



Datasets, tasks, and training methods for large-scale hypergraph learning

Sunwoo Kim¹ · Dongjin Lee² · Yul Kim³ · Jungho Park³ · Taeho Hwang³ · Kijung Shin^{1,2} 

Received: 10 February 2023 / Accepted: 28 June 2023 / Published online: 26 July 2023

© The Author(s), under exclusive licence to Springer Science+Business Media LLC, part of Springer Nature 2023

Abstract

Relations among multiple entities are prevalent in many fields, and hypergraphs are widely used to represent such group relations. Hence, machine learning on hypergraphs has received considerable attention, and especially much effort has been made in neural network architectures for hypergraphs (a.k.a., hypergraph neural networks). However, existing studies mostly focused on small datasets for a few single-entity-level downstream tasks and overlooked scalability issues, although most real-world group relations are large-scale. In this work, we propose new tasks, datasets, and scalable training methods for addressing these limitations. First, we introduce two pair-level hypergraph-learning tasks to formulate a wide range of real-world problems. Then, we build and publicly release two large-scale hypergraph datasets with tens of millions of nodes, rich features, and labels. After that, we propose PCL, a scalable learning method for hypergraph neural networks. To tackle scalability issues, PCL splits a given hypergraph into partitions and trains a neural network via contrastive learning. Our extensive experiments demonstrate that hypergraph neural networks can be trained for large-scale hypergraphs by PCL while outperforming 16 baseline models. Specifically, the performance is comparable, or surprisingly even better than that achieved by training hypergraph neural networks on the entire hypergraphs without partitioning.

Keywords Large-scale hypergraph datasets · Scalable hypergraph learning · Hypergraph neural networks · Contrastive learning · Partitioning

Sunwoo Kim, Dongjin Lee have contributed equally to this work.

Responsible editor: Charalampos Tsourakakis.

Extended author information available on the last page of the article

1 Introduction

Beyond pairwise relations among entities, understanding and modeling higher-order relations have received considerable attention (Benson et al. 2018; Lee et al. 2021; Ko et al. 2022; Luo et al. 2021; Torres et al. 2021; Yin et al. 2022; Li et al. 2022; Kim et al. 2022). A hypergraph, which generalizes a graph, is a data structure that is commonly used to model such higher-order relations (Konstantinova and Skorobogatov 2001; Klamt et al. 2009; Qu et al. 2018). While an edge in a graph joins only two entities, a hyperedge joins an arbitrary number of entities, which makes a hypergraph, which is a pair of a node set and a hyperedge set, inherently capable of capturing high-order relations.

Due to the omnipresence of hypergraphs, a number of machine learning tasks, such as node classification (Huang and Yang 2021; Yadati et al. 2019; Chien et al. 2021; Lee and Shin 2023) and hyperedge prediction (Yadati et al. 2020; Hwang et al. 2022; Zhang et al. 2020b), have been considered for hypergraph-structured data. One possible approach to tackle such tasks is to transform hypergraphs into ordinary graphs and apply existing graph representation models (e.g., graph neural networks) (Kipf and Welling 2017; Hamilton et al. 2017; Veličković et al. 2018). On the other hand, previous studies have shown that this transition can result in the loss of significant higher-order information, causing performance degradation in machine learning tasks (Dong et al. 2020; Feng et al. 2019; Yadati et al. 2020; Zhang et al. 2020b). These studies highlight the need for specialized representation models specifically designed for hypergraphs.

Numerous *hypergraph neural network models* (e.g., HGNN (Feng et al. 2019), NHP (Yadati et al. 2020), UniGCNII (Huang and Yang 2021), and AllSet (Chien et al. 2021)) have been developed in recent years. Despite these advancements, the evaluation of these models has been limited to small datasets, usually consisting of tens of thousands of nodes and a few single-entity level downstream tasks. Moreover, scalability for large-scale datasets has been overlooked in prior research, with the majority of studies focusing on enhancing the expressive power of these networks.

However, many real-world applications require predicting properties beyond the single-entity level, such as pairs or groups of entities. Examples include detecting collusion among users in peer review (Jecmen et al. 2023), recommending products to users (Koren et al. 2009), identifying the same users on online social networks (Zhang and Tong 2016), and predicting pairwise differences for chemical discovery (Tynes et al. 2021), to name a few. In addition, a massive number of group relations can be found across various domains, including over 200 million co-authorship relations in research papers (Zhang et al. 2019), co-appearances of hashtags in over 500 million posts on social media (Yang and Leskovec 2011), and over 10 million group discussions on an online Q & A platform (Benson et al. 2018).

To bridge the apparent gap between previous studies and practical applications, we present new tasks, datasets, and scalable training methods for large-scale hypergraph learning. Our contributions toward these goals are summarized as follows:

- *Novel Pair-level Tasks:* We present two new pair-level prediction tasks, hyper-edge disambiguation and local clustering, and demonstrate that they can be used to address various real-world problems.
- *Large-scale Datasets:* We construct and publicly release two large-scale hypergraph-structured datasets: AMINER and MAG, which contain 10 million nodes and 20 million nodes respectively. These datasets are equipped with rich features and labels.
- *Scalable Training Schemes:* We propose PCL (**P**artitioning-based **C**ontrastive **L**earning), a scalable learning method for hypergraph neural networks (HNNs).

In a nutshell, PCL partitions the input hypergraph and trains HNNs via contrastive learning while loading only one partition into memory at a time. PCL is also equipped with additional techniques to reduce information loss due to partitioning, and as a result, HNNs trained by PCL show surprisingly good performance in our experiments. Specifically, the performance is comparable to and often even better than that achieved by training HNNs on entire hypergraphs without partitioning, which is not scalable to large-scale datasets.

The structure of this paper is as follows. In Sect. 2, we provide some preliminaries and related work. In Sect. 3, we propose two pair-level prediction tasks with their mathematical formulation and real-world applications. In Sect. 4, we describe the two large-scale hypergraph datasets that we build. In Sect. 5, we present PCL, our scalable learning method for hypergraph neural networks. Using all the above, we perform experiments and report the results in Sect. 6. In Sect. 7, we offer a conclusion of our work.

2 Preliminaries and related work

2.1 Preliminaries

We first give some preliminaries on hypergraphs and hypergraph neural networks. See Table 1 for a list of the frequently-used symbols.

2.1.1 Hypergraphs and notations

A hypergraph generalizes a graph by allowing edges to join an arbitrary number of nodes. Consider a *hypergraph* $\mathcal{H} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$ where $\mathcal{V} = \{v_1, v_2, \dots, v_{|\mathcal{V}|}\}$ is the set of nodes, $\mathcal{E} = \{e_1, e_2, \dots, e_{|\mathcal{E}|}\}$ is the set of hyperedges, and $\mathbf{X} \in \mathbb{R}^{|\mathcal{V}| \times F}$ is the node feature matrix. Each hyperedge $e \in \mathcal{E}$ is a non-empty subset of \mathcal{V} (i.e., $e \subseteq \mathcal{V}$ and $e \neq \emptyset$), and we use $\mathbf{x}_i = \mathbf{X}[i, :]^T \in \mathbb{R}^F$ to denote the feature vector of the node v_i . The topological information in a hypergraph can also be represented in the form of a matrix called an *incidence matrix*. In the incidence matrix $\mathbf{H} \in \{0, 1\}^{|\mathcal{V}| \times |\mathcal{E}|}$ of \mathcal{H} , each (i, j) -th entry has a value of 1 if a node v_i is incident to a hyperedge e_j (i.e., $h_{ij} = 1$ if $v_i \in e_j$), or has a value of zero (i.e., $h_{ij} = 0$) otherwise. That is, $\mathcal{H} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$ can be denoted equivalently by $\mathcal{H} = (\mathbf{X}, \mathbf{H})$.

