

Graph Coarsening via Supervised Granular-Ball for Scalable Graph Neural Network Training

Shuyin Xia^{1,2,3}, Xinjun Ma^{2,3}, Zhiyuan Liu^{2,3}, Cheng Liu^{2,3}, Sen Zhao^{1,2,3*}, Guoyin Wang⁴,

¹Key Laboratory of Big Data Intelligent Computing

²Chongqing Key Laboratory of Computational Intelligence

³Chongqing University of Posts and Telecommunications, Chongqing, China

⁴National Center for Applied Mathematics in Chongqing, Chongqing Normal University, Chongqing 401331, China
{xiasy,zhaoseng}@cqupt.edu.cn; {s220201072,s230231079,s230201067}@stu.cqupt.edu.cn

Abstract

Graph Neural Networks (GNNs) have demonstrated significant achievements in processing graph data, yet scalability remains a substantial challenge. To address this, numerous graph coarsening methods have been developed. However, most existing coarsening methods are training-dependent, leading to lower efficiency, and they all require a predefined coarsening rate, lacking an adaptive approach. In this paper, we employ granular-ball computing to effectively compress graph data. We construct a coarsened graph network by iteratively splitting the graph into granular-balls based on a purity threshold and using these granular-balls as super vertices. This granulation process significantly reduces the size of the original graph, thereby greatly enhancing the training efficiency and scalability of GNNs. Additionally, our algorithm can adaptively perform splitting without requiring a predefined coarsening rate. Experimental results demonstrate that our method achieves accuracy comparable to training on the original graph. Noise injection experiments further indicate that our method exhibits robust performance. Moreover, our approach can reduce the graph size by up to 20 times without compromising test accuracy, substantially enhancing the scalability of GNNs. The code can be found in <https://github.com/mxjun17/Granular-Ball-Node-Classification>.

Introduction

In recent years, graph neural networks (GNNs) have shown significant potential across various domains, including social network analysis, drug discovery, finance, and material science (Fout et al. 2017; Ying et al. 2018; Li and Goldwasser 2019; Zhou et al. 2018; Wang 2017; Feng et al. 2022; Xia, Wang, and Gao 2023). The increasing prevalence of large-scale graph data offers richer information and larger training sets for learning algorithms. However, this also amplifies the challenge of processing such data, demanding substantial computational resources and escalating costs, particularly during model training and parameter tuning (Huang et al. 2021). A practical solution is to simplify or reduce the graph, which not only accelerates GNNs but




Figure 1: Comparing original graph training with coarsening training. The coarsening approach reduces graph complexity while preserving key structural and label information, improving graph neural network training efficiency.

also enhances graph data analysis tasks like storage, visualization, and retrieval.

Training graph neural networks (GNNs) often incurs significant computational costs. To mitigate this, dataset condensation or distillation has been explored, constructing smaller synthetic datasets for training (Jin et al. 2021, 2022). However, these methods may demand high network performance and require precise compression control to be effective. Graph coarsening offers a solution by reducing large graphs into smaller structures while preserving key attributes, based on algebraic multigrid theory. Widely used in graph partitioning, summarization, and machine learning (Stüben 2001; Chen, Saad, and Zhang 2022; Dhillon, Guan, and Kulis 2007; Liu et al. 2018; Shuman, Faraji, and Vandergheynst 2015), coarsening methods often focus on preserving spectral features (Wilson and Zhu 2008; Loukas and Vandergheynst 2018; Loukas 2019; Jin, Loukas, and JaJa 2020; Imre et al. 2020). Despite their benefits, the high memory and computational costs of spectral methods can limit their scalability. While applied in deep learning (Xie et al. 2020b; Fahrbach et al. 2020; Huang et al. 2021; Pang, Zhao,

*Corresponding author.

and Li 2021; Kumar et al. 2023; Zhang et al. 2022; Dickens et al. 2024), coarsening faces challenges like excessive information loss, which can degrade model performance.

Figure 1 illustrates the graph coarsening training process, highlighting the critical importance of effective coarsening, as it inevitably leads to the loss of some original graph information. While training-dependent graph coarsening methods like CMGC(Dickens et al. 2024) and GCOND(Jin et al. 2021) achieve high accuracy, they suffer from inefficiency and prolonged training times. In contrast, preprocessing-based methods like SCAL(Huang et al. 2021) offer better efficiency but lack the guidance of label information, resulting in slightly lower accuracy. Moreover, most existing graph coarsening techniques are not adaptive and require a predefined coarsening ratio.

To address these challenges, we propose a Supervised Granular-Ball Graph Coarsening (SGBGC) method, which optimizes information preservation by integrating node label information with graph structural characteristics and introduces an adaptive coarsening mechanism. This adaptability enhances the flexibility of graph coarsening. Our approach employs an innovative granular-ball generation mechanism, initially partitioning granular-balls based on node label similarity, followed by granularity adjustments to refine these granular-balls into high-quality super nodes. We conducted extensive experiments on real datasets to demonstrate the effectiveness of the SGBGC method. Our approach outperforms all existing graph coarsening architectures for node classification and achieves state-of-the-art results on five benchmark datasets for node classification. Detailed analyses and robustness studies further illustrate the superiority of our methods. In summary, the main contributions are:

- We propose a novel Supervised Granular-Ball Graph Coarsening (SGBGC) method that adaptively utilizes granular-balls of varying sizes to represent the sample space, maintaining crucial graph structural properties while enhancing the scalability of Graph Neural Network (GNN) models.
- SGBGC reduces time complexity by avoiding the need to calculate distances between all node pairs within the graph. Additionally, Our method has demonstrated superior noise resistance in robustness experiments, outperforming existing graph coarsening techniques.
- Extensive testing and comparisons on real-world datasets have confirmed the effectiveness and accuracy of SGBGC. Our method has achieved higher precision in experiments, validating its efficacy for GNN training and inference.

