



PDF Download  
3773966.3777968.pdf  
06 April 2026  
Total Citations: 0  
Total Downloads: 251

Published: 21 February 2026

Citation in BibTeX format

WSDM '26: The Nineteenth ACM  
International Conference on Web Search  
and Data Mining  
February 22 - 26, 2026  
ID, Boise, USA

Conference Sponsors:

SIGKDD  
SIGWEB  
SIGIR  
SIGMOD

DL Latest updates: <https://dl.acm.org/doi/10.1145/3773966.3777968>

RESEARCH-ARTICLE

## Multi-view Graph Condensation via Tensor Decomposition

**NICOLAS ROQUE DOS SANTOS**, University of California, Riverside, Riverside, CA, United States

**DAWON AHN**, University of California, Riverside, Riverside, CA, United States

**DIEGO MINATEL**, University of São Paulo, Sao Paulo, SP, Brazil

**ALNEU DE ANDRADE LOPES**, University of São Paulo, Sao Paulo, SP, Brazil

**EVANGELOS PAPALEXAKIS**, University of California, Riverside, Riverside, CA, United States

Open Access Support provided by:

University of California, Riverside

University of São Paulo

# Multi-view Graph Condensation via Tensor Decomposition

Nicolas Roque dos Santos  
University of California, Riverside  
Riverside, California, USA  
nicolasr@ucr.edu

Dawon Ahn  
University of California, Riverside  
Riverside, California, USA  
dahn017@ucr.edu

Diego Minatel  
University of São Paulo  
São Carlos, São Paulo, Brazil  
dminatel@usp.br

Alneu de Andrade Lopes  
University of São Paulo  
São Carlos, São Paulo, Brazil  
alneu@icmc.usp.br

Evangelos Papalexakis  
University of California, Riverside  
Riverside, California, USA  
epapalex@cs.ucr.edu

## Abstract

Graph Neural Networks (GNNs) have demonstrated remarkable results in various real-world applications, including drug discovery, object detection, social media analysis, recommender systems, and text classification. In contrast to their vast potential, training them on large-scale graphs presents significant computational challenges due to the resources required for their storage and processing. Graph Condensation has emerged as a promising solution to reduce these demands by learning a synthetic compact graph that preserves the essential information of the original one while maintaining the GNN's predictive performance. Despite their efficacy, current graph condensation approaches frequently rely on a computationally intensive bi-level optimization. Moreover, they fail to maintain a mapping between synthetic and original nodes, limiting the interpretability of the model's decisions. In this sense, a wide range of decomposition techniques have been applied to learn linear or multi-linear functions from graph data, offering a more transparent and less resource-intensive alternative. However, their applicability to graph condensation remains unexplored. This paper addresses this gap and proposes a novel method called Multi-view Graph Condensation via Tensor Decomposition (GCTD) to investigate to what extent such techniques can synthesize an informative smaller graph and achieve comparable downstream task performance. Extensive experiments on six real-world datasets demonstrate that GCTD effectively reduces graph size while preserving GNN performance, achieving up to a 4.0% improvement in accuracy on three out of six datasets and competitive performance on large graphs compared to existing approaches. Our code is available at <https://github.com/nicolarsantos/gctd>.

## CCS Concepts

• **Computing methodologies** → **Neural networks; Factorization methods**; • **Information systems** → **Data mining**.

## Keywords

Graph Condensation; Tensor Decomposition; Graph Neural Networks; Matrix Tri-factorization



This work is licensed under a Creative Commons Attribution 4.0 International License. *WSDM '26, Boise, ID, USA*

© 2026 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-2292-9/2026/02  
<https://doi.org/10.1145/3773966.3777968>

## ACM Reference Format:

Nicolas Roque dos Santos, Dawon Ahn, Diego Minatel, Alneu de Andrade Lopes, and Evangelos Papalexakis. 2026. Multi-view Graph Condensation via Tensor Decomposition. In *Proceedings of the Nineteenth ACM International Conference on Web Search and Data Mining (WSDM '26)*, February 22–26, 2026, Boise, ID, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3773966.3777968>

## 1 Introduction

Graph is a ubiquitous data structure used to model a broad spectrum of real-world systems, encompassing transportation networks, protein-protein interactions, social media, and epidemic spreading [6]. Over the past few years, Graph Neural Networks (GNNs) have become a pivotal tool for extracting representations from such data, supporting various downstream tasks. For example, they have been applied in the discovery of antibiotic compounds [46], improving estimated time of arrival (ETA) [9], and detecting the spread of fake news on social networks [45]. Despite their success, GNNs face significant scalability challenges when deployed on large-scale graphs [18, 28]. These issues become even more pronounced in scenarios like Neural Architecture Search or Continual Learning, where models require frequent retraining or incremental updates.

Numerous graph size reduction techniques have been proposed to mitigate these scaling issues [20]. Among these, condensation stands out as a promising one. Specifically, it aims to learn a smaller, representative graph that enables a GNN to achieve performance comparable to training on the original one. By reducing the graph, condensation significantly alleviates computational costs, making GNNs more feasible for large datasets. Existing methods typically generate condensed versions of the original graph using techniques such as Kernel Ridge Regression [49, 56] or matching gradients [28], training trajectories [64], eigenbasis [35], and distribution [34]. However, while such methods can compress graphs to an extreme degree without sacrificing much accuracy, they often rely on a complex bi-level optimization. Moreover, the need for multiple parameter initializations results in a costly triple-loop procedure. Additionally, these methods lack interpretability, losing the notion of how synthetic nodes relate to the original ones.

At the same time, there exists an extensive line of work in the literature that leverages matrix or tensor decomposition for graph-based tasks that learn simpler functions (*i.e.*, linear or multi-linear) from data in contrast to the nonlinear ones learned by the GNNs [22, 24, 32, 39]. Furthermore, they offer a more transparent view of their decisions, providing a more interpretable approach. Importantly,

the objectives of tensor decomposition and graph condensation are aligned: both aim to reduce the size and complexity of the original data while preserving its essential information, which in turn lowers the computational cost of downstream tasks. Nevertheless, to the best of our knowledge, no prior study has explored the application of tensor decomposition for synthesizing smaller graphs.

In this work, we address this gap by reframing the condensation task as a decomposition problem, investigating whether it can capture and transfer key information from a larger graph to a smaller one, thereby enabling a GNN to achieve performance comparable to that obtained when trained on the original graph. To achieve this, we propose a novel method called **Multi-view Graph Condensation via Tensor Decomposition (GCTD)**. Specifically, we construct a multi-view graph by augmenting the original adjacency matrix into a third-order tensor through random edge perturbation. We then decompose this tensor and exploit the optimized latent factors to obtain a synthetic graph. The core idea is that the model discovers latent co-clusters within the data, enabling it to group nodes into synthetic ones based on shared patterns [17].

Through extensive experiments on six real-world datasets, we demonstrate that GCTD effectively synthesizes smaller graphs, outperforming existing baselines on three out of six datasets with up to 4.0% improvement in accuracy. Additionally, we show that using multi-view decomposition leads to better performance than its single counterpart. Our contributions are summarized as follows:

- We propose a novel method for graph condensation that leverages a tensor decomposition technique to reduce the original graph;
- We employ multi-view augmented graphs and show through a comprehensive analysis that it improves upon single-view decomposition;
- We conduct an extensive analysis to showcase to what extent decomposition methods can condense graphs;

## 2 Related Work

**Graph Reduction.** Graph reduction encompasses many techniques developed to alleviate the computational cost of training graph-based algorithms by reducing the original graph into a smaller but informative one. Some examples of such methods include coreset [42, 50], sparsification [2, 29], and coarsening [13, 26, 37]. Specifically, coreset methods extract representative samples that reflect the dataset. Furthermore, while sparsification reduces the original graph by removing nodes and edges based on properties such as eigenvalues, pairwise distances, or cuts, coarsening techniques merge nodes into supernodes based on their similarity [20, 26, 37]. Although these techniques can preserve information from the graph, they might not be sufficient for the downstream task. Additionally, they face a significant performance loss when the reduction rate is large (e.g., 99.9%). Therefore, a new set of approaches known as condensation has emerged to tackle this problem.

**Graph Condensation.** Condensation is a technique that learns a small synthetic dataset that retains the essential patterns present in a larger one. Thus, when a model is trained on the reduced dataset, its performance is similar to when trained on the original one. These methods were originally developed for images [61, 62] and later

extended for graphs by Jin et al. [28], which proposed GCond, a method that synthesizes a graph by matching the gradients of a GNN trained on both the original and simplified graphs.

Since GCond, several strategies along the same line of research have emerged. For example, SGDD [57] and GCSR [36] incorporate the original structure information to generate the smaller graph. Conversely, SFGC [64] and GEOM [60] condense a graph by matching the training trajectories of the learned and original graphs. Rather than mimicking the training gradients or trajectories, GCDM [34] and GDEM [35] synthesize a smaller graph through distribution and eigenbasis matching, respectively.

One of the major issues most graph condensation techniques face is the presence of a costly bi-level optimization. Therefore, DosCond [27] tackles this problem by performing a single gradient matching step to synthesize a reduced graph. Moreover, KiDD [56] and SNTK [49] avoid the bi-level matching by reframing condensation within the perspective of the Kernel Ridge Regression method.