Table 1 Frequently used symbols

Notation	Definition
$\mathcal{H} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$	Hypergraph with nodes \mathcal{V} , hyperedges \mathcal{E} , and node features \mathbf{X}
$\mathcal{H} = (\mathbf{X}, \mathbf{H})$	Hypergraph with node features \mathbf{X} and the incidence matrix \mathbf{H}
$v_i \in \mathcal{V}, e_j \in \mathcal{E}$	Node and hyperedge
$\mathbf{X} \in \mathbb{R}^{ \mathcal{V} \times F}$	Node feature matrix
$\mathbf{H} \in \{0, 1\}^{ \mathcal{V} \times \mathcal{E} }$	Incidence matrix of \mathcal{H}
$Enc_\theta(\cdot)$	Hypergraph encoder
$C_k \subseteq \mathcal{V}$	Cluster of nodes
$\mathcal{P}_k = (\mathcal{P}_k^{\mathcal{V}}, \mathcal{P}_k^{\mathcal{E}})$	Partition with nodes $\mathcal{P}_k^{\mathcal{V}}$ and hyperedges $\mathcal{P}_k^{\mathcal{E}}$
$\mathcal{P}_k = (\mathcal{P}_k^{\mathbf{X}}, \mathcal{P}_k^{\mathbf{H}})$	Partition with node features $\mathcal{P}_k^{\mathbf{X}}$ and the incidence matrix $\mathcal{P}_k^{\mathbf{H}}$
$ A $	Cardinality (i.e., number of elements) of a set A

2.1.2 Hypergraph neural networks (HNNs)

Hypergraph neural networks (HNNs) use the hypergraph structure \mathbf{H} , node features \mathbf{X} , and (optionally) hyperedge features \mathbf{Y} to learn representations of nodes \mathbf{P} and/or of hyperedges \mathbf{Q} . Most of the modern HNNs (Feng et al. 2019; Yadati et al. 2019; Bai et al. 2021; Dong et al. 2020; Arya et al. 2020; Chien et al. 2021) follow a two-stage message aggregation strategy: node-to-hyperedge and hyperedge-to-node message aggregation. They iteratively update (a) The representation of a hyperedge by aggregating messages from its incident nodes and (b) The representation of a node by aggregating messages from its incident hyperedges. Let $\mathbf{P}^{(k)} \in \mathbb{R}^{|\mathcal{V}| \times F'}$ and $\mathbf{Q}^{(k)} \in \mathbb{R}^{|\mathcal{E}| \times F''}$ be the hidden representations of the nodes and hyperedges at the k th layer, respectively. Formally, the k -th layer of HNNs uses the following update rules:

$$\begin{aligned}
 \mathbf{q}_j^{(k)} &= f_{\mathcal{V} \rightarrow \mathcal{E}}^{(k)}\left(\mathbf{q}_j^{(k-1)}, \{\mathbf{p}_i^{(k-1)} : v_i \in e_j\}\right), \\
 \mathbf{p}_i^{(k)} &= f_{\mathcal{E} \rightarrow \mathcal{V}}^{(k)}\left(\mathbf{p}_i^{(k-1)}, \{\mathbf{q}_j^{(k)} : v_i \in e_j\}\right),
 \end{aligned}
 \tag{1}$$

where $f_{\mathcal{V} \rightarrow \mathcal{E}}(\cdot)$ and $f_{\mathcal{E} \rightarrow \mathcal{V}}(\cdot)$ are the message aggregation functions, and the initial node representation is identical to the node feature vector (i.e., for a node v_i , $\mathbf{z}_i^{(0)} = \mathbf{x}_i$ holds). Throughout this paper, we use HGNN (Feng et al. 2019) whose message aggregation strategies are as follows:

$$\mathbf{q}_j^{(k)} = \sum_{v_i \in e_j} \frac{\mathbf{p}_i^{(k-1)}}{\sqrt{d_i}}, \quad \mathbf{p}_i^{(k)} = \sigma\left(\frac{1}{\sqrt{d_i}} \sum_{e_j: v_i \in e_j} \frac{w_j \mathbf{q}_j^{(k)} \Theta^{(k)}}{\delta_j} + \mathbf{b}^{(k)}\right),
 \tag{2}$$

where (a) $\Theta^{(k)}$ is a learnable weight matrix, (b) $\mathbf{b}^{(k)}$ is a learnable bias, (c) d_i and δ_j are the degrees of each node v_i and hyperedge e_j , respectively, and (d) w_j is a positive weight assigned to each hyperedge $e_j \in \mathcal{E}$. For each node $v_i \in \mathcal{V}$, its degree is defined as $d_i = \sum_{j=1}^{|\mathcal{E}|} w_j h_{ij}$, and for each hyperedge $e_j \in \mathcal{E}$, its degree is defined as

$\delta_j = \sum_{i=1}^{|\mathcal{V}|} h_{ij}$. In this work, w_j is fixed to 1 for simplicity. In general, node and hyperedge representations at the last K -th layer are considered as final representations used for downstream tasks.

2.2 Related work

In this subsection, we review previous studies on hypergraph neural networks, scalable (hyper)graph learning, contrastive learning, and (hyper)graph partitioning, all of which are closely related to our work.

2.2.1 Hypergraph neural networks

There has been intense attention on designing message aggregation rules of hypergraph neural networks (HNN): $f_{\mathcal{V} \rightarrow \mathcal{E}}(\cdot)$ and $f_{\mathcal{E} \rightarrow \mathcal{V}}(\cdot)$ (see Eq. (1)). Many recent studies have focused on extending the applicability of graph neural networks (GNNs) to hypergraphs. Some approaches (Feng et al. 2019; Bai et al. 2021; Yadati et al. 2020) replace each hyperedge by a clique composed of its constituent nodes (i.e., clique-expansion) and employ GNN-based message passing on the resulting graph, which is called the clique-expanded graph. While these models are simple and effective, they suffer from undesired information loss due to structural distortion caused by clique expansion (Hein et al. 2013; Li and Milenkovic 2018; Dong et al. 2020). To mitigate such information loss, HNNH (Dong et al. 2020) utilizes an approach based on star-expansion, which does not lead to any information loss, with two different weight matrices for node- and hyperedge-side message aggregations. Several studies attempt to generalize the message-passing process in GNNs and HNNs in a unified form (Huang and Yang 2021; Zhang et al. 2022), and AllSet (Chien et al. 2021) generalizes message aggregation methods as multiset functions that are learned by DeepSets (Zaheer et al. 2017) or Set Transformer (Lee et al. 2019).

While many HNN models have been developed, the evaluation of their performance in most studies has primarily focused on single-entity-level prediction tasks, such as node classification (Feng et al. 2019; Huang and Yang 2021; Chien et al. 2021; Lee and Shin 2023), hyperedge classification (Dong et al. 2020), and hyperedge prediction (Yadati et al. 2020; Hwang et al. 2022). Although these tasks are commonly used as benchmarks for machine-learning models, it is important to note that many real-world applications may not inherently align with these tasks, as explained in greater detail in Sect. 3.3.