Related Works

Graph Coarsening and Condensation

During the training of Graph Neural Networks (GNNs), due to the typically high computational costs, methods such as graph coarsening and graph condensation are increasingly gaining attention. These methods aim to create smaller synthetic datasets for training GNNs, thereby reducing costs to some extent. Among these, SCAL (Huang et al. 2021) demonstrated scalable training of GNNs through graph

coarsening techniques, proving that even with a reduction of graph nodes to one-tenth of the original size, there was no significant impact on classification accuracy. FGC (Kumar et al. 2023) designed an innovative optimization-based graph coarsening framework, which takes the adjacency matrix and node features as inputs, and jointly learns the coarsened versions of the graph and feature matrices to ensure the desired properties while optimizing graph representation. (Jin et al. 2021, 2022) proposed the GCOND graph compression method, which not only reduces the number of nodes but also compresses node features and structural information in a supervised manner. CMGC (Dickens et al. 2024) focuses on preserving graph convolution operations while enhancing existing coarsening techniques by merging nodes with similar structures. It is also a training-dependent coarsening method, delivering state-of-the-art performance.

Granular-Ball Computing

Building on the theoretical foundations of traditional granularity computation and integrating the human cognitive mechanism of 'macro-first' (Chen 1982), Wang (Wang 2017) pioneered the concept of multi-granularity cognitive computation. Following this approach, Xia (Xia et al. 2019) proposed an efficient, robust and interpretable computational method known as granular-ball computing. Compared to traditional methods that input data at the finest granularity of a point, granular-ball computing uses granular-balls to cover and represent data and takes granular-balls as inputs, offering high efficiency, robustness. A granular-ball is defined by a point set covered by the ball, having a centroid and a radius in the standard formation, such as in granular-ball classification. Since its introduction, although the theory of granular-ball computing has not been developed for long, it has been widely applied across multiple domains of artificial intelligence, giving rise to various granular-ball generation techniques and related computational models. Specific examples include granular-ball classifiers (Xia et al. 2024b; Quadir and Tanveer 2024; Xia et al. 2024c), granular-ball clustering (Cheng et al. 2023; Xie et al. 2024c,b,a), granular-ball sampling methods (Xia et al. 2021), granular-ball rough set(Xia et al. 2023, 2020; Zhang et al. 2023), granular-ball three way decision(Yang et al. 2024; Xia et al. 2024a), and developments like granular-ball reinforcement learning (Liu et al. 2024). Furthermore, several applications have demonstrated the effectiveness of granular-ball representation, such as granular-ball representation in text adversarial defense(Wang et al. 2024a), label noise combating(Wang et al. 2024b; Dai et al. 2024).

Preliminaries

Graph Coarsening

We represent a graph with N nodes and M edges as $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$, where \mathcal{V} is the set of nodes, \mathcal{E} is the set of edges, \mathbf{X} represents the features of the nodes, and \mathbf{Y} represents the set of node labels. All graph data in this paper are undirected graphs. We use $\mathbf{A} \in \{0, 1\}^{N \times N}$ to denote the adjacency matrix of \mathcal{G} , i.e., the (i, j) -th entry in \mathbf{A} is 1 if and only if there is an edge between v_i and v_j . Given an




Figure 2: The overview of the SGBGC architecture. Our method consists of four stages: coarse partitioning, fine-grained splitting, granular-ball graph construction, and GNN training on the coarsened graph. In the coarse partitioning stage, the graph is divided into \sqrt{N} clusters, where N is the total number of nodes, with each cluster evenly divided based on label categories. The fine-grained splitting stage divides clusters further based on granular-ball quality until the specified threshold is reached. Finally, the granular-ball graph is constructed and used for GNN training to produce coarsened embeddings.

undirected graph \mathcal{G} , graph coarsening seeks to find a coarser graph $\mathcal{G}' = (\mathcal{V}', \mathcal{E}', \mathbf{X}')$ that approximates \mathcal{G} , where \mathcal{V}' is the set of super nodes obtained after coarsening, \mathcal{E}' is the set of super edges, \mathbf{X}' represents the features of the super nodes, and \mathbf{A}' is the adjacency matrix of the coarser graph. The node count in \mathcal{G}' is $n < N$ and the edge count is $m < M$. By computing a partition $\mathcal{P} = \{C_1, C_2, \dots, C_n\}$, we can construct \mathcal{G}' , where each cluster C_i corresponds to a super node in \mathcal{G}' , and the connections between these super nodes C_i and C_j form the super edges, i.e., a tightly connected subset of nodes from the original graph.

Traditional Granular-Ball Computing

The core concept of granular-ball computing involves partitioning the sample space into manageable segments using granular-balls. These granular-balls serve as the basic units for learning and facilitate multi-granularity representation and computation. The traditional granular-ball is defined by a center point and a radius, encompassing all points within the specified radius. Formally, a granular-ball centered at a point $\mathbf{c} \in \mathbb{R}^d$ with radius r is defined as:

$$\mathcal{GB} = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x} - \mathbf{c}\| \leq r\}, \quad (1)$$

where $\|\mathbf{x} - \mathbf{c}\|$ denotes the Euclidean distance between any point $\mathbf{x} \in \mathbb{R}^d$ and the center $\mathbf{c} \in \mathbb{R}^d$. The radius r is determined based on the distribution of data points within the cluster, providing a comprehensive representation of the data within that region. To ensure that the entire dataset $\mathcal{D} \subseteq \mathbb{R}^d$ is covered, the dataset is partitioned into n granular-balls. The collection of these granular-balls, denoted as $\mathcal{GB}_s = \{\mathcal{GB}_1, \mathcal{GB}_2, \dots, \mathcal{GB}_n\}$, is used such that $\mathcal{D} \subseteq \bigcup_{i=1}^n \mathcal{GB}_i$. This ensures that every data point in \mathcal{D} is included in at

least one granular-ball, facilitating complete coverage of the sample space. By using granular-balls instead of individual data points, granular computing simplifies data representation and reduces computational complexity. This structure forms a cluster that can be effectively used to represent spatial relationships in data.