Other approaches include CGC [15], which proposes a fast, training-free condensation method; DisCo [53], which introduces a disentangled, GNN-free framework for graph condensation; and SimGC [54], which employs an MLP alongside a heuristic to preserve the distribution of node features.

**Brief comparison.** In contrast to existing graph condensation methods, our approach leverages tensor decomposition rather than GNN-based learning. Moreover, while prior methods typically depend on a computationally expensive triple-loop optimization procedure (discussed in the next section), our method removes this requirement entirely. Finally, unlike single-view approaches, our model exploits a multi-view graph representation, highlighting the benefits of incorporating augmentations into condensation.

## 3 Preliminaries

In this section, we present the formal definition of graph condensation and then discuss a limitation with one of the main strategies adopted by prior work. Then, we describe the preliminaries on tensors and their decompositions.

### 3.1 Graph Condensation

A graph is denoted as  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , where  $\mathcal{V} = \{v_1, \dots, v_n\}$  comprises a set of  $N$  nodes and  $\mathcal{E}$  is the edge set. The graph's structure can be represented as an adjacency matrix  $\mathbf{A} \in \mathbb{R}^{N \times N}$ , where each entry  $A_{ij} = 1$  if there is an edge between nodes  $v_i$  and  $v_j$ , and 0 otherwise. Moreover,  $\mathbf{X} \in \mathbb{R}^{N \times d}$  represents the  $d$ -dimensional node feature matrix and  $\mathbf{Y} \in \{0, \dots, C-1\}^N$  are the node labels over  $C$  classes.

Given a target graph  $\mathcal{G}^{\mathcal{T}} = \{\mathbf{A}^{\mathcal{T}}, \mathbf{X}^{\mathcal{T}}, \mathbf{Y}^{\mathcal{T}}\}$ , the goal of graph condensation is to reduce  $\mathcal{G}^{\mathcal{T}}$  into a smaller graph  $\mathcal{G}^{\mathcal{S}} = \{\mathbf{A}^{\mathcal{S}}, \mathbf{X}^{\mathcal{S}}, \mathbf{Y}^{\mathcal{S}}\}$ , where  $\mathbf{A}^{\mathcal{S}} \in \mathbb{R}^{N' \times N'}$ ,  $\mathbf{X}^{\mathcal{S}} \in \mathbb{R}^{N' \times d}$ ,  $\mathbf{Y}^{\mathcal{S}} \in \{0, \dots, C-1\}^{N'}$ , and  $N' \ll N$ , such that a GNN trained on  $\mathcal{G}^{\mathcal{S}}$  achieves similar performance to when trained on the significantly larger graph  $\mathcal{G}^{\mathcal{T}}$ . Existing work for graph condensation approaches this task as a bi-level optimization problem [28, 57, 64]. Formally, it is defined as:

$$\begin{aligned} & \min_{\mathcal{G}^{\mathcal{S}}} \mathcal{L}(\text{GNN}_{\theta_{\mathcal{G}^{\mathcal{S}}}}(\mathbf{A}^{\mathcal{T}}, \mathbf{X}^{\mathcal{T}}), \mathbf{Y}^{\mathcal{T}}) \\ & \text{s.t. } \theta_{\mathcal{G}^{\mathcal{S}}} = \arg \min_{\theta} \mathcal{L}(\text{GNN}_{\theta}(\mathbf{A}^{\mathcal{S}}, \mathbf{X}^{\mathcal{S}}), \mathbf{Y}^{\mathcal{S}}). \end{aligned} \quad (1)$$

Here, the outer loop is responsible for optimizing the synthetic graph based on, for example, a gradient matching loss, while the inner loop trains a GNN parameterized by  $\theta$  on the synthetic dataset. Furthermore, this double loop is computed multiple times with different parameter initializations to ensure the entire process is not overfitting to a specific  $\theta$ . Consequently, it becomes a triple-loop procedure, making graph condensation models that follow this scheme computationally costly in terms of time, storage, and processing power. In addition to the aforementioned expensive step, most methods leverage a GNN to learn a condensed graph. This results in a black-box model, making it difficult to establish a clear connection between the original and synthetic nodes.

### 3.2 Tensors

**Notations.** Tensors are multi-dimensional arrays that generalize one-dimensional arrays (or vectors) and two-dimensional arrays (or matrices) to higher dimensions. Throughout this paper, we use boldface Euler script letters (e.g.,  $\mathcal{X}$ ) to denote tensors, boldface capitals (e.g.,  $\mathbf{A}$ ) to denote matrices, boldface lowercases (e.g.,  $\mathbf{a}$ ) to denote vectors, and unbolded letters (e.g.,  $A$ ,  $a$ ) to denote scalars and coefficients. We refer to a tensor’s dimension as its *mode* or *order*. The *slices* are two-dimensional sections of a tensor, defined by fixing all but two indices. For example, the  $k$ -th frontal slice of a third-order tensor  $\mathcal{X} \in \mathbb{R}^{I \times J \times K}$  is a matrix denoted as  $\mathbf{X}_k \in \mathbb{R}^{I \times J}$ . Moreover, the  $n$ -mode product defines the multiplication of a tensor with a matrix in mode  $n$ . For example, the  $n$ -mode product of a tensor  $\mathcal{X} \in \mathbb{R}^{I \times J \times K}$  with a matrix  $\mathbf{U} \in \mathbb{R}^{R \times I}$  along the first mode is denoted by  $\mathcal{X} \times_1 \mathbf{U} (\in \mathbb{R}^{R \times J \times K})$ .

**Tensor decomposition.** Tensor decomposition techniques aim to decompose tensors into low-rank latent components, facilitating data mining and analysis. Among the most widely used models are CANDECOMP/PARAFAC (CP) and Tucker decompositions, both of which have seen extensive development and application across a diverse range of fields [40, 44]. In this work, we leverage a variant of Tucker called RESCAL [38] and, thus, we describe them next.

Tucker decomposition approximates a given third mode tensor  $\mathcal{X} \in \mathbb{R}^{I \times J \times K}$  as follows:

$$\mathcal{X} \approx \mathcal{R} \times_1 \mathbf{U}^{(1)} \times_2 \mathbf{U}^{(2)} \times_3 \mathbf{U}^{(3)}, \quad (2)$$

where  $\times_1$ ,  $\times_2$ , and  $\times_3$  denote the  $n$ -mode product along the first, second, and third modes, respectively. Factor matrices  $\mathbf{U}^{(1)} \in \mathbb{R}^{I \times R_1}$ ,  $\mathbf{U}^{(2)} \in \mathbb{R}^{J \times R_2}$  and  $\mathbf{U}^{(3)} \in \mathbb{R}^{K \times R_3}$  are considered as the principal components in each mode. The core tensor  $\mathcal{R} \in \mathbb{R}^{R_1 \times R_2 \times R_3}$  captures the interaction between each component. Here,  $R_1$ ,  $R_2$ , and  $R_3$  denote the number of components in each mode and are smaller than  $I$ ,  $J$ ,  $K$ , respectively. Furthermore, it has been demonstrated that  $\mathcal{R}$  can be viewed as a compressed version of  $\mathcal{X}$  [31].

RESCAL was initially proposed for relational learning and is particularly useful when the frontal slices of the given tensors exhibit symmetry. It is a special case of the Tucker decomposition, as described in Equation 2, with a few key differences. Specifically, the core tensor  $\mathcal{R} \in \mathbb{R}^{R_1 \times R_2 \times K}$  is employed, and the factor matrix  $\mathbf{U}^{(3)}$  is set as an identity matrix. Additionally, the factor matrices  $\mathbf{U}^{(1)}$  and  $\mathbf{U}^{(2)}$  are identical and denoted by  $\mathbf{U}$ , as follows:

$$\mathcal{X} \approx \mathcal{R} \times_1 \mathbf{U} \times_2 \mathbf{U} \Leftrightarrow \mathbf{X}_k \approx \mathbf{U} \mathbf{R}_k \mathbf{U}^\top, \quad (3)$$

where  $\mathbf{R}_k \in \mathbb{R}^{R_1 \times R_2}$  indicates the relations between latent components. It is important to highlight that RESCAL does not compress the third mode of the tensor.

### 4 First attempt: Single-view Graph Modeling

Motivated by prior work employing Matrix Factorization (MF) for tasks such as clustering [11, 55], dimensionality reduction [1, 33], and compression [5, 47], all of which aim to simplify data complexity, our initial approach explored single-view graph condensation by factorizing the adjacency matrix  $\mathbf{A}^\mathcal{T}$  of the graph  $\mathcal{G}^\mathcal{T}$ .