2.2.2 Scalable (Hyper)graph learning

As real-world graphs grow larger, many studies have been conducted to scale graph neural networks (GNNs) to large graphs through parallelism (Zhang et al. 2020a; Zheng et al. 2020, 2022a), graph sampling (Hamilton et al. 2017; Chen et al. 2018a, b; Huang et al. 2018; Chiang et al. 2019; Zeng et al. 2019), and pre-computed convolutional filters (Wu et al. 2019; Rossi et al. 2020). Here, we focus on sampling-based

approaches. Graph sampling, which approximates local graph structures by subsampled ones suitable for computation, has been demonstrated to be effective for scalable graph learning. For instance, GraphSAGE (Hamilton et al. 2017) utilizes uniform sampling of a fixed-size set of neighboring nodes to approximate local connectivity. Similarly, FastGCN (Chen et al. 2018a) conducts node-level sampling independently for each layer while incorporating importance sampling to reduce variance. Instead of node-wise sampling, some works build mini-batches by graph-level sampling. In ClusterGCN (Chiang et al. 2019), non-overlapping clusters are computed at the pre-processing phase and form mini-batches cluster by cluster. Inter-cluster edges are simply disregarded, and this process enables ClusterGCN to avoid the “neighborhood expansion” problem (Chiang et al. 2019). GraphSAINT (Zeng et al. 2019) adopts a graph sampling approach, and specifically, it uses sampling algorithms for variance reduction with an additional normalization technique for unbiasedness.

The representation learning method for large hypergraphs remains largely underexplored. HyperNetVec, a scalable unsupervised hypergraph embedding method proposed by Maleki et al. (2022) leverages multi-level (hierarchical) embedding approaches that adopt existing graph embedding methods (Grover and Leskovec 2016; Hamilton et al. 2017) on a coarsened hypergraph. While HyperNetVec can handle hypergraphs with millions of hyperedges, it is a transductive embedding method. That is, it directly learns node embeddings and cannot be directly utilized for training HNNs. To the best of our knowledge, our study is the first to propose a scalable training approach for HNNs applicable to hypergraphs with millions of nodes and hyperedges. It should be noticed that our approach is scalable enough to be used with hypergraphs with tens of millions of nodes and hyperedges.

2.2.3 Contrastive learning (CL)

Due to its effectiveness and generality, contrastive learning (CL) has emerged as a novel solution for alleviating the label-scarcity issue in representation learning in various domains, including computer vision (Chen et al. 2020; He et al. 2020), natural language processing (Gao et al. 2021), graph learning (Veličković et al. 2018; You et al. 2020; Zhu et al. 2020, 2021; Hassani and Khasahmadi 2020; Luo et al. 2022), and hypergraph learning (Zhang et al. 2021; Xia et al. 2021; Yu et al. 2021; Lee and Shin 2023). The basic concept of CL is to (a) create two augmented views from the input data and (b) learn an encoder to maximize the agreement between the two views. That is, CL approaches aim to minimize (maximize) the distance between positive (negative) pairs. The learned representations can be utilized for various downstream tasks, such as node classification (Veličković et al. 2018; Zhu et al. 2020; Lee and Shin 2023) and recommendation (Liu et al. 2021; Xie et al. 2022).

Among contrastive learning for hypergraphs (Zhang et al. 2021; Xia et al. 2021; Yu et al. 2021; Lee and Shin 2023), S^2 -HHGR (Zhang et al. 2021) uses a coarse- and fine-grained node dropout for hypergraph augmentation, and it remedies a data sparsity issue for group recommendation. DHCN (Xia et al. 2021) employs session-level contrast for session recommendation. TriCL (Lee and Shin 2023) uses a tri-directional contrastive

loss, which combines node-, group-, and membership-level contrastive losses, resulting in better performance on several downstream tasks, compared to employing simply a node-level contrastive loss.

2.2.4 (Hyper)graph partitioning

(Hyper)graph partitioning is a fundamental task where the objective is to divide nodes into multiple groups (Girvan and Newman 2002; Çatalyürek and Aykanat 2011; Tsitsulin et al. 2020) to minimize the connectivity between groups, and it has extensive applications, including anomaly detection (Ahmed et al. 2021), molecular mining (Grunig et al. 2022), and face analysis (Wang et al. 2019).

Especially for hypergraph partitioning, multi-level approaches have received intensive attention (Çatalyürek and Aykanat 2011; Karypis et al. 1997; Schlag et al. 2023). A multi-level hypergraph partitioning algorithm consists of three phases: coarsening, initial partitioning, and uncoarsening. (a) *Coarsening*: A coarsened hypergraph $\mathcal{H}^{(c)}$ is formed by merging pairs of nodes in the input hypergraph \mathcal{H} . This procedure is recursively applied to the coarsened hypergraph. The final hypergraph is the coarsest one that meets predefined termination criteria. (b) *Initial partitioning*: The coarsest hypergraph is partitioned using any partitioning rules. (c) *Uncoarsening*: Partitions found in the second phase are successively projected back towards the original hypergraph \mathcal{H} . In this paper, we adopt PaToH (Çatalyürek and Aykanat 2011) for hypergraph partitioning due to its high-scalability. However, other approaches (Karypis et al. 1997; Schlag et al. 2023; Caldwell et al. 2000; Mayer et al. 2018) can be employed instead.

While the above hypergraph partitioning (clustering) studies make assumptions about static input hypergraphs and disjoint partitions, there have been efforts to relax these assumptions in ordinary graphs. Refer to (Xie et al. 2013; Rossetti and Cazabet 2018) for surveys on this topic. Notably, such extensions include overlapping clustering, which allows entities to belong to multiple clusters (Yang and Leskovec 2013; Shchur and Günnemann 2019), clustering in time-varying graphs (Ruan et al. 2021), and overlapping clustering in time-varying graphs (Mutakin et al. 2022). Recently, overlapping clustering is addressed using graph neural networks (Shchur and Günnemann 2019) and also considered in hypergraphs (Concisciani et al. 2022).

3 Proposed hypergraph learning tasks

In this section, we propose two new hypergraph learning tasks: *hyperedge disambiguation* and *local clustering*. For each task, we first provide its mathematical formulation and then discuss its application to real-world problems. Lastly, we describe how our proposed tasks differ from commonly considered benchmark tasks, with a focus on advantages in addressing potential issues in practice.

3.1 Hyperedge disambiguation

3.1.1 Mathematical description

Consider a hypergraph $\mathcal{H} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$ and a subset $S \subseteq \mathcal{E}$ of hyperedges. Every hyperedge $e_i \in S$ is split (disjointly or with partial overlap) into two hyperedges e_{i1} and e_{i2} so that $e_{i1} \cup e_{i2} = e_i$ holds. The hyperedges in S are replaced by these split hyperedges, which results in a new hypergraph $\mathcal{H}' = (\mathcal{V}, \mathcal{E}')$ where

$$\mathcal{E}' = (\mathcal{E} \setminus S) \cup \bigcup_{e_i \in S} \{e_{i1}, e_{i2}\}. \tag{3}$$

Based on this setting, *hyperedge disambiguation* is defined as to predict whether a given pair of hyperedges in the given hypergraph $\mathcal{H}' = (\mathcal{V}, \mathcal{E}', \mathbf{X})$ are split hyperedges or not. Formally,

$$Y(\{e_i, e_j\}) = \begin{cases} 1 & e_i \cup e_j \in S \\ 0 & \text{otherwise} \end{cases}, \tag{4}$$

where $e_i, e_j \in \mathcal{E}'$. The goal of this task is to learn a function $f : \binom{\mathcal{E}'}{2} \mapsto \mathbb{R} \in [0, 1]$ to approximate $Y(\cdot)$ in Eq. (4). In this paper, we consider this task in a semi-supervised setting where a small amount of the ground-truth split pairs of hyperedges are given. However, this problem can also be considered in other settings.¹

3.1.2 Real-world applications

The hyperedge disambiguation task can be applied to many real-world applications. Below, we introduce two examples: *researcher disambiguation* and *user identification*.