Methodology

Overview

Figure 2 presents an overview of our proposed SGBGC framework. In contrast to traditional granular-ball computing, which primarily focuses on abstract data granulation, in graph-based scenarios, a granular-ball \mathcal{GB} integrates both structural and feature-specific properties of graph data.

Definition 1: Granular-Ball on Graph

Consider an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{X}, \mathbf{A})$, where \mathcal{V} is the set of nodes with $|\mathcal{V}| = N$ representing the total number of nodes, \mathcal{E} is the set of edges with $|\mathcal{E}| = M$ indicating the total number of edges, \mathbf{X} is the matrix of node features, and \mathbf{A} is the adjacency matrix of the graph. A granular-ball \mathcal{GB} is defined as a connected subgraph of \mathcal{G} , represented by:

$$\mathcal{GB} = (\mathcal{V}_{gb}, \mathcal{E}_{gb}, \mathbf{X}_{gb}, \mathbf{A}_{gb}) \quad (2)$$

Here, $\mathcal{V}_{gb} \subseteq \mathcal{V}$ represents the subset of nodes in the granular-ball, $\mathcal{E}_{gb} \subseteq \mathcal{E}$ represents the subset of edges among nodes within \mathcal{V}_{gb} , $\mathbf{X}_{gb} \subseteq \mathbf{X}$ represents the subset of features corresponding to \mathcal{V}_{gb} , and \mathbf{A}_{gb} represents the adjacency matrix corresponding to \mathcal{V}_{gb} . The granular-ball \mathcal{GB} must be a connected subgraph, which means for any two nodes $v_i, v_j \in \mathcal{V}_{gb}$, there does not exist a situation where $d(v_i, v_j) = \infty$. Here, $d(v_i, v_j)$ represents the shortest path

distance (in terms of the number of edges) between nodes v_i and v_j in the graph. The center node c of a granular-ball is defined as the node within \mathcal{V}_{gb} with the maximum degree:

$$c = \arg \max_{v \in \mathcal{V}_{gb}} \deg_{gb}(v), \quad (3)$$

where $\deg_{gb}(v)$ is the degree of node v within the granular-ball \mathcal{GB} , defined as the number of edges connected to v within \mathcal{E}_{gb} .

Definition 2: The Purity of the Granular-Ball

The purity T of a granular-ball \mathcal{GB} is defined as the ratio of the count of the most frequent label to the total number of nodes within the ball. This ratio quantifies the uniformity of labeling within \mathcal{GB} . Mathematically, it is given by:

$$T = \frac{\max(n_i)}{n}, \quad i = 1, 2, \dots, k, \quad (4)$$

where n_i is the number of nodes with label i in \mathcal{GB} , n is the total number of nodes in \mathcal{GB} , and k is the number of distinct labels in the dataset. Here, T falls strictly within the range $\frac{1}{k} < T \leq 1$, where $T = 1$ indicates that all nodes in \mathcal{GB} share the same label, representing maximum purity. The lower bound $\frac{1}{k}$ occurs when the node labels are distributed as evenly as possible across the k labels, representing the least purity achievable under the constraint that one label must be most frequent.

Problem Definition

In our framework, the original graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$ is transformed into a granular-ball graph, denoted as $\tilde{\mathcal{G}} = (\tilde{\mathcal{V}}, \tilde{\mathcal{E}}, \tilde{\mathbf{X}})$, by coarsening the graph into n granular-balls. Each granular-ball \mathcal{GB}_j corresponds to a super node in $\tilde{\mathcal{G}}$, forming the set $\tilde{\mathcal{V}} = \{\mathcal{GB}_1, \mathcal{GB}_2, \dots, \mathcal{GB}_n\}$ where n is the total number of super nodes. The goal is to obtain a coarsened graph $\tilde{\mathcal{G}}$ that reduces the complexity of the original graph while retaining essential structural and feature information. In this coarsened graph $\tilde{\mathcal{G}}$, the adjacency matrix $\tilde{\mathbf{A}}$ is defined such that an element $\tilde{\mathbf{A}}_{ij}$ equals 1 if there is an edge between super nodes \mathcal{GB}_i and \mathcal{GB}_j ; otherwise, $\tilde{\mathbf{A}}_{ij} = 0$. Each node v_i in the original graph has a label y_i , and in the node classification task, the classifier is initially aware of the labels of a subset of nodes \mathcal{V}_L . The goal of node classification is to infer the labels of nodes in $\mathcal{V} \setminus \mathcal{V}_L$ by learning a classification function from the coarsened graph $\tilde{\mathcal{G}}$, effectively utilizing the reduced graph structure to predict labels with reduced computational resources while maintaining or even enhancing prediction accuracy.

Supervised Granular-Ball Graph Coarsening

The SGBGC method comprises two main stages: initial coarse partitioning and fine-grained binary splitting. In the initial stage, the goal is to simplify the graph structure by establishing a uniform distribution of granular-balls based on node labels and their degree of connectivity. Starting with the highest degree nodes as centers, $\alpha = \sqrt{N}$ centers are chosen, where N is the total number of nodes. The value of \sqrt{N} is an empirical choice, referenced from previous work

(Xie et al. 2020a; Bezdek and Pal 1998), which has been shown to balance the trade-off between computational complexity and accuracy. These α centers are then evenly distributed among the k label categories, resulting in each category initially having α/k centers. This setup forms the preliminary granular-ball distribution,

$$\mathcal{GB}_{\text{init}} = \{\mathcal{GB}_1, \mathcal{GB}_2, \dots, \mathcal{GB}_\alpha\}, \quad (5)$$

which covers the entire graph dataset. The choice of high-degree nodes ensures that the initial centers are well-connected, facilitating effective coverage of the graph. Each \mathcal{GB}_i ($1 \leq i \leq \alpha$) in $\mathcal{GB}_{\text{init}}$ undergoes binary splitting in the fine-grained phase. For each \mathcal{GB}_i , the algorithm selects the two highest-degree nodes as centers and splits \mathcal{GB}_i into two new granular-balls, \mathcal{GB}_{c_1} and \mathcal{GB}_{c_2} , based on the shortest path criterion:

$$c_1 = \arg \max_{v \in \mathcal{V}_{gb_i}} \deg_{gb_i}(v), \quad c_2 = \arg \max_{v \in \mathcal{V}_{gb_i} \setminus \{c_1\}} \deg_{gb_i}(v). \quad (6)$$

