large Scale. *arXiv preprint arXiv:2211.00495* (2022).

[21] Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2018. Allennlp: A deep semantic natural language processing platform. *arXiv preprint arXiv:1803.07640* (2018).

[22] Floris Geerts. 2023. A Query Language Perspective on Graph Learning. In *PODS*.

[23] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. *NeurIPS* (2017).

[24] Xueting Han, Zhenhuan Huang, Bang An, and Jing Bai. 2021. Adaptive transfer learning on graph neural networks. In *KDD*.

[25] Mohammad Hashemi, Shengbo Gong, Juntong Ni, Wenqi Fan, B Aditya Prakash, and Wei Jin. 2024. A Comprehensive Survey on Graph Reduction: Sparsification, Coarsening, and Condensation. *arXiv preprint arXiv:2402.03358* (2024).

[26] Weihua Hu, Matthias Fey, Hongyu Ren, Maho Nakata, Yuxiao Dong, and Jure Leskovec. 2021. OGB-LSC: A Large-Scale Challenge for Machine Learning on Graphs. *NeurIPS* (2021).

[27] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. 2020. Open graph benchmark: Datasets for machine learning on graphs. *NeurIPS* (2020).

[28] Zengfeng Huang, Shengzhong Zhang, Chong Xi, Tang Liu, and Min Zhou. 2021. Scaling up graph neural networks via graph coarsening. In *SIGKDD*.

[29] Ahmad Maroof Karimi, Yinghui Wu, Mehmet Koyuturk, and Roger H French. 2021. Spatiotemporal graph neural network for performance prediction of photovoltaic power systems. In *AAAI*.

[30] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).

[31] Manoj Kumar, Anurag Sharma, Shashwat Saxena, and Sandeep Kumar. 2023. Featured graph coarsening with similarity guarantees. PMLR.

[32] N Jesper Larsson and Alistair Moffat. 2000. Off-line dictionary-based compression. *Proc. IEEE* 88, 11 (2000), 1722–1732.

[33] Dongyue Li, Tao Yang, Lun Du, Zhezhi He, and Li Jiang. 2021. AdaptiveGCN: Efficient GCN through adaptively sparsifying graphs. In *CIKM*.

[34] Xin Liu, Mingyu Yan, Lei Deng, Guoqi Li, Xiaochun Ye, Dongrui Fan, Shirui Pan, and Yuan Xie. 2022. Survey on graph neural network acceleration: An algorithmic perspective. *arXiv preprint arXiv:2202.04822* (2022).

[35] Yajiong Liu, Yanfeng Zhang, Qiange Wang, Hao Yuan, Xin Ai, and Ge Yu. 2025. NeutronSketch: An in-depth exploration of redundancy in large-scale graph neural network training. *Knowledge-Based Systems* 309 (2025), 112786.

[36] Zirui Liu, Kaixiong Zhou, Zhimeng Jiang, Li Li, Rui Chen, Soo-Hyun Choi, and Xia Hu. 2023. DSpar: An Embarrassingly Simple Strategy for Efficient GNN Training and Inference via Degree-Based Sparsification. *TMLR* (2023).

[37] Francois Lorrain and Harrison C White. 1971. Structural equivalence of individuals in social networks. *The Journal of mathematical sociology* 1, 1 (1971), 49–80.

[38] Andreas Loukas and Pierre Vandergheynst. 2018. Spectrally approximating large graphs with smaller graphs. In *ICML*.

[39] Andrew Kachites McCallum, Kamal Nigam, Jason Rennie, and Kristie Seymore. 2000. Automating the construction of internet portals with machine learning. *Information Retrieval* (2000).

[40] Hongwu Peng, Deniz Gurevin, Shaoyi Huang, Tong Geng, Weiwen Jiang, Orner Khan, and Caiwen Ding. 2022. Towards sparsification of graph neural networks. In *2022 IEEE 40th International Conference on Computer Design (ICCD)*. 272–279.

[41] Hao Peng, Hongfei Wang, Bowen Du, Md Zakirul Alam Bhuiyan, Hongyuan Ma, Jianwei Liu, Lihong Wang, Zeyu Yang, Linfeng Du, Senzhang Wang, et al. 2020. Spatial temporal incidence dynamic graph neural networks for traffic flow forecasting. *Information Sciences* (2020).

[42] Patrick Reiser, Marlen Neubert, André Eberhard, Luca Torresi, Chen Zhou, Chen Shao, Houssam Metni, Clint van Hoesel, Henrik Schopmans, Timo Sommer, et al. 2022. Graph neural networks for materials science and chemistry. *Communications Materials* (2022).

[43] Vıctor Garcia Satorras, Emiel Hoogeboom, and Max Welling. 2021. E (n) equivariant graph neural networks. In *ICML*.

[44] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. 2008. The graph neural network model. *IEEE transactions on neural networks* (2008).

[45] Weijing Shi and Raj Rajkumar. 2020. Point-GNN: Graph Neural Network for 3D Object Detection in a Point Cloud. In *CVPR*.

[46] Shyam A Tailor, Javier Fernandez-Marques, and Nicholas D Lane. 2020. Degree-quant: Quantization-aware training for graph neural networks. *arXiv preprint arXiv:2008.05000* (2020).

[47] Petar Veličković. 2023. Everything is connected: Graph neural networks. *Current Opinion in Structural Biology* (2023).

[48] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, Yoshua Bengio, et al. 2017. Graph attention networks. *stat* (2017).

[49] Duncan J Watts and Steven H Strogatz. 1998. Collective dynamics of 'small-world'networks. *nature* (1998).

[50] Fengli Xu, Quanming Yao, Pan Hui, and Yong Li. 2021. Automorphic equivalence-aware graph neural network. *NeurIPS* 34 (2021), 15138–15150.

[51] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2019. How powerful are graph neural networks?. In *ICLR*.

[52] Hanqing Zeng, Muhan Zhang, Yinglong Xia, Ajitesh Srivastava, Andrey Malevich, Rajgopal Kannan, Viktor Prasanna, Long Jin, and Ren Chen. 2021. Decoupling the depth and scope of graph neural networks. *NeurIPS* (2021).

[53] Hanqing Zeng, Hongkuan Zhou, Ajitesh Srivastava, Rajgopal Kannan, and Viktor Prasanna. 2019. Graphsaint: Graph sampling based inductive learning method. *arXiv preprint arXiv:1907.04931* (2019).

[54] Muhan Zhang and Yixin Chen. 2018. Link prediction based on graph neural networks. *NeurIPS* (2018).

[55] Cheng Zheng, Bo Zong, Wei Cheng, Dongjin Song, Jingchao Ni, Wenchao Yu, Haifeng Chen, and Wei Wang. 2020. Robust graph representation learning via neural sparsification. In *ICML*.

[56] WANG Zhili, DI Shimin, CHEN Lei, and ZHOU Xiaofang. 2024. Search to fine-tune pre-trained graph neural networks for graph-level tasks. In *ICDE*.

[57] Hongkuan Zhou, Ajitesh Srivastava, Hanqing Zeng, Rajgopal Kannan, and Viktor Prasanna. 2021. Accelerating large scale real-time GNN inference using channel pruning. *arXiv preprint arXiv:2105.04528* (2021).