In this work, we use a variant of MF called Matrix Tri-factorization (MTF), which decomposes a matrix into the product of three lower-dimensional factors [11]. Specifically, we approximate  $\mathbf{A}^\mathcal{T} \approx \mathbf{U} \mathbf{R} \mathbf{V}^\top$ , where  $\mathbf{U}$  and  $\mathbf{V}$  capture the row and column spaces of  $\mathbf{A}^\mathcal{T}$ , while  $\mathbf{R}$  encodes the interactions between them. However, given the nature of the undirected graphs, we adopt a symmetric formulation of MTF where  $\mathbf{V} = \mathbf{U}$ , yielding the following reconstruction:

$$\mathbf{A}^\mathcal{T} \approx \mathbf{U} \mathbf{R} \mathbf{U}^\top, \quad (4)$$

where  $\mathbf{U} \in \mathbb{R}^{N \times N'}$  and  $\mathbf{R} \in \mathbb{R}^{N' \times N'}$ . Here,  $N$  is the number of nodes in the original graph,  $N' = rN$  is the number of nodes in the condensed graph,  $r$  is the condensation ratio, and  $N' \ll N$ . Note that if the core tensor in Equation 3 is a matrix, Equation 4 is equivalent to the Equation 3.

**Reconstruction objective.** After computing  $\mathbf{A}^\mathcal{T}$ , we optimize both  $\mathbf{U}$  and  $\mathbf{R}$  using gradient descent, employing mean squared error (MSE) as the reconstruction loss function. Additionally, we identified empirically that normalizing the MSE by the sum of the squared elements of the original matrix  $\mathbf{A}$  enhances convergence speed and stability. Therefore, the final loss function is defined as:

$$\mathcal{L}_{rec} = \frac{\frac{1}{n} \sum_{i,j} (a_{ij} - \hat{a}_{ij})^2}{\sum_{i,j} a_{ij}^2}, \quad (5)$$

where  $a_{ij}$  and  $\hat{a}_{ij}$  represent elements from the original and reconstructed matrices, respectively.

**Uncovering the condensed graph.** The reconstruction from Equation 4 provides a low-rank approximation of the original adjacency matrix, grouping together nodes with similar patterns and capturing their relationships through  $\mathbf{U}$  and  $\mathbf{R}$ , respectively. The interaction between the rows of  $\mathbf{U}$ , encoded by  $\mathbf{R}$ , are then used as the adjacency matrix  $\mathbf{A}^\mathcal{S}$  of the condensed graph.

In order to extract the synthetic nodes from this decomposition, we apply the K-Means<sup>1</sup> algorithm to the rows of the factor matrix  $\mathbf{U}$ . Formally, we minimize the following error function:

$$\begin{aligned} \mathcal{L}_{\text{K-Means}} &= \sum_{i=1}^I \sum_{j=1}^J \delta_{ij} \|\mathbf{u}_i - \mu_j\|^2 \\ \text{s.t. } \delta_{ij} &= \begin{cases} 1, & \text{if } j = \arg \min_k \|\mathbf{u}_i - \mu_k\| \\ 0, & \text{otherwise} \end{cases} \end{aligned} \quad (6)$$

where  $\mathbf{u}_i$  is a row from the factor matrix,  $\mu_j$  corresponds to a cluster centroid, and  $\delta_{ij} \in \{0, 1\}$  indicates to which of the  $J$  clusters  $\mathbf{u}_i$  is assigned to. We initialize the centroids randomly and the number

<sup>1</sup>We use the implementation provided by the Faiss library, available at faiss.ai.

of clusters is set to match the number of columns of  $\mathbf{U}$ , denoted  $N'$ , which corresponds to the size of  $\mathcal{G}^S$ . This process determines which nodes from the original graph  $\mathcal{G}^T$  should be grouped together in the reduced graph, providing the necessary information to assign class labels, splits, and features to the synthetic nodes.

Given the assignments for cluster  $j$ , the split of the corresponding synthetic node is chosen as the most frequent split among the original nodes  $i$  for which  $\delta_{i,j} = 1$ . The class label is then determined by considering only the nodes from that split. For example, if five nodes are assigned to cluster  $j$  with splits train, val, train, test, train and classes 0, 2, 1, 0, 0, the synthetic node will be assigned to the training split and to class 0, which is the most frequent class among the training nodes. Finally, the features of the synthetic node are computed by averaging the embeddings of the nodes assigned to it that share both its split and class. Throughout this process, we prioritize underrepresented splits and classes to ensure the condensed graph preserves the original distribution.

**Remarks on single-view decomposition.** While single-view decomposition enables the synthesis of a condensed graph, we extend this approach to a multi-view setting, as recent studies show that multi-view augmentation improves model performance, generalization, and robustness [12, 52, 63]. Moreover, as demonstrated in Section 6.2, this strategy yields notable gains in GNN performance on condensed graphs compared to the single-view case. Thus, in the following section, we introduce our method, which leverages tensor decomposition to condense multi-view augmented graphs.

## 5 Proposed Method

We introduce our proposed method, GCTD. We begin by outlining the process for obtaining multi-view graphs, followed by a detailed explanation of the decomposition procedure. Afterward, we describe the steps involved in deriving the condensed graph. Figure 1 illustrates the complete pipeline of our method.

**Multi-view augmentation.** Data augmentation (DA) methods have been widely explored across various domains, including Graph Machine Learning, Computer Vision, and Natural Language Processing [14, 43, 63]. They aim to generate variations of existing data, thereby alleviating the burden of collecting and labeling additional samples. Existing work has demonstrated that synthetic samples produced through DA enhance generalization, improving downstream task performance [7, 41, 48]. Furthermore, the variations created by DA enable models to effectively learn to distinguish signal from noise [63].

The utilization of DA to generate multiple views of a dataset and jointly extract information from them has been investigated in several studies [4, 58]. Additionally, it has also been shown that employing a multi-view graph not only improves robustness against adversarial attacks but also reduces the variance of the GNNs' predictions on augmented graphs [52].

Motivated by the aforementioned work on augmentation and multi-view graphs, we generate multiple views of the graph by augmenting its adjacency matrix  $\mathbf{A}^T$  to obtain perturbed versions of its original structure. Specifically, we construct a set of  $K$  matrices,  $S = \{s_1, \dots, s_K\}$ , where each matrix represents an augmented

version of  $\mathbf{A}^T$  with edges randomly removed and added based on prespecified probabilities.

For each perturbed matrix  $s_k$ , edges are first dropped randomly with probability  $p_r$ . This results in a reduced edge count of  $M' = M(1 - p_r)$ , where  $M$  is the original number of edges. Subsequently, new edges are added randomly with probability  $p_a$ , drawn from the set of potential edges not already present in the graph. The total number of edges after this step is updated to  $M' = M' + p_a \cdot \left(\frac{N(N-1)}{2} - M'\right)$ , where  $\frac{N(N-1)}{2}$  represents the total number of possible edges in an undirected graph with  $N$  nodes. In practice, we set  $p_r$  and  $p_a$  to small values (i.e.,  $0.05 \leq p_r, p_a \leq 0.2$ ) to ensure that the resulting perturbed graphs remain with a similar size to the original one, with  $M' \approx M$  on average.

The perturbed matrices are then combined to form a tensor  $\mathcal{X}$ , where the first slice corresponds to  $\mathbf{A}^T$ , and the remaining ones consist of the  $K$  matrices in  $S$ . The final structure of  $\mathcal{X}$  is illustrated in Step 1 of Figure 1. It is important to note that, to mitigate computational overhead during the augmentation process, all matrices in  $S$  are precomputed.

**Decomposing a multi-view graph.** We reconstruct the multi-view graph  $\mathcal{X}$  using the tensor decomposition RESCAL [38]. Formally, we compute the following operation for each slice  $k$  of  $\mathcal{X}$ :

$$\tilde{\mathbf{X}}_k = \mathbf{U}\mathbf{R}_k\mathbf{U}^\top, \quad (7)$$

where  $\mathbf{U} \in \mathbb{R}^{N \times N'}$  is the factor matrix capturing latent components in the given graph and  $\mathbf{R}_k$  is the  $k$ -th frontal slice of the core tensor  $\mathcal{R} \in \mathbb{R}^{N' \times N' \times K}$ , indicating relations between latent components existing in the  $k$ -th view of the graph. Note that  $N'$  is the size of the condensed graph, and  $K$  is the total number of augmented graphs.


The formulation presented in Equation 7 consists of an operation that produces  $k$  dense matrices, which is computationally expensive for large graphs. To make things worse, keeping them in memory during runtime is resource-intensive since we have to store multiple  $N \times N$  matrices. To address this issue, we instead adopt the sparse version of the decomposition where we reconstruct only the observed entries of the original tensor as follows:

$$\tilde{x}_{ijk} = \mathbf{u}_i^\top \mathbf{R}_k \mathbf{u}_j. \quad (8)$$

Since this operation involves only observed entries (i.e., nonzero values), providing negative samples to the model is essential to prevent overfitting to the training data. Therefore, we randomly generate negative examples in 1:1 proportion to the available nonzero values, ensuring a balanced representation of both positive and negative samples. It is noteworthy that this process is precomputed to minimize the computational overhead during the decomposition. In addition to performing the sparse operation on the nonzero values, we adopt mini-batching in our model's reconstruction step to further reduce the memory footprints.