The nodes in \mathcal{GB}_i are assigned to the nearest center c_1 or c_2 based on the shortest path:

$$v \in \mathcal{GB}_{c_1} \iff d(v, c_1) \leq d(v, c_2), \quad (7)$$

$$v \in \mathcal{GB}_{c_2} \iff d(v, c_2) < d(v, c_1). \quad (8)$$

The splitting continues iteratively until the purity of each child granular-ball meets a predefined threshold $T = 1$, meaning that each granular-ball must contain nodes with the same label. High purity T in a granular-ball indicates strong label consistency, preserving or enhancing the discriminative properties of the original graph. Conversely, low purity T suggests a heterogeneous label mix, potentially impairing model performance. Therefore, in preliminary experiments, the purity threshold is typically set high to encourage clear decision boundaries. This strict purity threshold ensures that the granular-balls are homogeneous in terms of node labels, which helps maintain label consistency and improves the performance of downstream classification tasks. Algorithm 3 details this process, dynamically adjusting granular-ball sizes to capture graph data complexity.

Constructing the Granular-Ball Graph The construction of a granular-ball graph, $\tilde{\mathcal{G}}$, involves a transformation from the set of granular-balls derived from the original graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$. Each granular-ball is represented as a super node in $\tilde{\mathcal{G}}$, which initially contains no nodes or edges, i.e., $\tilde{\mathcal{G}} = \emptyset$. Node indices from \mathcal{V} are associated with their respective granular-balls based on their labels and connectivity, defining a mapping $f : \mathcal{V} \rightarrow \{\mathcal{GB}_1, \mathcal{GB}_2, \dots, \mathcal{GB}_n\}$, where each \mathcal{GB}_i is a granular-ball. This mapping is used to establish edges between super nodes in $\tilde{\mathcal{G}}$ as follows:

$$(u, v) \in \mathcal{E} \Rightarrow (f(u), f(v)) \in \tilde{\mathcal{E}} \text{ if } f(u) \neq f(v). \quad (9)$$

This condition ensures that an edge is added between two different super nodes in $\tilde{\mathcal{G}}$ if the corresponding original nodes u and v belong to different granular-balls. This connects granular-balls that share a direct link in the original graph, avoiding self-loops and ensuring that $\tilde{\mathcal{G}}$ represents the higher-level structure of \mathcal{G} . This methodology enhances the representation of graph by focusing on broader connectivity patterns while preserving essential graph properties.

Algorithm 1: Training GNN with Granular-Ball Graph Coarsening

Input: Graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$, Labels \mathbf{Y}
Output: Trained weight matrix \mathbf{W}^*

- 1: Apply the SGBGC algorithm on \mathcal{G} , and output a partition \mathcal{P} .
 - 2: Construct the Granular-Ball Graph $\tilde{\mathcal{G}}$ using partition \mathcal{P} .
 - 3: Compute the feature matrix of $\tilde{\mathcal{G}}$, $\tilde{\mathbf{X}} = \mathcal{P}^\top \mathbf{X}$.
 - 4: Calculate the labels of $\tilde{\mathcal{G}}$, $\tilde{\mathbf{Y}} = \arg \max(\mathcal{P}\mathbf{Y})$.
 - 5: Train parameter \mathbf{W} to minimize the loss $\ell(\text{GNN}_{\tilde{\mathcal{G}}}(\mathbf{W}), \tilde{\mathbf{Y}})$.
 - 6: Obtain an optimal weight matrix \mathbf{W}^* .
 - 7: **Return** \mathbf{W}^* .
-

Training GNN with SGBGC

Next, we incorporate the coarsened graph into the GNN models for training. We denote a GNN model based on graph G as $\text{GNN}_G(\mathbf{W})$. Given a loss function ℓ , such as cross-entropy, the loss of model is expressed as $\ell(\text{GNN}_G(\mathbf{W}), \mathbf{Y})$. We train a GNN model using G_{gb} and the corresponding labels $\tilde{\mathbf{Y}}$, with the objective to minimize the loss function $\ell(\text{GNN}_{G_{\text{gb}}}(\mathbf{W}), \tilde{\mathbf{Y}})$ on the granular-ball coarsened graph, where \mathbf{W} represents the model parameters. Ultimately, the optimal parameters \mathbf{W}^* obtained from training on the coarsened graph are used in the GNN model on the original graph G for downstream tasks.

In the granular-ball coarsened graph, each node is a super node, corresponding to a cluster of nodes in the original graph. The feature vector of each super node is the mean of all node feature vectors in the cluster, i.e., $\tilde{\mathbf{X}} = \mathbf{P}^\top \mathbf{X}$. We also set each label of super node accordingly, i.e., $\tilde{\mathbf{Y}} = \mathbf{P}^\top \mathbf{Y}$. However, super nodes may contain nodes from multiple categories. In such cases, we select the dominant label, i.e., the category with the highest purity, as the label for the super node by performing an argmax operation on $\tilde{\mathbf{Y}}$. During GCN model training, we can define the granular-ball graph convolution as:

$$\mathbf{Z}_{\text{gb}} = \tilde{\mathbf{D}}_{\text{gb}}^{-1/2} \tilde{\mathbf{A}}_{\text{gb}} \tilde{\mathbf{D}}_{\text{gb}}^{-1/2} \tilde{\mathbf{X}}_{\text{gb}} \mathbf{W}^*, \quad (10)$$