Researcher disambiguation is a task to identify whether a given pair of researchers are in fact the same researcher or not. Due to namesakes and other reasons, one researcher can be represented as multiple individuals in a system, which complicates searching for experts, surveying related papers, and recommending scholarly papers (Milojević 2013; Caron and van Eck 2014; Sanyal et al. 2021). This problem can naturally be formulated as a hyperedge disambiguation task on a publication-author hypergraph where nodes denote publications and each hyperedge represents a (potentially partial) set of publications authored by one researcher. Then, predicting whether a given pair of hyperedges are split ones or not corresponds to predicting whether the two sets of publications are authored in fact by the same researcher or not, as visually depicted in Fig. 1a.

¹ For example, in an unsupervised setting without any given positive example pairs, one can additionally split edges in \mathcal{E}' to use them as positive example pairs for training.




Fig. 1 Example applications of the proposed tasks. **a** In a researcher disambiguation task, we aim to predict whether a given pair of researchers (modeled as hyperedges) are identical or not. **b** In a household matching task, we aim to predict whether a given pair of devices (modeled as nodes) are owned by people in the same household or not

The second task is *user identification*, which is to identify whether a given pair of electronic devices (or accounts for a web service) are owned by the same user (Deng et al. 2019; Malatras et al. 2017). If such a pair can be identified, the data (e.g., behavior log) from both can be used together to improve the quality of services (e.g., personalization and recommendation). To handle this problem, we can build a hypergraph where nodes denote devices (or accounts) and each hyperedge represents the set of devices (or accounts) that each device (or account) communicates with. Then, predicting whether a given pair of hyperedges are split ones or not corresponds to predicting whether the two devices (or accounts) are owned by the same user or not.

3.2 Local clustering

3.2.1 Mathematical description

Consider a hypergraph $\mathcal{H} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$ and (disjoint or partially overlapping) clusters of nodes $\mathcal{C} = \{C_1, C_2, \dots, C_{|C|}\}$ where each cluster is a subset of nodes (i.e., $C_k \subseteq \mathcal{V}, \forall C_k \in \mathcal{C}$). Based on this setting, *local clustering* is defined as to predict whether a given pair of nodes in a given hypergraph \mathcal{H} belongs together to the same cluster or not. Formally,

$$Y(\{v_i, v_j\}) = \begin{cases} 1 & \exists C_k \in \mathcal{C} \text{ such that } \{v_i, v_j\} \subseteq C_k \\ 0 & \text{otherwise} \end{cases}, \tag{5}$$

where $v_i, v_j \in \mathcal{V}$. The goal of this task is to find a function $f : \binom{\mathcal{V}}{2} \mapsto \mathbb{R} \in [0, 1]$ to approximate $Y(\cdot)$ in Eq. (5). In this paper, we consider this task in a semi-supervised setting where the members of a few clusters are given, while any information about the other clusters (e.g., the number of them) is not provided. However, this problem can also be considered in other settings.²

² For example, it can be considered in unsupervised settings especially when the clustering membership is strongly correlated with node features and/or topological information.

3.2.2 Real-world applications

The local clustering task can be used to formulate a wide range of real-world problems. Here, we provide two examples: *sub-field detection* and *household matching*.

The first task is *sub-field detection*, which is to classify whether a given pair of publications belong to the same sub-field or not. Note that the boundary of a sub-field is often unclear, and thus a single publication may belong to multiple sub-fields. This problem is formulated as a local clustering task on hypergraphs where a node denotes a publication, a hyperedge represents the set of publications that are co-cited by one publication, and clusters indicate sub-fields.

The second task is *household matching*, a task to predict whether a given pair of electronic devices are owned by people in the same household or not. By grouping devices from the same households, recommendations and advertisements can be better personalized. For example, consider a hypergraph where nodes denote devices and each hyperedge indicates the set of devices access to the same IP within a specific time interval (e.g., set of devices access to IP 123.456.789 between 9 and 10 AM on May 16, 2023). If each cluster is formed by the devices from the same household, inferring whether two nodes belong to the same cluster or not corresponds to inferring whether two devices are from the same household or not, as visually depicted in Fig. 1b. Note that household matching is not formulated accurately as an entity classification task, since (1) It is hard to know in advance the number of households, which is required and fixed in entity classification tasks, and (2) The number of households changes over time.

3.3 Differences from existing tasks

In this subsection, we discuss some difficulties when formulating the aforementioned real-world applications as commonly-used single-entity-level tasks (e.g. node classification). Note that our proposed pair-level tasks mitigate such difficulties.

3.3.1 Entity classification

One may suspect that the aforementioned applications can also be formulated as single-entity-level tasks. However, formulating these applications as our tasks offer several advantages over formulating them as an entity classification task: (1) Our tasks do not require the number of labels, which is equivalent to the number of split hyperedges and the number of clusters in our tasks, in advance, and (2) A model for our tasks does not need to be retrained whenever the number of labels changes in time-evolving data.

For example, one may consider formulating the household matching problem as a node classification problem where each class indicates a household. Here, it is important to note that the number of such classes is typically impossible to know in

advance, and it also changes ceaselessly over time. Moreover, no example can be provided for some classes during training. However, in node-classification formulation, the output dimension of classifiers should be predetermined by the number of classes, and they cannot output classes that are not observed during training.

On the other hand, our formulation as the local clustering task is free from such unrealistic requirements. That is, we do not require the number of households in advance and require only some pairs of devices from some households. Moreover, machine learning models trained for the local clustering task can naturally be generalized to households unseen during training.

Formulating the researcher disambiguation problem as a hyperedge classification task also encounters the same issues above. The number of real identities, which change over time, is impossible to know in advance. However, considering the researcher disambiguation problem as a hyperedge disambiguation task is free from the above-mentioned issues.

3.3.2 Overlapping clustering

The proposed local clustering problem is also closely related to the overlapping clustering problem, whose goal is to identify clusters where an object (e.g., a node) can belong to multiple clusters (see Sect. 2.2.4). The problem has been considered mostly in ordinary graphs and more importantly in transductive settings. That is, clusters are identified based on the current state of a graph, and once the graph changes, one needs to re-run the entire clustering process from scratch. Thus, the task is not proper to be used for evaluating HNNs (especially their generalization capabilities), which are particularly useful in inductive settings. Moreover, in many cases, the number of ground truth clusters is required, while it is unavailable in many realistic settings.

4 Proposed large-scale hypergraph datasets

In this section, we introduce two large-scale hypergraph-structured datasets. We build them by processing raw data from AMINER (Tang et al. 2008) and Microsoft Academic Graph (MAG) (Sinha et al. 2015). First, we describe how the node features and labels are obtained from the raw data, and then we present how the hypergraph topologies are constructed. Figure 2 provides a visual description of the hypergraph construction process.

4.1 Extracting basic information

The source of the datasets is Open Academic Graph³ (Zhang et al. 2019), and it provides two large academic bibliographic datasets: AMINER and MAG. These bibliographic

³ Available at <https://www.aminer.cn/oag-2-1>.