After reconstruction, we optimize  $\mathcal{R}$  and  $\mathbf{U}$  using the same objective as in the single-view decomposition, with a slight modification to Equation 5 to account for tensor entries:

$$\mathcal{L}_{rec} = \frac{\frac{1}{n} \sum_{i,j,k} (x_{ijk} - \tilde{x}_{ijk})^2}{\sum_{i,j,k} x_{ijk}^2}, \quad (9)$$



**Figure 1: Pipeline of GCTD.** We construct a tensor by augmenting the graph’s adjacency matrix  $A^T$  and stacking them together in the third dimension with  $A^T$ . Then, we apply non-negative RESCAL to the given tensor to extract low-rank structures  $U$  and a multi-view condensed graph  $\mathcal{R}$ . Lastly, we obtain a condensed graph by aggregating  $\mathcal{R}$  along the third mode, and we compute the feature and label for each synthetic node by applying K-Means to  $U$ .

where  $x_{ijk}$  and  $\tilde{x}_{ijk}$  are entries from the original and reconstructed tensor, respectively.

**Nonnegativity and sparsity constraint.** Given the nonnegative nature of most real-world graphs, we impose a constraint in our method to maintain the nonnegative characteristics of the data. It is important to note that we introduce hard constraints instead of soft regularization-based ones that prior work did [66]. Specifically, we first ensure that the random initialization of the factor matrix and core tensor comprises only nonnegative values by applying the absolute function to them (*i.e.*,  $U \leftarrow |U|$  and  $\mathcal{R} \leftarrow |\mathcal{R}|$ ). Additionally, since every optimization step of GCTD can shift some values of  $U$  and  $\mathcal{R}$  to negative, we apply the nonlinear activation function ReLU to them in every decomposition epoch to ensure they remain nonnegative, as shown in Algorithm 1. Besides guaranteeing the constraint mentioned above, applying ReLU inherently induces sparsity in the learned graph since it sets a portion of the values to zero, reducing the storage necessary for the condensed graph.

**Computing the structure and features of a condensed graph.** While it is not possible to make uniqueness arguments for GCTD similar to those for CP [44], by imposing nonnegativity and sparsity to  $\mathcal{R}$  and  $U$ , we are constraining the otherwise non-unique Tucker-like model to a smaller space of solutions [25]. These imposed constraints guide GCTD toward behaving close to a constrained CP model, enabling it to uncover latent co-clusters from the data effectively [17]. Thus, we can interpret that the computed factor  $U$  co-clusters the original nodes based on similar patterns they share, and the core tensor  $\mathcal{R}$  captures the relations between them.

We specifically uncover the condensed graph from  $\mathcal{R}$  and  $U$  similarly to the process we described in the single-view decomposition approach. The only existing difference is that a tensor is the latent component that captures the interactions between the entities in the factors, instead of a matrix. Therefore, after deriving the information for each synthetic node, we compute the structure of the condensed graph  $\mathcal{G}^S$ . This is done by averaging each  $\mathcal{R}_{i,j,:}$  from the core tensor to form the adjacency matrix  $A^S$ . Since the aggregation process of the core tensor may result in a non-symmetric matrix,

we add the corresponding off-diagonal values to symmetrize  $A^S$ .

**Algorithm.** We summarize the step-by-step of GCTD in Algorithm 1. Our method takes as input the observed entries of the multi-view graph along with the negative samples. The factor matrix and core tensor are randomly initialized, and the absolute function is applied to ensure their nonnegativity. From lines 6 to 9 of the algorithm, the adjacency matrix is decomposed and the latent components are optimized while maintaining their nonnegativity throughout the process. Convergence is determined either after 200 epochs or when the variation in reconstruction loss between consecutive steps falls below  $10^{-7}$ , as previously described. Finally, the condensed graph is uncovered by averaging the core tensor and applying KMeans clustering to the factor matrix, extracting the necessary information to generate the features, classes, and splits.

---

**Algorithm 1:** Graph Condensation via Tensor Decomposition (GCTD)

---

- 1 **Input:** Pre-computed multi-view graph  $\mathcal{X}$  with negative samples.
  - 2 **Output:** Condensed graph  $\mathcal{G}^S = (A^S, X^S, Y^S)$
  - 3 Initialize matrix  $U$  and core tensor  $\mathcal{R}$  randomly.
  - 4  $U \leftarrow |U|$ ;  $\mathcal{R} \leftarrow |\mathcal{R}|$
  - 5 **while** *convergence is not achieved* **do**
  - 6     Reconstruct  $\mathcal{X}$  according to Equation 8.
  - 7     Compute the reconstruction error  $\mathcal{L}_{rec}$  according to Equation 9.
  - 8     Update  $U \leftarrow U - \eta \nabla_U \mathcal{L}_{rec}$  and  $\mathcal{R} \leftarrow \mathcal{R} - \eta \nabla_{\mathcal{R}} \mathcal{L}_{rec}$
  - 9      $U \leftarrow \text{ReLU}(U)$ ;  $\mathcal{R} \leftarrow \text{ReLU}(\mathcal{R})$
  - 10 Compute  $A_{ij}^S \leftarrow \text{Average}(\mathcal{R}_{i,j,:})$ ,  $\forall i, j$
  - 11 Cluster  $U$  rows according to Equation 6.
  - 12 Average embeddings of the nodes assigned by K-Means to get  $X^S$ .
  - 13 Assign to each synthetic node the dominant split and class among its assigned nodes, focusing on underrepresented classes first.
- 

The split of a synthetic node is chosen based on the most frequent split among its assigned nodes, with preference given to underrepresented splits according to the target distribution. To assign a class label, we first filter the assigned nodes to those within the selected split and examine their class labels. Instead of simple

majority voting, we compute the class frequencies and prioritize underrepresented classes. The first class whose current count is below its target proportion is selected; if all are satisfied, a class is chosen randomly among the candidates. This strategy helps maintain class balance in the condensed graph. Finally, the feature representation of the synthetic node is computed by averaging the embeddings of its assigned nodes.

**Complexity Analysis.** GCTD’s time complexity consists of three components: tensor decomposition, clustering, and adjacency matrix generation. For tensor decomposition, operations are performed only on observed entries (Equation 8), with each entry requiring  $O(N'^2)$  time. With approximately  $M$  observed entries per view and  $K$  views ( $2 \leq K \leq 5$ ), the decomposition cost is  $O(N'^2MK)$ . For the sparse graphs typically used in graph condensation,  $M \ll N^2$  and  $N' \ll N$ . KMeans clustering incurs  $O(TN'^2N)$  time, where  $T$  denotes the number of iterations ( $T = 20$ ), and constructing the condensed adjacency matrix by averaging along the third mode adds  $O(N'^2)$ . Thus, the overall complexity is  $O(N'^2MK) + O(TN'^2N) + O(N'^2)$ , which remains practical since graph condensation explicitly targets small  $N'$ .

The space complexity is dominated by the factor matrix  $\mathbf{U}$  and the core tensor  $\mathcal{R}$ , requiring  $O(N \times N')$  and  $O(N'^2 \times K)$  memory, respectively. Since only observed entries are stored, the original graph requires  $O(M)$  space, while the multi-view augmented graphs require  $O(M')$ . In practice,  $M' \approx M$ , as the augmentation probabilities  $p_r$  and  $p_a$  are set between 0.05 and 0.2.

## 6 Experiments

In this section, we present the experimental results of our proposed method for graph condensation, GCTD.

### 6.1 Experimental Settings

**Baselines.** We compared our condensation method against twelve baselines, which include three graph coreset methods (Random, Herding [50], and K-Center [42]), one coarsening method [26], two graph distillation methods: SGDD [57], and GDEM [35], and six graph condensation methods: GCond [28], SFGC [64], SNTK [49], GCSR [36], GCDM [34], and CGC [15].

It is worth noting that although graph coarsening techniques [10, 21, 26, 37] are conceptually related to graph condensation, numerous studies have consistently shown that coarsening tends to underperform condensation methods on downstream tasks at high condensation ratios [15, 28, 35, 57, 60]. Consequently, our evaluation focuses on distillation and condensation baselines, which is consistent with the empirical comparison practices adopted in the majority of recent works in this area.

**Datasets.** Following prior work in graph condensation [28, 57, 64], we evaluate our method on six datasets that range from a few thousand to hundreds of thousands of nodes: four transductive graphs (Cora, Citeseer, Pubmed, and Ogbn-arxiv) and two inductive graphs (Flickr and Reddit). Cora, Citeseer, and Pubmed are taken from the PyTorch Geometric library<sup>2</sup>, Flickr and Reddit from GraphSAINT [59], and Ogbn-arxiv from the Open Graph Benchmark [23]. To

<sup>2</sup><https://pytorch-geometric.readthedocs.io/>

maintain consistency with prior work, all experiments utilize the official public splits for each dataset. Furthermore, we adopt the same dataset-specific condensation ratios employed in previous studies to enable a direct and fair comparison. A summary of dataset characteristics can be found in Table 1.

**Table 1: Statistics of the datasets. The first four datasets are transductive, while the last two are inductive.**

Dataset	Nodes	Edges	Classes	Features	Train/Val/Test
Cora	2,708	5,429	7	1433	140/500/1000
Citeseer	3,327	4,732	6	3,703	120/500/1000
Pubmed	19,717	44,338	3	500	50/500/1000
Ogbn-arxiv	169,343	1,166,243	40	128	90,941/29,799/48,603
Flickr	89,250	899,756	7	500	44,625/22,312/22,313
Reddit	232,965	57,307,946	210	602	153,932/23,699/55,334

**Hyperparameter Settings.** In the condensation step of GCTD, we considered the decomposition converged either after 200 epochs or when the absolute difference in reconstruction error between epochs  $t$  and  $t + 1$  is below  $10^{-7}$ . Following Jin et al. [28], during evaluation, we train a 2-layer GCN with 256 hidden units and no dropout for 600 epochs on the condensed graph, selecting the model with the lowest validation loss for performance evaluation on the original test set. Moreover, we reported the average accuracy and standard deviations of ten runs (condensation and evaluation steps).