where \mathbf{Z}_{gb} denotes the output features of the granular-ball graph nodes, $\tilde{\mathbf{A}}_{\text{gb}} = \tilde{\mathbf{A}}_{\text{gb}} + \tilde{\mathbf{I}}_{\text{gb}}$, with $\tilde{\mathbf{A}}_{\text{gb}}$ being the adjacency matrix of the granular-ball graph and $\tilde{\mathbf{I}}_{\text{gb}}$ an identity matrix to incorporate self-connections. $\tilde{\mathbf{D}}_{\text{gb}}$ is the degree matrix of $\tilde{\mathbf{A}}_{\text{gb}}$, $\tilde{\mathbf{X}}_{\text{gb}}$ are the input features of the granular-ball graph nodes, and \mathbf{W}_{gb} is the matrix of weights to be learned. Algorithm 1 details the use of granular-ball graph coarsening within the GNN framework.

Experiments

Experimental Setup

Datasets We evaluate our methods on five standard node classification datasets: Cora, Citeseer, Pubmed(Kipf and Welling 2016), Co-CS, and Co-Phy(Shchur et al. 2018). Each dataset features nodes represented as documents with




Figure 3: Comparison of time costs among different coarsening methods.

sparse bag-of-words feature vectors, and edges representing citation links between these documents.

Baselines To evaluate the effectiveness of SGBGC in graph coarsening, we compare our approach against three established baselines: FGC (Kumar et al. 2023), SCAL (Huang et al. 2021), Gcond (Jin et al. 2021), and CMGC (Dickens et al. 2024). FGC utilizes training to derive a coarsened graph, optimizing for feature and connection preservation. Gcond focuses on maintaining the spectral properties of the graph, beneficial for applications requiring eigenstructure preservation. CMGC aims to retain the graph convolution operations and improves the existing coarsening methods by merging structurally similar nodes. SCAL, similar to SGBGC, pre-processes the graph to produce a coarsened version by incorporating various coarsening techniques, specifically VNGC, VEGC(Loukas 2019), GSGC, and JCGC(Ron, Safro, and Brandt 2011).

Implementation Details All models were implemented using Python and PyTorch Geometric, with experiments conducted on an Intel(R) Xeon(R) W-2245 CPU @3.90GHz and an NVIDIA GeForce RTX 3090 GPU. The dataset splits were handled according to the methodologies described in references (Huang et al. 2021), and accuracy (ACC) was used as the evaluation metric. Results reported are averages over 20 runs, along with standard deviations to indicate variability. To ensure fairness in comparisons, our models utilized the same network architectures as the baselines. We uniformly use Adam optimizer with learning rates of 0.01 and set the weight decay to 0.0005. For the coarse GCN, the number of training epochs are 200 and the number of hidden layers are 64 and the early stopping is set to 10.

Experiment Results

Node Classification Performance In a transductive setting, using GCN as the base model, as shown in Table 1, it is evident that our SGBGC method not only maintains accuracy comparable to the original graph but also achieves significant improvements over existing unsupervised graph coarsening methods. The GCOND, FGC, and CMGC methods require training for coarsening, which is highly time-

Table 1: Performance comparison of SGBGC and other coarsening methods under different ratios (r) for node classification in GNNs.(mean±std%, the best results are bolded, the second results are underlined, and ”-” means out of memory).

DataSet	Ratio (r)	VNGC	VEGC	JCGC	GSGC	GCOND	FGC	CMGC	SGBGC
Cora	0.5	77.28±1.18	78.36±1.75	79.85±1.24	81.31±1.59	-	86.22±0.15	88.40±0.67	87.85±0.64
	0.3	78.33±1.78	76.97±1.86	77.79±1.84	80.05±0.86	76.49±0.71	85.39±0.27	<u>85.96±0.84</u>	86.29±0.79
	0.1	69.19±0.73	62.76±1.94	67.44±1.25	68.32±1.08	80.54±1.61	81.66±0.36	<u>82.50±0.93</u>	82.99±1.00
Citeseer	0.5	73.16±0.64	71.31±1.10	71.91±1.56	71.23±1.69	-	<u>76.28±0.21</u>	77.33±0.85	75.25±1.02
	0.3	70.12±1.16	42.32±2.08	67.58±1.12	66.13±2.03	71.26±3.01	73.16±0.27	74.85±0.92	75.73±0.88
	0.1	60.97±2.79	41.88±1.18	52.68±2.35	55.24±1.52	69.67±2.37	72.87±1.13	<u>73.52±1.01</u>	73.69±0.93
Pubmed	0.3	83.95±0.21	82.54±0.35	82.13±0.64	83.99±0.27	-	83.36±0.05	85.80±0.30	86.28±0.20
	0.1	81.55±0.28	82.49±0.50	81.83±0.54	82.96±0.34	-	82.44±0.07	<u>83.75±0.45</u>	84.04±0.36
	0.05	78.12±3.00	<u>82.48±0.42</u>	81.51±1.31	81.81±0.31	80.23±0.08	81.64±0.29	82.37±0.53	82.50±0.38
Co-CS	0.3	56.05±5.36	53.61±4.19	46.87±4.59	52.76±4.78	-	89.94±0.04	75.54±2.35	93.09±0.13
	0.1	37.67±2.62	33.69±2.96	33.14±0.36	39.79±0.96	-	<u>88.47±0.06</u>	72.57±3.04	92.22±0.22
	0.05	23.22±0.12	34.05±1.16	27.09±3.46	33.08±0.41	82.76±0.94	<u>83.85±1.35</u>	83.34±1.87	91.21±0.30
Co-Phy	0.3	90.76±2.04	90.66±2.77	86.33±2.79	89.02±3.87	-	89.94±0.15	96.30±0.13	96.35±0.09
	0.1	76.00±0.94	77.73±2.29	73.61±2.41	69.86±0.09	-	89.23±0.28	<u>95.76±0.16</u>	95.88±0.11
	0.05	72.18±1.76	76.99±1.90	72.42±1.50	69.84±0.15	85.06±0.33	81.91±0.15	<u>95.36±0.21</u>	95.61±0.14