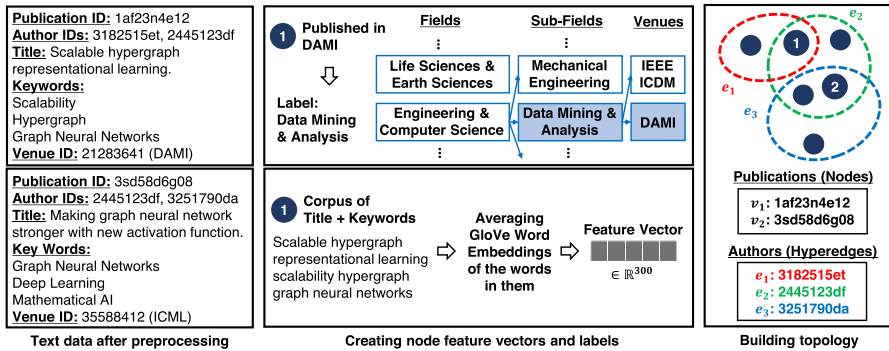


Fig. 2 Examples for the construction process of the AMINER and MAG datasets. **(Left)** From raw text data, we extract the **a** publication ID, **b** author IDs, **c** title, **d** keywords, and **e** venue ID of each publication. **(Middle)** We obtain the 300-dimensional features for each publication by averaging the embeddings of all words in its title and keywords. Each publication is labelled based on the sub-field it belongs to. **(Right)** We create nodes corresponding to publications and hyperedges corresponding to authors

datasets consist of huge raw text data containing information about over 100 million authors and publications. From this raw data, we first extract the (a) publication ID, (b) author IDs, (c) the title, (d) keywords, and (e) the venue ID of each publication.

4.2 Building node features

We construct a 300-dimensional feature vector for each publication, which corresponds to a node. To this end, we average the embeddings of all the words appearing in its title and keywords, where the word embeddings are obtained by the pre-trained GloVe (Pennington et al. 2014) model⁴. The node features are created separately for AMINER and MAG.

4.3 Building node labels

We classify publications based on the academic fields they belong to. Since the original raw data does not contain any information about the publications’ academic fields, we utilize the venue name (e.g., conference and journal) of each publication to infer its field. Specifically, we use the category hierarchy provided by Google Scholar⁵ for mapping venue names to academic fields. The hierarchy consists of three levels: fields, sub-fields, and venues. For example, the journal *Data Mining and Knowledge Discovery* (DAMI) belongs to the sub-field *Data Mining & Analysis*, which in turn belongs to the field *Engineering & Computer Science*. We label each publication

⁴ Available at <https://github.com/UKPLab/sentence-transformers>.

⁵ Available at https://scholar.google.com/citations?view_op=top_venues&hl=en&vq=eng. In this taxonomy, a single venue is associated with a single sub-field.

based on the sub-field (i.e., the middle level) it belongs to. Note that, in the category that we use, each venue (node) is matched with exactly one sub-field (label).

4.4 Building hypergraph structures

The raw data are far from complete with many missing values (e.g., missing authors) and meaningless information (e.g., incomprehensible words). For the quality of the created hypergraphs, we first filter out all publications that satisfy at least one of the following three conditions:

- *It does not have author information (i.e., names or IDs),*
- *Its title and keywords consist of less than three words,*
- *It is not matched with any academic field.*

Then, we build hypergraphs using the remaining publications.

As mentioned in Sect. 4.2, publications and authors are represented by nodes and hyperedges, respectively. To be specific, for each author ID, the publications that contain the ID as an author form a hyperedge. In this way, we create two co-authorship hypergraphs, AMiner and MAG, which contain over 20 million hyperedges and 30 million hyperedges, respectively. Numerically, the MAG dataset contains **158×** more nodes and **129×** more hyperedges than the Trivago dataset, one of the commonly used hypergraph datasets. Table 2 compares the statistics of the proposed datasets and some existing datasets. Note that, the AMiner and MAG datasets can be used not only for the proposed pair-level tasks but also for common downstream tasks, such as node classification, hyperedge prediction, and clustering.

5 Proposed scalable hypergraph learning method

In this section, we introduce PCL (**P**artitioning-based **C**ontrastive **L**earning), a scalable hypergraph learning algorithm for pair-level tasks. We first provide an overview of PCL together with the rationale of its components. Next, we describe the details of PCL. At last, we give the complexity analysis of PCL. A pictorial description of PCL is provided in Fig. 3.

5.1 Challenges and main ideas


We analyze two challenges encountered in training representation learning models for pair-level downstream tasks on large-scale hypergraphs. Then, we present our proposed solutions for overcoming these problems.

Table 2 Statistics of five real-world hypergraph datasets. MAG has 158× more nodes and 129× more hyperedges than Trivago, a commonly used hypergraph dataset

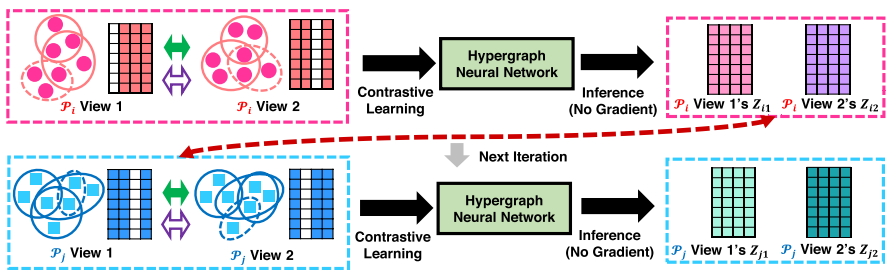
Source	Dataset	$ V $	$ E $	# features	# classes
Existing	DBLP	41,302	22,363	1,425	6
	Trivago*	172,738	233,202	300	160
Transformed**	OGBN-MAG	736,389	1,134,649	128	349
Proposal	AMINER	13,262,573	22,552,647	300	257
	MAG	27,320,375	30,175,013	300	257

* In the Trivago dataset, we use the first 300 left singular vectors of the incidence matrix as the node features.

** Transformed from a heterogeneous graph



(a) Notations & Symbols



(b) Proposed methods: Basic PCL and PCL+PINS

Fig. 3 Overview of Basic PCL and PCL+PINS, the proposed large-scale hypergraph learning methods. We consider each partition (\mathcal{P}_i and \mathcal{P}_j in the example) as a mini-batch and perform contrastive learning on each mini-batch. In Basic PCL, after a single contrastive learning iteration is done on a partition (\mathcal{P}_i), it moves to another partition (\mathcal{P}_j) and starts the next iteration. In PCL+PINS, after a model is updated on a partition (\mathcal{P}_i), PCL+PINS extracts the node embeddings of \mathcal{P}_i 's two views (\mathbf{Z}_{i1} and \mathbf{Z}_{i2} in our example) and uses \mathbf{Z}_{i1} and \mathbf{Z}_{i2} to compute inter-partition CL loss in the next partition \mathcal{P}_j

Challenge 1. Considerable Time and Space Cost: How can we train hypergraph neural network models on large-scale hypergraph datasets? The commonly-used full-batch training is typically not possible, since such hypergraphs cannot be entirely loaded into GPU memory, and this causes the out-of-memory problem. To mitigate this issue, a mini-batch training method (Hamilton et al. 2017) loads only the subgraph from which each node aggregates messages into GPU memory at a time. Although it reduces memory requirements, it introduces significant computational

overhead to extract such subhypergraphs for each node and load them repeatedly into GPU memory. Even worse, such subhypergraphs are often large since real-world hypergraphs tend to have a small diameter (Do et al. 2020; Ko et al. 2022).