To optimize the core tensor  $\mathcal{R}$  and the factor matrix  $\mathbf{U}$ , we used the Adam optimizer [30]. We tuned the learning rates for the reconstruction and evaluation steps within  $\{0.1, 0.01, 0.001, 0.0001\}$ , and the weight decay in a range of  $\{0, 0.01, 0.001, 0.0001\}$ . Additionally, the random edge additions and removals used to generate each augmented view of the original adjacency matrix are tuned across  $\{0.05, 0.1, 0.15, 0.2\}$ . Lastly, the number of views was tuned in a range of  $\{1, 2, 3, 4, 5\}$ , where the value 1 corresponds to a single-view decomposition. All hyperparameter optimization is performed using the Bayesian optimizer provided by Weights and Biases [3].


### 6.2 Experimental Results

**Comparison with baselines.** To begin our analysis, we evaluate the performance of our method against selected baselines across six datasets, using three condensation ratios commonly utilized in the literature. We present the average performance across ten runs and the corresponding standard deviation in Table 2. Notably, our method yields superior performance on the Citeseer, Cora, and Pubmed datasets. For example, it achieves a 4.0%, 3.4%, and 2.7% improvement in accuracy across the three condensation ratios applied to Citeseer. Moreover, GCTD exhibits lossless performance in all settings for Citeseer, Pubmed, and Flickr, as well as in the 1.3% and 2.6% ratios for Cora.

Analyzing the results on Flickr, a challenging dataset in our evaluation, we observe that GCTD performs robustly, securing the second-best results across all condensation ratios and achieving lossless performance. Similarly, GCTD performs competitively on Reddit, consistently ranking among the top-performing methods. For instance, at the 0.05% condensation ratio, it achieves 91.3% accuracy (the second-best result overall) and maintains comparable

**Table 2: Node classification accuracy (%) of the baselines and our proposed method on six datasets with three different condensation ratios. We report the average of ten runs and the corresponding standard deviation. The best and second-best results are in bold and underlined, respectively. OOM stands for out of memory, which in our case was 49GB.**

Datasets	Ratio (%)	Traditional Methods				Condensation and Distillation Methods								Full Dataset	
		Random	Herding	K-Center	Coarse	GCond	SFGC	SGDD	SNTK	GCSR	GCDM	GDEM	CGC		GCTD
Cora	1.3	63.6±3.7	67.0±1.3	64.0±2.3	31.2±0.2	79.8±1.3	77.7±1.8	79.1±1.3	78.4±1.4	79.9±0.7	79.1±0.9	68.0±0.1	<b>82.6±0.3</b>	<u>81.4±1.6</u>	81.4±0.6
	2.6	72.8±1.1	73.4±1.0	73.2±1.2	65.2±0.6	80.1±0.6	79.3±0.8	79.0±1.9	79.7±0.9	80.6±0.8	80.5±0.3	72.8±0.8	<u>81.2±0.6</u>	<b>84.0±0.4</b>	
	5.2	76.8±0.1	76.8±0.1	76.7±0.1	70.6±0.1	79.3±0.3	79.4±0.5	80.2±0.8	80.5±0.6	<u>81.2±0.9</u>	80.2±0.5	77.4±0.6	<b>82.1±0.9</b>	79.9±0.3	
Citeseer	0.9	54.4±4.4	57.1±1.5	52.4±2.8	52.2±0.4	70.5±1.2	66.3±2.4	71.5±0.9	66.1±3.0	70.2±1.1	<u>72.8±0.3</u>	72.3±0.3	72.5±0.5	<b>76.8±0.4</b>	71.7±0.4
	1.8	64.2±1.7	66.7±1.0	64.3±1.0	59.0±0.5	70.6±0.9	69.0±1.1	71.2±0.7	69.2±1.2	71.7±0.9	71.7±0.2	72.6±0.6	<u>73.1±0.2</u>	<b>76.5±2.5</b>	
	3.6	69.1±0.1	69.0±0.1	69.1±0.1	65.3±0.5	69.8±1.4	70.8±0.4	70.9±1.2	71.0±0.6	<u>74.0±0.4</u>	72.5±0.5	72.6±0.5	71.5±0.3	<b>76.7±0.2</b>	
Pubmed	0.08	69.5±0.5	73.0±0.7	69.0±0.6	18.1±0.1	78.3±0.2	76.4±1.2	77.1±0.5	<u>78.9±0.7</u>	77.1±1.1	77.1±0.3	77.7±0.7	77.3±0.1	<b>79.9±0.2</b>	77.1±0.3
	0.15	73.8±0.8	75.4±0.7	73.7±0.8	28.7±4.1	77.1±0.3	77.5±0.4	78.0±0.3	<u>79.3±0.3</u>	77.5±1.9	76.8±0.6	78.4±1.8	76.0±0.5	<b>79.4±2.8</b>	
	0.3	77.9±0.4	77.9±0.4	77.8±0.5	42.8±4.1	78.4±0.3	77.9±0.3	77.5±0.5	<u>79.4±0.3</u>	78.0±0.5	78.1±0.3	78.2±0.8	77.8±0.2	<b>80.0±1.0</b>	
Ogbn-arxiv	0.05	47.1±3.9	52.4±1.8	47.2±3.0	35.4±0.3	59.2±1.1	59.0±1.8	59.6±0.5	58.7±1.7	60.6±1.1	<u>63.8±0.3</u>	63.7±0.8	<b>64.1±0.4</b>	58.2±1.7	71.3±0.1
	0.25	57.3±1.1	58.6±1.2	56.8±0.8	43.5±0.2	63.2±0.3	64.6±0.3	61.7±0.3	64.2±0.5	65.4±0.8	<b>66.7±0.4</b>	63.8±0.6	<u>66.2±0.1</u>	<b>57.9±2.1</b>	
	0.5	60.0±0.9	60.4±0.8	60.3±0.4	50.4±0.1	64.0±0.4	65.2±0.8	58.7±0.6	65.1±0.7	65.9±0.6	<b>67.6±0.3</b>	64.1±0.3	<u>67.0±0.2</u>	57.8±0.2	
Flickr	0.1	41.8±2.0	42.5±1.8	42.0±0.7	41.9±0.2	46.5±0.4	45.5±0.8	46.1±0.3	45.4±0.4	46.6±0.3	44.8±0.5	<b>49.9±0.8</b>	47.1±0.1	<u>48.4±0.8</u>	47.1±0.1
	0.5	44.0±0.4	43.9±0.9	43.2±0.1	44.5±0.1	47.1±0.1	46.0±0.4	45.9±0.4	46.0±0.4	46.6±0.2	46.4±0.2	<b>49.4±1.3</b>	47.1±0.0	<u>48.1±1.4</u>	
	1	44.6±0.2	44.4±0.6	44.1±0.4	44.6±0.1	47.1±0.1	46.1±0.3	46.4±0.2	46.2±0.4	46.8±0.2	46.7±0.2	<b>49.9±0.6</b>	46.8±0.1	<u>48.0±0.9</u>	
Reddit	0.05	46.1±4.4	53.1±2.5	46.6±2.3	40.9±0.5	88.0±1.8	80.0±3.1	84.2±0.7	OOM	90.5±0.2	91.2±0.1	<b>92.9±0.3</b>	90.7±0.1	<u>91.3±0.1</u>	94.1±0.0
	0.1	58.0±2.2	62.7±1.0	53.0±3.3	42.8±0.8	89.6±0.7	84.6±1.6	80.6±0.4	OOM	91.2±0.2	<u>92.4±0.0</u>	<b>93.1±0.2</b>	91.0±0.1	<u>90.1±0.9</u>	
	0.2	66.3±1.9	71.0±1.6	58.5±2.1	47.4±0.9	90.1±0.5	87.9±1.2	84.1±0.3	OOM	92.2±0.1	<u>92.7±0.1</u>	<b>93.2±0.4</b>	90.5±0.0	90.8±0.5	



**Figure 2: Accuracy scores achieved by our proposed method on graphs with varying numbers of views. The values following each dataset name represent the condensation ratio applied in this ablation study. The experiments were run ten times and we report the average accuracy and the respective error bars.**

performance at higher ratios. These findings highlight the effectiveness of our method in transductive and inductive settings, as well as in large datasets (i.e., Flickr and Reddit). Although results on Ogbn-arxiv do not surpass the strongest baselines, GCTD maintains competitive performance and demonstrates robustness across diverse datasets, indicating potential for further development.

**On the number of graph views.** Next, we present an experiment we conducted to analyze the impact of the number of augmented views used in condensing the original graph. In this evaluation, we assessed our method on graphs with varying counts, ranging from 1 to 5. It is important to note that the single-view case corresponds to our first attempt, discussed in Section 4. Figure 2 shows how the number of views influences overall performance.