Figure 4: Comparison of different noise rates on Citeseer with 0.3 coarsening ratio, including original graph and various coarsening methods.

consuming. Additionally, GCOND often runs out of memory capacity at higher coarsening rates. In contrast, SGBGC is less time-intensive while still ensuring accuracy across most coarsening rate settings. Notably, at a coarsening rate of 0.1, our method demonstrates a substantial performance boost, achieving an accuracy of 82.99%, which is markedly higher than that of other methods. This indicates that even under extreme graph compression, our method preserves high accuracy, showcasing robust information retention capabilities. In the Citeseer dataset, our method demonstrates superior performance across most tested coarsening rates. Specifically, at a coarsening rate of 0.1, our method achieved an accuracy of 73.6%, which represents a significant improvement over other methods. Although at a coarsening rate of 0.5 our method is slightly outperformed by CMGC,

Table 2: Comparison of classification performance with adaptive granular-ball coarsening rates. FULL represents performance on uncoarsened datasets.

Datasets	Cora	Citeseer	Pubmed	Co-CS	Co-Phy
Ratio (r)	0.45	0.49	0.42	0.41	0.36
SGBGC	87.57	75.38	86.29	93.28	96.39
FULL	88.38	76.48	86.40	93.78	96.65

the overall results still confirm the generalizability and robustness of our approach across different graph datasets. Results from the Pubmed dataset also support the superiority of our method, particularly at a coarsening rate of 0.3, where the accuracy reached 86.28%, showing minimal performance loss compared to the original graph. On the Co-CS and Co-Phy datasets, our method exhibited remarkable performance improvements. Especially in the Co-CS dataset, even at very low coarsening rates (such as 0.1 and 0.05), the accuracy of SGBGC exceeded 91.21%, far surpassing other methods. These results emphasize the powerful capability of our method in handling large-scale graph data with complex structures. By dynamically leveraging label information and a purity-based granular-ball splitting strategy, our method effectively enhances the quality of graph coarsening while also ensuring efficient computational performance.

Adaptive approach Existing graph coarsening methods often lack support for adaptive coarsening, requiring a pre-determined coarsening rate to proceed. This limitation can hinder their flexibility and scalability. In contrast, our SGBGC method introduces an adaptive coarsening mechanism, where the coarsening rate is not pre-specified but is instead determined dynamically through granular-ball splitting. This allows SGBGC to adjust autonomously, enhancing scalability and adaptability to various graph sizes. The




Figure 5: The memory usage of APPNP and granular-ball coarsened APPNP.




Figure 6: Parametric Analysis of Initial Centers

results in Table 2 demonstrate that under adaptive conditions, SGBGC achieves performance very close to that of the original graph, validating the effectiveness of our approach.

Computational Efficiency of SGBGC For a fair comparison of coarsening methods, we set the coarsening rates as follows: 0.3 for the Cora and Citeseer datasets, and 0.05 for the PubMed, Co-CS, and Co-Phy datasets. Figure 3 shows the performance results. On the Cora dataset, SGBGC required only 4.92 seconds, which is 51% faster than VNGC and 46% faster than VEGC. On the Citeseer dataset, SGBGC completed processing in 3.47 seconds, which is 66% faster than both VNGC and GSGC.

For larger datasets, the advantages of SGBGC become more pronounced. On the Co-CS dataset, SGBGC completed processing in 20.18 seconds, outperforming VNGC by 44% and GSGC by over 9 times. Similarly, on the Co-Phy dataset, SGBGC finished in 62.72 seconds, nearly 17 times faster than GSGC. These results clearly indicate that the SGBGC method consistently outperforms other coarsening methods in terms of speed, especially on larger datasets. The superior performance of SGBGC is likely due to its reliance on local computations (i.e., granular-ball subgraphs) rather than global computations, making it highly efficient for large-scale datasets. Due to its extremely slow processing time, Gcond and CMGC were excluded from our comparisons; for example, Gcond took nearly two hours and CMGC over a thousand seconds on the Cora dataset.

Memory Usage Figure 5 shows the memory usage of APPNP at different coarsening ratios. Since parameter space is negligible compared to input tensor size, coarse APPNP’s memory usage is roughly proportional to the coarsening ratio. Our method significantly reduces memory consumption by several orders of magnitude on medium to large datasets.

Robustness Analysis Under Different Noise Rates Figure 4 evaluates various coarsening methods under 5%, 10%, 15%, and 20% noise on the Citeseer dataset with a 0.3 coarsening ratio. SGBGC consistently outperformed other methods and the noisy original graph. At 5% noise, SGBGC achieved 75.20% accuracy, surpassing the original graph’s 73.21%. Even at higher noise levels, SGBGC maintained superior accuracy, significantly better than both the original graph and the next best method, VNGC. SGBGC’s robustness likely stems from granular-ball computing, which mitigates label noise through node clustering based on purity thresholds. This approach preserves structural and feature integrity, enhancing node classification accuracy. These results underscore SGBGC’s superior noise tolerance, making it a robust choice for real-world applications.

Parameter Analysis We experimented with different initial center numbers, as shown in Figure 6. The results indicate that our method maintains consistent performance across varying settings, demonstrating strong stability. Experimental results show that in most cases, when the number of initial centers is too large, the accuracy will be slightly improved, but the efficiency will be correspondingly slower. For example, on the Pubmed dataset, coarsening takes 29.7 seconds with \sqrt{N} centers, 35.0 seconds with $2\sqrt{N}$, and 78.4 seconds with $4\sqrt{N}$, with only minor accuracy gains. Thus, \sqrt{N} was chosen as the optimal value, informed by both our experiments and empirical evidence from the previous work (Xie et al. 2020a), which has been shown to balance the trade-off between computational complexity and accuracy.