Solution 1: Our solution to the scalability issue in large-scale hypergraph representation learning is *partitioning* (or equivalently node clustering). The process involves dividing the input hypergraph into smaller partitions, treating each partition as a distinct mini-batch. Partitioning-based approaches have proved efficient in training models (i.e., graph neural networks) for representation learning in ordinary graphs, mitigating the neighborhood expansion problem (Chiang et al. 2019). Despite their efficiency, partitioning-based approaches have not been considered for training HNNs on large-scale hypergraphs. Motivated by this efficiency, we split the entire hypergraph into K partitions. Here, we use a well-known hypergraph partitioning algorithm PaToH⁶ (Çatalyürek and Aykanat 2011), while any partitioning method can be used instead. We load a single partition into GPU memory at a time, which requires K times more processes of loading and unloading partitions into GPU memory. Alternatively, one can put two different partitions into GPU memory, but it requires at most $K(K - 1)/2$ times more processes, which are computationally too costly for large K .


Challenge 2. Information Loss Caused by Partitioning: Although the partitioning approach can mitigate the scalability issues, it has several limitations in handling pair-level downstream tasks presented in Sect. 3.

- Some hyperedges do not exist in any partitioned hypergraph if their constituent nodes belong to different partitions, and thus such hyperedges cannot be utilized in the training phase (see e_1 of Fig. 4a).
- Certain node pairs and hyperedge pairs that are given as labeled pairs may be split into different partitions (see $\{v_1, v_5\}$ of Fig. 4a). In such scenarios, these divided pairs cannot be used as labeled pairs for the partitioning-based approach to train HNNs.
- A single partition may not represent the entire hypergraph adequately since partitioning algorithms tend to group similar nodes together in the same partition (Chiang et al. 2019; Guo et al. 2021). In other words, each partition may exhibit bias.

Solution 2: Our countermeasure for the second challenge is three-fold.

- The first one is to leverage a *contrastive learning* (CL) framework (Chen et al. 2020; Lee and Shin 2023) to train hypergraph encoders (i.e., HNNs) on partitioned hypergraphs. The basic concept of CL is to train a hypergraph encoder by

⁶ PaToH is a balanced partitioning method. It ensures that all generated partitions are of similar sizes (Çatalyürek and Aykanat 2011), specifically satisfying $|\mathcal{P}_k^v| \leq \frac{(1+\epsilon)}{|\mathcal{P}|} \sum_{i=1}^{|\mathcal{P}|} |\mathcal{P}_i^v|$, $\forall k = 1, \dots, |\mathcal{P}|$. As shown in Table 8 in Sect. 6.3.5, partitions obtained by PaToH from real-world hypergraphs are well balanced. Specifically, the standard deviation of the number of nodes in each partition is less than 0.5% of the average number of nodes per partition.



(a) Hypergraph and resulting partitions: depicting difficulties of representation learning on partitions.

(b) An example of P-IOS.

Fig. 4 (a) Assume the two nodes v_1 and v_5 belong to the same cluster, and the two hyperedges e_1 and e_2 are split from one original hyperedge. Note that v_1 and v_5 belong to different partitions, and $e_1 = \{v_2, v_5\}$ does not even exist in any partitioned hypergraph since $v_2 \in \mathcal{P}_1^V$ and $v_5 \in \mathcal{P}_2^V$. Thus, $\{v_1, v_5\}$ and $\{e_1, e_2\}$ cannot be used as labeled pairs when one trains a model on partitioned hypergraphs in a supervised manner. (b) In order to mitigate information loss caused by partitioning, P-IOS restores each hyperedge by adding its missing nodes back to the hyperedge

(a) creating two augmented views from the input data and (b) letting the encoder maximize the agreement between the two views. Note that CL does not require any label supervision, and this characteristic ensures that the encoder is properly trained even in the case where labeled pairs are split or lost.

- The second one is **PINS**, which mitigates the bias problem caused by hypergraph partitioning. Since it processes a single partition as a mini-batch, a hypergraph encoder has limited exposure to the differences between nodes in different partitions. To enable the hypergraph encoder to incorporate information from other partitions, PINS utilizes the node representations from the previous partition.
- The last one is **P-IOS**, restores each hyperedge by adding its missing nodes back to the hyperedge, to alleviate information loss caused by partitioning.

5.2 Details of proposed methods

In this subsection, we explain the details of the three proposed solutions discussed in the previous subsection. We first introduce **Basic PCL**, a contrastive learning (CL) framework with partitioned hypergraphs. Then, we introduce PCL+PINS, a method that uses the embeddings of the previous partitions as negative samples. In addition, we provide PCL+P-IOS, a method that augments each partition without additional hyperedges. Lastly, we describe how a classifier of downstream tasks is trained with pre-trained hypergraph encoders.

Algorithm 1 Basic PCL and PCL+PINS: Partitioned contrastive learning for hypergraph

Input: (1) Hypergraph encoder Enc_θ , projection head $Proj_\Theta$
 (2) Partitioned hypergraph $\mathcal{P} = \{\mathcal{P}_1, \dots, \mathcal{P}_K\}$
 (3) Feature augmentation prob. p_v , hyperedge augmentation prob. p_e
 (4) Number of epochs EP , Optimizer $Optim$, Number of negative samples N
 (5) PCL+PINS scalar λ (if $\lambda = 0$; then it is Basic PCL)

Output: Pretrained hypergraph encoder Enc_θ

```

1 for  $ep \leftarrow 1, \dots, EP$  do
2   Shuffle the order of partitions in  $\mathcal{P}$ 
3   foreach  $\mathcal{P}_k \in \mathcal{P}$  do
4      $\mathcal{P}'_{k1} \leftarrow Au(\mathcal{P}_k, p_v, p_e)$ ;  $\mathcal{P}'_{k2} \leftarrow Au(\mathcal{P}_k, p_v, p_e)$            ▷ Augmentation
5      $\mathbf{P}_{k1} \leftarrow Enc_\theta(\mathcal{P}'_{k1})$ ;  $\mathbf{P}_{k2} \leftarrow Enc_\theta(\mathcal{P}'_{k2})$            ▷ Node embedding
6      $\mathbf{Z}_{k1} \leftarrow Proj_\Theta(\mathbf{P}_{k1})$ ;  $\mathbf{Z}_{k2} \leftarrow Proj_\Theta(\mathbf{P}_{k2})$            ▷ Projection
7      $Loss \leftarrow \mathcal{L}_{CL,k}$  (use  $\mathbf{Z}_{k1}$ ,  $\mathbf{Z}_{k2}$ , and  $N$ )           ▷ Contrastive loss in Eq (10)
8     if  $(\lambda > 0) \wedge (ep > 1)$  then
9        $Loss \leftarrow Loss + \lambda \sum_{i=1, j=1}^{2,2} \mathcal{L}_{I,k}(i, j)$            ▷ Inter-partition loss in Eq (11)
          (use  $\mathbf{Z}'_1$  and  $\mathbf{Z}'_2$  as node embeddings of two views of the previous partition)
10    Update  $Enc_\theta$  and  $Proj_\Theta$  by  $Optim(Loss, \theta, \Theta)$            ▷ Back propagation
11    if  $\lambda > 0$  then
12       $\mathbf{Z}'_1 \leftarrow Proj_\Theta(Enc_\theta(\mathcal{P}'_{k1}))$ ;  $\mathbf{Z}'_2 \leftarrow Proj_\Theta(Enc_\theta(\mathcal{P}'_{k2}))$    ▷ Only for PINS
13  return  $Enc_\theta$ 
    
```

5.2.1 Basic PCL

Basic PCL (**P**artitioning-based **C**ontrastive **L**earning) is a contrastive learning framework that trains HNNs on partitioned hypergraphs. An example of Basic PCL is illustrated in Fig. 3, and its pseudocode is provided in Algorithm 1. We denote the set of all partitions as $\mathcal{P} = \{\mathcal{P}_1, \dots, \mathcal{P}_K\}$, where each partition is defined as $\mathcal{P}_k = (\mathcal{P}_k^V, \mathcal{P}_k^E)$, where $\mathcal{P}_k^V \subseteq \mathcal{V}$ and $\mathcal{P}_k^E = \{e \cap \mathcal{P}_k^V : e \in \mathcal{E}, e \cap \mathcal{P}_k^V \neq \emptyset\}$. Here, K denotes the number of partitions⁷. Note that \mathcal{P}_k can also be expressed as $(\mathcal{P}_k^X, \mathcal{P}_k^H)$, where \mathcal{P}_k^X denotes the node features of \mathcal{P}_k^V , and \mathcal{P}_k^H denotes the incidence matrix of \mathcal{P}_k .