Our results indicate that, except for Cora, using a multi-view augmented graph yields better performance than the single-view one. For example, on the Ogbn-arxiv dataset, we observed a 31.5% performance improvement when using three views instead of just

one, highlighting the effectiveness of using augmented views.

**Performance across different GNNs.** We also evaluated the performance of various GNN architectures on the condensed graphs generated by GCTD and compared them to GCond, SFGC, and GCDM, which represent diverse approaches to graph condensation. Consistent with the evaluation protocol of Jin et al. [28] and other follow-up works, we trained 2-layer versions of APPNP [16], ChebyNet [8], SGC [51], and GraphSAGE [18] using the synthesized graphs and measured their performance on the original graph’s test set. For this evaluation, we selected datasets across three size categories: Citeseer and Cora (small), Pubmed (medium), and Flickr (large), with condensation ratios of 1.8%, 2.6%, 0.15%, and 0.5%, respectively. We report the average accuracy over ten runs in Table 3.


Focusing on GCTD, we observe that APPNP, SGC, and GCN generally achieve strong results, often surpassing the average performance across datasets. While APPNP performs best on Cora and

**Table 3: Accuracy (%) of different GNN architectures on the graphs condensed by our method, GCond, and SFGC. The reported values are an average of ten runs. Best and second-best averages are in bold and underlined, respectively.**

Datasets (r%)	Methods	Models					Avg.
		GCN	APPNP	Cheby	SGC	SAGE	
Cora (2.6)	GCond	80.1	78.5	76.0	79.3	78.4	78.5
	SFGC	79.3	78.8	79.0	79.1	80.0	79.2
	GCDM	80.5	79.7	75.5	77.7	77.5	78.2
	GCTD	84.0	85.5	76.2	82.5	74.6	<b>80.5</b>
Citeseer (1.8)	GCond	70.6	69.6	68.3	70.3	66.2	69.0
	SFGC	69.0	70.5	71.8	71.8	71.7	71.0
	GCDM	71.7	73.6	66.1	73.6	71.1	<u>71.2</u>
	GCTD	76.5	73.8	75.2	77.4	73.5	<b>75.3</b>
Pubmed (0.08)	GCond	77.1	76.8	75.9	77.1	76.9	76.8
	SFGC	77.5	76.3	77.7	77.8	76.3	<u>77.1</u>
	GCDM	77.1	78.3	71.5	74.9	77.0	75.8
	GCTD	79.9	79.2	78.0	79.8	79.0	<b>79.2</b>
Flickr (0.5)	GCond	47.1	45.9	42.8	46.1	46.2	45.6
	SFGC	46.0	40.7	45.4	42.5	47.0	44.3
	GCDM	46.4	46.0	42.4	45.8	42.6	44.6
	GCTD	48.1	46.9	42.7	45.4	47.2	<b>46.1</b>

SGC leads on Citeseer and Pubmed, GCN consistently stays competitive. GraphSAGE also stands out on Flickr, achieving the highest score among the GNNs for that dataset. In contrast, ChebyNet underperforms, with below-average results on most datasets.

Analyzing the baselines, we observe that GCTD outperforms all methods in terms of average accuracy across datasets. For instance, on Citeseer, GCTD improves over GCDM (the second-best method) by 4.1%. On Pubmed, it surpasses SFGC by 2.1%, and on Cora, it leads by 1.3%. Even on Flickr, where all methods struggle, GCTD still achieves the best average, showing a 0.5% gain over GCond.



**Figure 3: Accuracy of GCTD with K-Means and Argmax as the method employed to compute the synthetic node assignments from factor matrix U. In this ablation study, we used Citeseer, Cora, Pubmed, and Flickr with the condensation ratio set to 1.8%, 2.6%, 0.15%, and 0.5%, respectively.**

**Ablation on the synthetic node assignment.** We conducted an ablation study to evaluate the effectiveness of the synthetic node assignment method we leveraged. Specifically, we used the *argmax* operation on each row of the factor U, assigning the selected node to the corresponding synthetic node. We then compared the performance of this approach to that of K-Means, the default method

in GCTD. In this ablation, we follow the same settings used in the previous experiment, employing Citeseer, Cora, Pubmed, and Flickr with 1.8%, 2.6%, 0.15%, and 0.5% condensation ratios, respectively.

As shown in Figure 3, while *argmax* provides a straightforward method for generating synthetic nodes, its simplicity can sometimes limit its ability to produce high-quality synthetic nodes. Specifically, we can observe that only selecting the maximum value in each row as the cluster membership is insufficient to capture the necessary information for the condensed graph. On the other hand, even though KMeans is computationally complex compared to *argmax*, it offers greater performance. It typically obtains better clustering of nodes by effectively grouping them into synthetic nodes based on the optimized factor matrix. This matrix, derived during the decomposition process, captures co-clustered nodes exhibiting similar patterns, leading to more meaningful synthetic representations [17].

**Table 4: Running time analysis of the proposed method and the SGC version of the baselines. The experiments were performed on a single NVIDIA RTX A6000. OOM occurred at 49GB of memory usage.**

Dataset (r%)	GCond	SGDD	SFGC	GCDM	SNTK	GCTD
Citeseer (0.9)	70.8	140.9	63.0	131.3	2.9	32.1
Cora (1.3)	90.3	115.1	62.4	174.7	3.1	42.8
Pubmed (0.08)	57.1	2200.5	61.6	75.4	4.2	48.1
Ogbn-arxiv (0.05)	725.0	2073.3	238.0	480.0	925.4	965.6
Flickr (0.1)	470.4	361.5	42.9	225.0	148.7	222.9
Reddit (0.05)	3780	10350	2420	1617.1	OOM	2515.3

**Running time.** In Table 4, we compare the running time of our method against GCond, SGDD, SFGC, GCDM, and SNTK. To ensure a fair comparison, we adopted the SGC implementation for all baselines, as it offers faster runtime than their GCN counterparts. Overall, our method strikes a balance between speed and performance. It consistently outperforms SGDD across all datasets, significantly reducing condensation time. For instance, on Citeseer, our method is over four times faster than SGDD and twice as fast as SFGC, one of the fastest baseline models. On larger datasets our method remains considerably faster than SGDD.

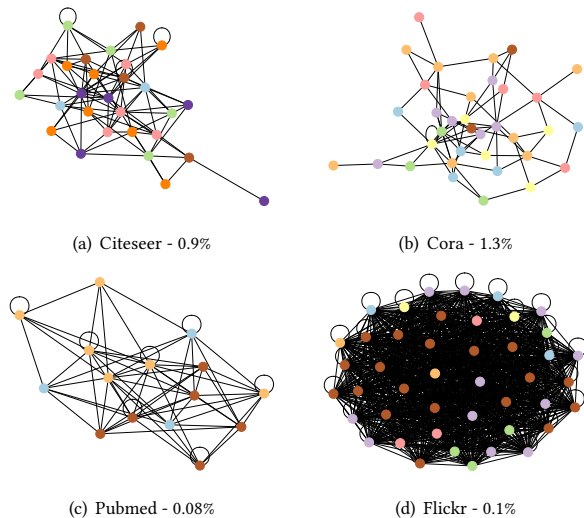
When compared to GCond and GCDM, our method shows better efficiency on smaller datasets and Flickr, while also outperforming GCond on Reddit. It is, however, slightly slower than GCDM on Reddit. Notably, SFGC remains among the fastest methods due to its structure-free condensation approach, which aligns with the findings of Han et al. [19] that matrix multiplication with the adjacency matrix is the most time-consuming operation in GNNs. Lastly, while SNTK performs well on smaller graphs, its runtime advantage diminishes as graph size increases. On Reddit, it runs out of memory (49GB in our setup) due to its expensive kernel computation. These results highlight that our method scales well across diverse datasets, delivering competitive performance in graph condensation without incurring excessive computational costs.

**Visualization of the condensed graphs.** We visualize the condensed graphs generated with the smallest condensation ratio for Citeseer, Cora, Pubmed, and Flickr in Figure 4. Several key observations emerge from these visualizations. First, GCTD produces

**Table 5: Comparison between the condensed graphs generated by GCTD and original graphs.**

	Citeseer (0.9%)		Cora (1.3%)		Pubmed (0.08%)		Flickr (0.1%)		Ogbn-arxiv (0.05%)	
	Whole	GCTD	Whole	GCTD	Whole	GCTD	Whole	GCTD	Whole	GCTD
Accuracy	71.7	76.8	81.4	81.4	77.1	79.9	47.1	48.4	71.3	58.2
#Nodes	3,327	30	2,708	35	19,717	15	44,625	44	169,343	84
#Edges	4,732	93	5,429	72	44,338	67	218,140	990	1,116,243	37
Sparsity	0.09%	22.91%	0.15%	12.1%	0.01%	63.8%	0.02%	100%	0.09%	1.1%
Storage	47.1 MB	0.51 MB	14.9 MB	0.23 MB	40 MB	0.05 MB	86.8 MB	0.17 MB	100.4 MB	0.08 MB

graphs with well-defined structures and introduces self-loops for certain nodes. Additionally, it reduces the homophily present in the original graphs, indicating a diminished reliance on this property. This is particularly intriguing, given that the GNNs we used typically depend on graph homophily for optimal performance, as discussed in prior work [65]. Finally, GCTD generates a complete graph for Flickr, implying that the GNN now relies on information from every node to compute individual node representations. As a result, the GNN relies on the weights and features of nodes to differentiate between information from various neighbors.