Conclusion and Future Work

In this work, we introduce Supervised Granular-Ball Graph Coarsening (SGBGC), a novel method for supervised graph coarsening that significantly enhances the scalability of GNNs by compressing the original graph over tenfold without losing accuracy. SGBGC is an efficient and adaptive approach requiring minimal hyperparameters, and it can also operate in a non-adaptive form based on the coarsening ratio. Experiments show that SGBGC outperforms several advanced methods in both precision and efficiency, notably accelerating GNN training in node classification tasks. While our current architecture is focused on node classification as a downstream task, future work will explore other applications such as graph classification and link prediction to further broaden the applicability of our graph coarsening method.

Acknowledgments

We sincerely appreciate the dedication and efforts of everyone who contributed to this work. Thank you all for your hard work. This work was supported in part by the National Natural Science Foundation of China under Grant Nos. 62222601, 62450043, 62176033, 62221005 and 61936001, key cooperation project of Chongqing municipal education commission No.HZ2021008.

References

- Bezdek, J. C.; and Pal, N. R. 1998. Some new indexes of cluster validity. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 28(3): 301–315.
- Chen, J.; Saad, Y.; and Zhang, Z. 2022. Graph coarsening: from scientific computing to machine learning. *SeMA Journal*, 79(1): 187–223.
- Chen, L. 1982. Topological structure in visual perception. *Science*, 218(4573): 699–700.
- Cheng, D.; Li, Y.; Xia, S.; Wang, G.; Huang, J.; and Zhang, S. 2023. A Fast Granular-Ball-Based Density Peaks Clustering Algorithm for Large-Scale Data. *IEEE Transactions on Neural Networks and Learning Systems*, 1–14.
- Dai, D.; Zhu, H.; Xia, S.; and Wang, G. 2024. Granular-ball Representation Learning for Deep CNN on Learning with Label Noise. *arXiv preprint arXiv:2409.03254*.
- Dhillon, I. S.; Guan, Y.; and Kulis, B. 2007. Weighted graph cuts without eigenvectors a multilevel approach. *IEEE transactions on pattern analysis and machine intelligence*, 29(11): 1944–1957.
- Dickens, C.; Huang, E.; Reganti, A.; Zhu, J.; Subbian, K.; and Koutra, D. 2024. Graph coarsening via convolution matching for scalable graph neural network training. In *Companion Proceedings of the ACM on Web Conference 2024*, 1502–1510.
- Fahrbach, M.; Goranci, G.; Peng, R.; Sachdeva, S.; and Wang, C. 2020. Faster graph embeddings via coarsening. In *international conference on machine learning*, 2953–2963. PMLR.
- Feng, S.; Tan, Z.; Wan, H.; Wang, N.; Chen, Z.; Zhang, B.; Zheng, Q.; Zhang, W.; Lei, Z.; Yang, S.; et al. 2022. Twibot-22: Towards graph-based twitter bot detection. *Advances in Neural Information Processing Systems*, 35: 35254–35269.
- Fout, A.; Byrd, J.; Shariat, B.; and Ben-Hur, A. 2017. Protein interface prediction using graph convolutional networks. *Advances in neural information processing systems*, 30.
- Huang, Z.; Zhang, S.; Xi, C.; Liu, T.; and Zhou, M. 2021. Scaling up graph neural networks via graph coarsening. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, 675–684.
- Imre, M.; Tao, J.; Wang, Y.; Zhao, Z.; Feng, Z.; and Wang, C. 2020. Spectrum-preserving sparsification for visualization of big graphs. *Computers & Graphics*, 87: 89–102.
- Jin, W.; Tang, X.; Jiang, H.; Li, Z.; Zhang, D.; Tang, J.; and Yin, B. 2022. Condensing graphs via one-step gradient matching. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 720–730.
- Jin, W.; Zhao, L.; Zhang, S.; Liu, Y.; Tang, J.; and Shah, N. 2021. Graph condensation for graph neural networks. *arXiv preprint arXiv:2110.07580*.
- Jin, Y.; Loukas, A.; and JaJa, J. 2020. Graph coarsening with preserved spectral properties. In *International Conference on Artificial Intelligence and Statistics*, 4452–4462. PMLR.
- Kipf, T. N.; and Welling, M. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Kumar, M.; Sharma, A.; Saxena, S.; and Kumar, S. 2023. Featured graph coarsening with similarity guarantees. In *International Conference on Machine Learning*, 17953–17975. PMLR.
- Li, C.; and Goldwasser, D. 2019. Encoding social information with graph convolutional networks for political perspective detection in news media. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2594–2604.
- Liu, J.; Jianye, H.; Ma, Y.; and Xia, S. 2024. Unlock the Cognitive Generalization of Deep Reinforcement Learning via Granular Ball Representation. In *Forty-first International Conference on Machine Learning*.
- Liu, Y.; Safavi, T.; Dighe, A.; and Koutra, D. 2018. Graph summarization methods and applications: A survey. *ACM computing surveys (CSUR)*, 51(3): 1–34.
- Loukas, A. 2019. Graph reduction with spectral and cut guarantees. *Journal of Machine Learning Research*, 20(116): 1–42.
- Loukas, A.; and Vandergheynst, P. 2018. Spectrally approximating large graphs with smaller graphs. In *International conference on machine learning*, 3237–3246. PMLR.
- Pang, Y.; Zhao, Y.; and Li, D. 2021. Graph pooling via coarsened graph infomax. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2177–2181.
- Quadir, A.; and Tanveer, M. 2024. Granular ball twin support vector machine with pinball loss function. *IEEE Transactions on Computational Social Systems*.
- Ron, D.; Safro, I.; and Brandt, A. 2011. Relaxation-based coarsening and multiscale graph organization. *Multiscale Modeling & Simulation*, 9(1): 407–423.
- Shchur, O.; Mumme, M.; Bojchevski, A.; and Günnemann, S. 2018. Pitfalls of graph neural network evaluation. *arXiv preprint arXiv:1811.05868*.
- Shuman, D. I.; Faraji, M. J.; and Vandergheynst, P. 2015. A multiscale pyramid transform for graph signals. *IEEE Transactions on Signal Processing*, 64(8): 2119–2134.
- Stüben, K. 2001. A review of algebraic multigrid. *Numerical Analysis: Historical Developments in the 20th Century*, 331–359.
- Wang, G. 2017. DGCC: data-driven granular cognitive computing. *Granular Computing*, 2(4): 343–355.