In a nutshell, Basic PCL receives a set of partitions \mathcal{P} as an input and treats each partition $\mathcal{P}_k \in \mathcal{P}$ as a mini-batch (line 3) for contrastive learning. It trains a hypergraph encoder (i.e., hypergraph neural network) aiming to minimize a contrastive learning loss, which will be further elaborated. Details are as follows:

First, Basic PCL creates two augmented views of an input partition \mathcal{P}_k (line 4). It mainly utilizes a masking-based augmentation strategy (Lee and Shin 2023), which

⁷ One can set K based on the available amount of space (low K takes more memory consumption in general). Note that the performance of the proposed method is not significantly affected by K , which will be demonstrated in Sect. 6.3.5.

corrupts node features and membership (each group and its members). For node feature augmentation, Basic PCL employs a single random binary mask, with each entry sampled from a Bernoulli distribution $B(1 - p_v)$. This mask is utilized to set certain columns of \mathcal{P}_k^X , which is a feature matrix of the current partition, to zero vectors. Similarly, for membership augmentation, Basic PCL randomly removes nodes from the hyperedges that they belong to. Specifically, Basic PCL generates a binary mask of size $nz(\mathcal{P}_k^H)$ (i.e., the number of nonzero entries of an incidence matrix \mathcal{P}_k^H of the current partition) with each entry sampled from a Bernoulli distribution $B(1 - p_e)$. This mask is leveraged to remove certain node-hyperedge memberships. We use \mathcal{P}'_{k1} and \mathcal{P}'_{k2} to denote the two resulting views created through this process, as in Eq. (6):

$$\mathcal{P}'_{k1} = (\mathcal{P}'_{k1X}, \mathcal{P}'_{k1H}) \leftarrow Au(\mathcal{P}_k, p_v, p_e); \mathcal{P}'_{k2} = (\mathcal{P}'_{k2X}, \mathcal{P}'_{k2H}) \leftarrow Au(\mathcal{P}_k, p_v, p_e), \tag{6}$$

where Au denotes the augmentation process described above.

Then, Basic PCL obtains node representations of the two views by sequentially passing the two views through a hypergraph encoder Enc_θ (see Sect. 2.1.2 for details) and a projection head $Proj_\Theta$ (line 5-6):

$$\mathbf{Z}_{k1} = Proj_\Theta(Enc_\theta(\mathcal{P}'_{k1})); \mathbf{Z}_{k2} = Proj_\Theta(Enc_\theta(\mathcal{P}'_{k2})), \tag{7}$$

where $Proj_\Theta$ is a one-layer MLP with the ReLU (Nair and Hinton 2010) activation function. Lastly, PCL employs the InfoNCE loss (Oord et al. 2018) for parameter updates to make the representations of the corresponding nodes in the two views similar and those of different nodes dissimilar⁸:

$$\ell(\mathbf{z}_{k1,i}, \mathbf{z}_{k2,i}) = -\log \frac{\exp(s(\mathbf{z}_{k1,i}, \mathbf{z}_{k2,i})/\tau)}{\sum_{t=1}^{|\mathcal{P}_k^V|} \exp(s(\mathbf{z}_{k1,i}, \mathbf{z}_{k2,t})/\tau)}, \tag{8}$$

where $\mathbf{z}_{ka,i}$ is a i^{th} row vector of \mathbf{Z}_{ka} , τ is a temperature parameter, and s is a similarity function. As the similarity function s , we use the cosine similarity, i.e., $s(\mathbf{u}, \mathbf{v}) = (\mathbf{u}^T \mathbf{v}) / (\|\mathbf{u}\| \cdot \|\mathbf{v}\|)$. As shown in the denominator of Eq. (8), for each node, its similarity with all other nodes of \mathcal{P}_k^V is computed. For computational efficiency, instead of using the entire \mathcal{P}_k^V , we obtain a uniform sample $\mathcal{P}'_k \subset (\mathcal{P}_k^V \setminus \{v_i\})$, where $|\mathcal{P}'_k| = \min(|\mathcal{P}_k^V|, N)$ and N is a hyperparameter, and use it as negative samples. That is,

$$\ell(\mathbf{z}_{k1,i}, \mathbf{z}_{k2,i}) = -\log \frac{\exp(s(\mathbf{z}_{k1,i}, \mathbf{z}_{k2,i})/\tau)}{\exp(s(\mathbf{z}_{k1,i}, \mathbf{z}_{k2,i})/\tau) + \sum_{v_i \in \mathcal{P}'_k} \exp(s(\mathbf{z}_{k1,i}, \mathbf{z}_{k2,t})/\tau)}. \tag{9}$$

In practice, we symmetrize Eq. (8) so that the final CL loss (line 7) becomes:

⁸ Note that other self-supervised losses (e.g., (Addanki et al. 2021)) can be used alternatively.

$$\mathcal{L}_{CL,k} = \frac{1}{2|\mathcal{P}_k^v|} \sum_{i=1}^{|\mathcal{P}_k^v|} (\ell(\mathbf{z}_{k1,i}, \mathbf{z}_{k2,i}) + \ell(\mathbf{z}_{k2,i}, \mathbf{z}_{k1,i})). \tag{10}$$

By using Eq. (10), Basic PCL updates Enc_θ and $Proj_\Theta$ via gradient descent.

5.2.2 PCL+PINS

PCL+PINS (**P**revious part**I**tion's **N**egative **S**amples) encourages a hypergraph encoder Enc_θ to learn dissimilarity between nodes at different partitions (see Fig. 3). A pseudocode of PCL+PINS is provided in Algorithm 1. As Basic PCL does, PCL+PINS treats each partition as a mini-batch for contrastive learning. However, while Basic PCL only utilizes contrastive loss obtained within each partition, PCL+PINS incorporates information from other partitions.

In PCL+PINS, before moving to the next partition (after updating the encoder in the current partition), the encoder extracts node embeddings of two augmented views of the current partition \mathcal{P}_k (denoted by \mathbf{Z}_{k1} and \mathbf{Z}_{k2}) (line 12). For the next partition \mathcal{P}_t , PCL+PINS uses not only the CL loss in Eq. (10), which is obtained within a partition \mathcal{P}_t , but also an inter-partition CL loss to maximize the dissimilarity between the node embeddings from a view of \mathcal{P}_t and those from a view of \mathcal{P}_k . Specifically, the inter-partition CL loss between a view \mathcal{P}'_{t1} of \mathcal{P}_t and a view \mathcal{P}'_{k1} of \mathcal{P}_k is defined as:

$$\mathcal{L}_{I,t}(1, 1) = \frac{1}{|\mathcal{P}_t^v|} \sum_{i=1}^{|\mathcal{P}_t^v|} \log \left(\sum_{j \in \mathcal{P}'_{k,s}} \exp \left(\frac{s(\mathbf{z}_{t1,i}, \mathbf{z}_{k1,j})}{\tau} \right) \right), \tag{11}$$

where $\mathcal{P}'_{k,s} \subset \mathcal{P}_k^v$ are sampled nodes from \mathcal{P}_k for negative samples. We consider all possible pairs of a view from \mathcal{P}_t and a view from \mathcal{P}_k , and based on them we define $\mathcal{L}_{I,t}(1, 1)$, $\mathcal{L}_{I,t}(1, 2)$, $\mathcal{L}_{I,t}(2, 1)$, and $\mathcal{L}_{I,t}(2, 2)$, accordingly. Then, the final loss of PCL+PINS for each partition \mathcal{P}_t is defined as:

$$\mathcal{L} = \mathcal{L}_{CL,t} + \lambda(\mathcal{L}_{I,t}(1, 1) + \mathcal{L}_{I,t}(1, 2) + \mathcal{L}_{I,t}(2, 1) + \mathcal{L}_{I,t}(2, 2)), \tag{12}$$

where $\mathcal{L}_{CL,t}$ is the loss within \mathcal{P}_t in Eq. (10), and λ is a hyperparameter that controls the strength of the inter-partition CL loss.