**Figure 4: Visualization of condensed graphs generated by GCTD. Each node represents a synthetic node, with its color indicating the corresponding class.**

**Comparison with original graphs.** Table 5 presents a comparison of various properties between the original graphs and the condensed graphs generated by GCTD. The results show that GCTD achieves comparable performance on Citeseer, Cora, Pubmed, and Flickr, while significantly reducing the number of nodes and edges, as well as requiring less storage. Furthermore, the condensed graphs are denser than their original counterparts, and a notable behavior is observed on Flickr, where the learned graph is complete.

## 7 Conclusion

In this paper, we introduced GCTD, a novel framework for graph condensation through the decomposition of multi-view augmented graphs. Moreover, we also showed the extent to which tensor decomposition methods can generate smaller graphs while maintaining the performance of GNNs on downstream tasks. We evaluated GCTD on six real-world datasets spanning transductive and inductive learning settings. Additionally, we conducted an in-depth analysis to demonstrate the effectiveness of our approach.

Our results indicate that GCTD outperforms existing methods on three of the six datasets and achieves competitive performance on the larger datasets. Furthermore, our proposed method performs well on transductive and inductive datasets. For future work, we plan to explore different decomposition techniques to generate condensed graphs. In addition, we will explore different ways of performing augmentation to improve the quality of the synthesized graph. Lastly, we plan to investigate whether replacing the current hard assignment with a soft membership approach can improve the effectiveness of GCTD.

## Acknowledgments

Research was supported in part by the National Science Foundation under CAREER grant no. IIS 2046086, grant no. 2431569 and CREST Center for Multidisciplinary Research Excellence in Cyber-Physical Infrastructure Systems (MECIS) grant no. 2112650, and by the Agriculture and Food Research Initiative Competitive Grant no. 2020-69012-31914 from the USDA National Institute of Food and Agriculture. This study was also supported by the Coordination for the Improvement of Higher Education Personnel (CAPES) through the Institutional Internationalization Program (PRINT), Call No. 41/2017, and partially funded by CAPES (Finance Code 001), the São Paulo Research Foundation (FAPESP) [grants 20/09835-1, 22/02176-8, and 22/09091-8], and the National Council for Scientific and Technological Development (CNPq) [grants 303588/2022-5 and 406417/2022-9]. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the funding agencies. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes not withstanding any copyright notation here on.

## References

- [1] Mohammadreza Babae, Stefanos Tsoukalas, Maryam Babae, Gerhard Rigoll, and Mihai Datcu. 2016. Discriminative Nonnegative Matrix Factorization for

- dimensionality reduction. *Neurocomputing* 173 (2016), 212–223.
- [2] Joshua Batson, Daniel A. Spielman, Nikhil Srivastava, and Shang-Hua Teng. 2013. Spectral sparsification of graphs: theory and algorithms. *Commun. ACM* 56, 8 (aug 2013), 87–94.
  - [3] Lukas Biewald. 2020. Experiment Tracking with Weights and Biases. <https://www.wandb.com/> Software available from wandb.com.
  - [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*. PMLR, 1597–1607.
  - [5] Minsik Cho, Vinod Muthusamy, Brad Nemanich, and Ruchir Puri. 2019. Gradzip: Gradient compression using alternating matrix factorization for large-scale deep learning. In *NeurIPS*.
  - [6] Luciano da Fontoura Costa, Osvaldo N Oliveira Jr, Gonzalo Travieso, Francisco Aparecido Rodrigues, Paulino Ribeiro Villas Boas, Lucas Antiquiera, Matheus Palhares Viana, and Luis Enrique Correa Rocha. 2011. Analyzing and modeling real-world phenomena with complex networks: a survey of applications. *Advances in Physics* 60, 3 (2011), 329–412.
  - [7] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. 2019. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 113–123.
  - [8] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. 2016. Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering. In *NeurIPS*. Curran Associates, Inc., 3844–3852.
  - [9] Austin Derrrow-Pinion, Jennifer She, David Wong, Oliver Lange, Todd Hester, Luis Perez, Marc Nunkesser, Seongjae Lee, Xueying Guo, Brett Wiltshire, Peter W. Battaglia, Vishal Gupta, Ang Li, Zhongwen Xu, Alvaro Sanchez-Gonzalez, Yujia Li, and Petar Velickovic. 2021. ETA Prediction with Graph Neural Networks in Google Maps. In *CIKM*. ACM.
  - [10] Charles Dickens, Edward Huang, Aishwarya Reganti, Jiong Zhu, Karthik Subbian, and Danaï Koutra. 2024. Graph Coarsening via Convolution Matching for Scalable Graph Neural Network Training. In *Companion Proceedings of the ACM Web Conference 2024* (Singapore, Singapore) (*WWW '24*). Association for Computing Machinery, New York, NY, USA, 1502–1510.
  - [11] Chris Ding, Tao Li, Wei Peng, and Haesun Park. 2006. Orthogonal nonnegative matrix t-factorizations for clustering. In *KDD* (Philadelphia, PA, USA). Association for Computing Machinery, New York, NY, USA, 126–135.
  - [12] Kaize Ding, Zhe Xu, Hanghang Tong, and Huan Liu. 2022. Data augmentation for deep graph learning: A survey. *ACM SIGKDD Explorations Newsletter* 24, 2 (2022), 61–77.
  - [13] Nicolas Roque dos Santos, Diego Minatel, Alan Demétrius Baria Valejo, and Alneu de Andrade Lopes. 2024. Semi-Supervised Coarsening of Bipartite Graphs for Text Classification via Graph Neural Network. In *2024 IEEE 11th International Conference on Data Science and Advanced Analytics (DSAA)*. 1–10. doi:10.1109/DSAA61799.2024.10722822
  - [14] Steven Y. Feng, Varun Gangal, Jason Wei, Sarath Chandar, Soroush Vosoughi, Teruko Mitamura, and Eduard Hovy. 2021. A Survey of Data Augmentation Approaches for NLP. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*. Association for Computational Linguistics, Online, 968–988.
  - [15] Xinyi Gao, Guanhua Ye, Tong Chen, Wentao Zhang, Junliang Yu, and Hongzhi Yin. 2025. Rethinking and Accelerating Graph Condensation: A Training-Free Approach with Class Partition. In *Proceedings of the ACM on Web Conference 2025* (Sydney NSW, Australia) (*WWW '25*). Association for Computing Machinery, New York, NY, USA, 4359–4373.
  - [16] Johannes Gastegger, Aleksandar Bojchevski, and Stephan Günnemann. 2019. Predict then Propagate: Graph Neural Networks meet Personalized PageRank. In *International Conference on Learning Representations (ICLR)*.
  - [17] Ekta Gujral and Evangelos E. Papalexakis. 2018. SMACD: Semi-supervised Multi-Aspect Community Detection. In *Proceedings of the 2018 SIAM International Conference on Data Mining (SDM)*. 702–710.
  - [18] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. *NeurIPS* 30 (2017).
  - [19] Xiaotian Han, Tong Zhao, Yozen Liu, Xia Hu, and Neil Shah. 2023. MLPinit: Embarrassingly Simple GNN Training Acceleration with MLP Initialization. In *The Eleventh International Conference on Learning Representations*.
  - [20] 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. *International Joint Conference on Artificial Intelligence (IJCAI)* (2024).
  - [21] Minguo He, Zhewei Wei, Zengfeng Huang, and Hongteng Xu. 2021. BernNet: Learning Arbitrary Graph Spectral Filters via Bernstein Approximation. In *Advances in Neural Information Processing Systems*, A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (Eds.).
  - [22] Keith Henderson, Brian Gallagher, Tina Eliassi-Rad, Hanghang Tong, Sugato Basu, Leman Akoglu, Danaï Koutra, Christos Faloutsos, and Lei Li. 2012. RoIX: structural role extraction & mining in large graphs. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Beijing, China) (*KDD '12*). Association for Computing Machinery, New York, NY, USA, 1231–1239.
  - [23] 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. *Advances in neural information processing systems* 33 (2020), 22118–22133.
  - [24] Chenqing Hua, Guillaume Rabusseau, and Jian Tang. 2022. High-order pooling for graph neural networks with tensor decomposition. *Advances in Neural Information Processing Systems* 35 (2022), 6021–6033.
  - [25] Kejun Huang, Nicholas D. Sidiropoulos, and Ananthram Swami. 2014. Non-Negative Matrix Factorization Revisited: Uniqueness and Algorithm for Symmetric Decomposition. *IEEE Transactions on Signal Processing* 62, 1 (2014), 211–224. doi:10.1109/TSP.2013.2285514
  - [26] Zengfeng Huang, Shengzhong Zhang, Chong Xi, Tang Liu, and Min Zhou. 2021. Scaling up graph neural networks via graph coarsening. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*. 675–684.
  - [27] Wei Jin, Xianfeng Tang, Haoming Jiang, Zheng Li, Danqing Zhang, Jiliang Tang, and Bing Yin. 2022. Condensing graphs via one-step gradient matching. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 720–730.
  - [28] Wei Jin, Lingxiao Zhao, Shichang Zhang, Yozen Liu, Jiliang Tang, and Neil Shah. 2022. Graph Condensation for Graph Neural Networks. In *International Conference on Learning Representations*.
  - [29] David R. Karger. 1994. Random sampling in cut, flow, and network design problems. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing* (Montreal, Quebec, Canada) (*STOC '94*). Association for Computing Machinery, New York, NY, USA, 648–657.
  - [30] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *ICLR, 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
  - [31] Tamara G Kolda and Brett W Bader. 2009. Tensor decompositions and applications. *SIAM review* 51, 3 (2009), 455–500.
  - [32] Da Kuang, Chris Ding, and Haesun Park. [n. d.]. *Symmetric Nonnegative Matrix Factorization for Graph Clustering*. 106–117.
  - [33] Zechao Li, Jing Liu, and Hanqing Lu. 2013. Structure preserving non-negative matrix factorization for dimensionality reduction. *Computer Vision and Image Understanding* 117, 9 (2013), 1175–1189.
  - [34] Mengyang Liu, Shanchuan Li, Xinshi Chen, and Le Song. 2022. Graph condensation via receptive field distribution matching. *arXiv preprint arXiv:2206.13697* (2022).
  - [35] Yang Liu, Deyu Bo, and Chuan Shi. 2024. Graph Distillation with Eigenbasis Matching. In *Forty-first International Conference on Machine Learning*.
  - [36] Zhanyu Liu, Chaoly Zeng, and Guanjie Zheng. 2024. Graph Data Condensation via Self-expressive Graph Structure Reconstruction. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (Barcelona, Spain) (*KDD '24*). Association for Computing Machinery, New York, NY, USA, 1992–2002.
  - [37] Andreas Loukas. 2019. Graph reduction with spectral and cut guarantees. *Journal of Machine Learning Research* 20, 116 (2019), 1–42.
  - [38] Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. 2011. A three-way model for collective learning on multi-relational data. In *Proceedings of the 28th International Conference on International Conference on Machine Learning* (Bellevue, Washington, USA) (*ICML '11*). Omnipress, Madison, WI, USA, 809–816.
  - [39] Evangelos E Papalexakis, Leman Akoglu, and Dino Ienco. 2013. Do more views of a graph help? community detection and clustering in multi-graphs. In *Proceedings of the 16th International Conference on Information Fusion*. IEEE, 899–905.
  - [40] Stephan Rabanser, Oleksandr Shchur, and Stephan Günnemann. 2017. Introduction to tensor decompositions and their applications in machine learning. *arXiv preprint arXiv:1711.10781* (2017).
  - [41] Yu Rong, Wenbing Huang, Tingyang Xu, and Junzhou Huang. 2020. DropEdge: Towards Deep Graph Convolutional Networks on Node Classification. In *International Conference on Learning Representations*.
  - [42] Ozan Sener and Silvio Savarese. 2018. Active Learning for Convolutional Neural Networks: A Core-Set Approach. In *International Conference on Learning Representations*.
  - [43] Connor Shorten and Taghi M Khoshgoftaar. 2019. A survey on image data augmentation for deep learning. *Journal of big data* 6, 1 (2019), 1–48.
  - [44] Nicholas D. Sidiropoulos, Lieven De Lathauwer, Xiao Fu, Kejun Huang, Evangelos E. Papalexakis, and Christos Faloutsos. 2017. Tensor Decomposition for Signal Processing and Machine Learning. *IEEE Transactions on Signal Processing* 65, 13 (2017), 3551–3582. doi:10.1109/TSP.2017.2690524
  - [45] Chenguang Song, Kai Shu, and Bin Wu. 2021. Temporally evolving graph neural network for fake news detection. *Information Processing & Management* 58, 6 (2021), 102712.
  - [46] Jonathan M Stokes, Kevin Yang, Kyle Swanson, Wengong Jin, Andres Cubillos-Ruiz, Nina M Donghia, Craig R MacNair, Shawn French, Lindsey A Carfrae, Zohar Bloom-Ackermann, et al. 2020. A deep learning approach to antibiotic discovery. *Cell* 180, 4 (2020), 688–702.