- Wang, Z.; Li, J.; Xia, S.; Lin, L.; and Wang, G. 2024a. Text Adversarial Defense via Granular-Ball Sample Enhancement. In *Proceedings of the 2024 International Conference on Multimedia Retrieval*, 348–356.
- Wang, Z.; Zhang, T.; Xia, S.; Lin, L.; and Wang, G. 2024b. GBRAIN: Combating Textual Label Noise by Granular-ball based Robust Training. In *Proceedings of the 2024 International Conference on Multimedia Retrieval*, 357–365.
- Wilson, R. C.; and Zhu, P. 2008. A study of graph spectra for comparing graphs and trees. *Pattern Recognition*, 41(9): 2833–2841.
- Xia, D.; Wang, G.; Zhang, Q.; Yang, J.; and Xia, S. 2024a. Three-way approximations fusion with granular-ball computing to guide multi-granularity fuzzy entropy for feature selection. *IEEE Transactions on Fuzzy Systems*.
- Xia, S.; Lian, X.; Wang, G.; Gao, X.; Chen, J.; and Peng, X. 2024b. GBSVM: An Efficient and Robust Support Vector Machine Framework via Granular-Ball Computing. *IEEE Transactions on Neural Networks and Learning Systems*, 1–15.
- Xia, S.; Lian, X.; Wang, G.; Gao, X.; Hu, Q.; and Shao, Y. 2024c. Granular-Ball Fuzzy Set and Its Implement in SVM. *IEEE Transactions on Knowledge and Data Engineering*, 36(11): 6293–6304.
- Xia, S.; Liu, Y.; Ding, X.; Wang, G.; Yu, H.; and Luo, Y. 2019. Granular ball computing classifiers for efficient, scalable and robust learning. *Information Sciences*, 483: 136–152.
- Xia, S.; Wang, C.; Wang, G.; Gao, X.; Ding, W.; Yu, J.; Zhai, Y.; and Chen, Z. 2023. GBRs: A Unified Granular-Ball Learning Model of Pawlak Rough Set and Neighborhood Rough Set. *IEEE Transactions on Neural Networks and Learning Systems*.
- Xia, S.; Wang, G.; and Gao, X. 2023. Granular ball computing: an efficient, robust, and interpretable adaptive multi-granularity representation and computation method. *arXiv preprint arXiv:2304.11171*.
- Xia, S.; Zhang, H.; Li, W.; Wang, G.; Giem, E.; and Chen, Z. 2020. GBNRS: A novel rough set algorithm for fast adaptive attribute reduction in classification. *IEEE Transactions on Knowledge and Data Engineering*, 34(3): 1231–1242.
- Xia, S.; Zheng, S.; Wang, G.; Gao, X.; and Wang, B. 2021. Granular ball sampling for noisy label classification or imbalanced classification. *IEEE Transactions on Neural Networks and Learning Systems*, 34(4): 2144–2155.
- Xie, J.; Dai, M.; Xia, S.; Zhang, J.; Wang, G.; and Gao, X. 2024a. An Efficient Fuzzy Stream Clustering Method Based on Granular-Ball Structure. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*, 901–913. IEEE.
- Xie, J.; Hua, C.; Xia, S.; Cheng, Y.; Wang, G.; and Gao, X. 2024b. W-GBC: An Adaptive Weighted Clustering Method Based on Granular-Ball Structure. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*, 914–925. IEEE.
- Xie, J.; Xiang, X.; Xia, S.; Jiang, L.; Wang, G.; and Gao, X. 2024c. MGNR: A Multi-Granularity Neighbor Relationship and Its Application in KNN Classification and Clustering Methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–17.
- Xie, J.; Xiong, Z.-Y.; Dai, Q.-Z.; Wang, X.-X.; and Zhang, Y.-F. 2020a. A new internal index based on density core for clustering validation. *Information Sciences*, 506: 346–365.
- Xie, Y.; Yao, C.; Gong, M.; Chen, C.; and Qin, A. K. 2020b. Graph convolutional networks with multi-level coarsening for graph classification. *Knowledge-Based Systems*, 194: 105578.
- Yang, J.; Liu, Z.; Xia, S.; Wang, G.; Zhang, Q.; Li, S.; and Xu, T. 2024. 3WC-GBNRS++: A novel three-way classifier with granular-ball neighborhood rough sets based on uncertainty. *IEEE Transactions on Fuzzy Systems*.
- Ying, R.; He, R.; Chen, K.; Eksombatchai, P.; Hamilton, W. L.; and Leskovec, J. 2018. Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, 974–983.
- Zhang, Q.; Wu, C.; Xia, S.; Zhao, F.; Gao, M.; Cheng, Y.; and Wang, G. 2023. Incremental learning based on granular ball rough sets for classification in dynamic mixed-type decision system. *IEEE Transactions on Knowledge and Data Engineering*, 35(9): 9319–9332.
- Zhang, Z.; Liu, Q.; Hu, Q.; and Lee, C.-K. 2022. Hierarchical graph transformer with adaptive node sampling. *Advances in Neural Information Processing Systems*, 35: 21171–21183.
- Zhou, J.; Cui, G.; Hu, S.; Zhang, Z.; Yang, C.; Liu, Z.; Wang, L.; Li, C.; and Sun, M. 2018. Graph Neural Networks: A Review of Methods and Applications. *arXiv e-prints*, arXiv-1812.