- [47] Sridhar Swaminathan, Deepak Garg, Rajkumar Kannan, and Frederic Andres. 2020. Sparse low rank factorization for deep neural network compression. *Neurocomputing* 398 (2020), 185–196.
- [48] Brandon Trabucco, Kyle Doherty, Max A Gurinas, and Ruslan Salakhutdinov. 2024. Effective Data Augmentation With Diffusion Models. In *The Twelfth International Conference on Learning Representations*.
- [49] Lin Wang, Wenqi Fan, Jiatong Li, Yao Ma, and Qing Li. 2024. Fast graph condensation with structure-based neural tangent kernel. In *Proceedings of the ACM on Web Conference 2024*. 4439–4448.
- [50] Max Welling. 2009. Herding dynamical weights to learn. In *Proceedings of the 26th annual international conference on machine learning*. 1121–1128.
- [51] Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. 2019. Simplifying graph convolutional networks. In *ICML*. PMLR, 6861–6871.
- [52] Zhebin Wu, Lin Shu, Ziyue Xu, Yaomin Chang, Chuan Chen, and Zibin Zheng. 2022. Robust Tensor Graph Convolutional Networks via T-SVD based Graph Augmentation. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (Washington DC, USA) (*KDD '22*). Association for Computing Machinery, New York, NY, USA, 2090–2099.
- [53] Zhenbang Xiao, Yu Wang, Shunyu Liu, Bingde Hu, Huiqiong Wang, Mingli Song, and Tongya Zheng. 2025. Disentangled Condensation for Large-scale Graphs. In *Proceedings of the ACM on Web Conference 2025* (Sydney NSW, Australia) (*WWW '25*). Association for Computing Machinery, New York, NY, USA, 4494–4506.
- [54] Zhenbang Xiao, Yu Wang, Shunyu Liu, Huiqiong Wang, Mingli Song, and Tongya Zheng. 2024. Simple Graph Condensation. In *Machine Learning and Knowledge Discovery in Databases. Research Track*, Albert Bifet, Jesse Davis, Tomas Krilavicius, Meelis Kull, Eirini Ntoutsi, and Indrè Žliobaitė (Eds.). Springer Nature Switzerland, Cham, 53–71.
- [55] Wei Xu, Xin Liu, and Yihong Gong. 2003. Document clustering based on non-negative matrix factorization. In *Proceedings of the 26th annual international ACM SIGIR conference on Research and development in informaion retrieval*. 267–273.
- [56] Zhe Xu, Yuzhong Chen, Menghai Pan, Huiyuan Chen, Mahashweta Das, Hao Yang, and Hanghang Tong. 2023. Kernel Ridge Regression-Based Graph Dataset Distillation. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (Long Beach, CA, USA) (*KDD '23*). Association for Computing Machinery, New York, NY, USA, 2850–2861.
- [57] Beining Yang, Kai Wang, Qingyun Sun, Cheng Ji, Xingcheng Fu, Hao Tang, Yang You, and Jianxin Li. 2023. Does Graph Distillation See Like Vision Dataset Counterpart?. In *NeurIPS*.
- [58] Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, and Yang Shen. 2020. Graph contrastive learning with augmentations. *Advances in neural information processing systems* 33 (2020), 5812–5823.
- [59] Hanqing Zeng, Hongkuan Zhou, Ajitesh Srivastava, Rajgopal Kannan, and Viktor Prasanna. 2020. GraphSAINT: Graph Sampling Based Inductive Learning Method. In *ICLR*.
- [60] Yuchen Zhang, Tianle Zhang, Kai Wang, Ziyao Guo, Yuxuan Liang, Xavier Bresson, Wei Jin, and Yang You. 2024. Navigating Complexity: Toward Lossless Graph Condensation via Expanding Window Matching. *ICML 2024* (2024).
- [61] Bo Zhao and Hakan Bilen. 2023. Dataset condensation with distribution matching. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 6514–6523.
- [62] Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. 2021. Dataset Condensation with Gradient Matching. In *International Conference on Learning Representations*.
- [63] Tong Zhao, Yozen Liu, Leonardo Neves, Oliver Woodford, Meng Jiang, and Neil Shah. 2021. Data augmentation for graph neural networks. In *Proceedings of the aaai conference on artificial intelligence*, Vol. 35. 11015–11023.
- [64] Xin Zheng, Miao Zhang, Chunyang Chen, Quoc Viet Hung Nguyen, Xingquan Zhu, and Shirui Pan. 2023. Structure-free Graph Condensation: From Large-scale Graphs to Condensed Graph-free Data. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- [65] Jiong Zhu, Yujun Yan, Lingxiao Zhao, Mark Heimann, Leman Akoglu, and Danai Koutra. 2020. Beyond homophily in graph neural networks: Current limitations and effective designs. *Advances in neural information processing systems* 33 (2020), 7793–7804.
- [66] Unai Zulaika, Aitor Almeida, and Diego Lopez-de Ipina. 2023. Regularized online tensor factorization for sparse knowledge graph embeddings. *Neural Computing and Applications* 35, 1 (2023), 787–797.