In summary, PCL+PINS uses node embeddings from the previous partitions as negative samples for the current partition. It is important to note that \mathbf{Z}_{k1} and \mathbf{Z}_{k2} are obtained from a partition already residing in GPU memory, and does not require any additional data loading process. In addition, \mathbf{Z}_{k1} and \mathbf{Z}_{k2} are not outdated since they are extracted after updating the encoder.

5.2.3 PCL+P-IOS

PCL+P-IOS (**P**artitions with the **I**nclusion of **O**utside **S**et) restores each hyperedge by adding its constituent nodes missed during partitioning back to the hyperedge, to

mitigate information loss caused by partitioning. An example of P-IOIS is illustrated in Fig. 4b.

Formally, for a given set of hypergraph partitions \mathcal{P} , PCL+P-IOIS builds a new hypergraph partition set $\tilde{\mathcal{P}}$, which is defined as:

$$\tilde{\mathcal{P}} = \left\{ \tilde{\mathcal{P}}_k = \left(\tilde{\mathcal{P}}_k^{\mathcal{V}}, \tilde{\mathcal{P}}_k^{\mathcal{E}} \right) : k \in \{1, \dots, K\} \right\}, \tag{13}$$

where $\tilde{\mathcal{P}}_k^{\mathcal{E}} = \{e \in \mathcal{E} : e \cap \mathcal{P}_k \neq \emptyset\}$, and $\tilde{\mathcal{P}}_k^{\mathcal{V}} = \bigcup_{e \in \tilde{\mathcal{P}}_k^{\mathcal{E}}} e$. The training strategy of PCL+P-IOIS is the same as that of Basic PCL except that PCL+P-IOIS uses $\tilde{\mathcal{P}}$, as input hypergraph partitions, instead of \mathcal{P} .

5.2.4 Classifier for pair-level downstream tasks

Note that the proposed methods (Basic PCL, PCL+PINS, and PCL+P-IOIS) are for obtaining representations of nodes. Lastly, we present how PCL makes predictions for node-pair-level tasks and hyperedge-pair-level tasks with acquired representations. We denote node representations obtained by a hypergraph encoder (trained by one of the proposed methods) by $\mathbf{P}_v \in \mathbb{R}^{|\mathcal{V}| \times F'}$. It should be noticed that the representations are the output of the hypergraph encoder Enc_θ not the projection head $Proj_\theta$, which is used only for contrastive learning.

For a node pair-level task, PCL directly uses \mathbf{P}_v as node features. To make a prediction for a pair of nodes $\{v_i, v_j\}$, PCL uses a classifier $h_v : (\mathbb{R}^{F'} \times \mathbb{R}^{F'}) \mapsto [0, 1]$. Specifically, h_v is structured as follows:

$$h_v(\mathbf{p}_{v,i}, \mathbf{p}_{v,j}) = \sigma(W_{12}(R(W_{11}\mathbf{p}_{v,i} + b_{11}) \otimes R(W_{11}\mathbf{p}_{v,j} + b_{11})) + b_{12}), \tag{14}$$

where \otimes denotes the element-wise product; σ denotes the sigmoid function; R denotes the ReLU activation function (Nair and Hinton 2010); W_{11} and W_{12} are learnable weight matrices; and b_{11} and b_{12} are learnable bias matrices. All parameters (i.e., W_{11} , W_{12} , b_{11} , and b_{12}) are updated using a classification loss. Note that in the proposed local clustering task, this prediction is used to identify whether the pair of nodes $\{v_i, v_j\}$ belong to the same cluster or not.

For a hyperedge pair-level task, PCL first creates hyperedge features by aggregating node representations that belong to the corresponding hyperedge. Formally, hyperedge features $\mathbf{P}_e \in \mathbb{R}^{|\mathcal{E}| \times F'}$ (or $\mathbf{P}_e \in \mathbb{R}^{|\mathcal{E}'| \times F'}$ in our hyperedge disambiguation task) are computed as follows:

$$\mathbf{p}_{e,i} = \sum_{v_k \in e_i} p_{v,k}, \quad \forall e_i \in \mathcal{E} \text{ (or } \forall e_i \in \mathcal{E}' \text{ in hyperedge disambiguation)}. \tag{15}$$

Although we have adopted a summation for the aggregation function, other permutation invariant functions (e.g., average and maximum) can also be utilized. To make a prediction for a pair of hyperedges $\{e_i, e_j\}$, PCL uses a classifier $h_e : (\mathbb{R}^{F'} \times \mathbb{R}^{F'}) \mapsto [0, 1]$. Specifically, h_e is structured as follows:

$$h_e(\mathbf{p}_{e,i}, \mathbf{p}_{e,j}) = \sigma(W_{22}(R(W_{21}\mathbf{p}_{e,i} + b_{21}) \otimes R(W_{21}\mathbf{p}_{e,j} + b_{21})) + b_{22}), \tag{16}$$

where \otimes denotes an element-wise product; σ denotes the sigmoid function; R denotes the ReLU activation function (Nair and Hinton 2010), W_{21} and W_{22} are learnable weight matrices; and b_{21} and b_{22} are learnable bias matrices. All parameters (W_{21} , W_{22} , b_{21} , and b_{22}) are updated using a classification loss. Note that in the proposed hyperedge disambiguation task, this prediction is used to identify whether a given pair of hyperedges $\{e_i, e_j\}$ are in fact split from one hyperedge.

When training the classifiers h_v and h_e , we freeze the parameters of the hypergraph encoder Enc_θ , and thus the node representations \mathbf{P}_v and hyperedge representations \mathbf{P}_e are also frozen. However, it is also possible to fine-tune Enc_θ while training the h_v and h_e for downstream tasks.

5.3 Complexity analysis

We examine the time complexity and memory requirement of PCL, focusing on a forward pass of Basic PCL. Note that obtaining node embeddings with a hypergraph neural network incurs time complexity of $\mathcal{O}(|\mathcal{V}| + \sum_{e_i \in \mathcal{E}} |e_i|)$, as outlined in Eqs. (1) and (2).

We first present the time complexity of a forward pass of the HCL (hypergraph contrastive learning), a contrastive learning framework based on the entire hypergraph dataset. As stated above, the time complexity of obtaining node embeddings of two views using a hypergraph encoder is $\mathcal{O}(|\mathcal{V}| + \sum_{e_i \in \mathcal{E}} |e_i|)$. After this, node embeddings of two views are fed into a projection head, which takes $\mathcal{O}(|\mathcal{V}|)$ time. At last, it computes the contrastive loss between each node and negative samples, whose time complexity is $\mathcal{O}(|\mathcal{V}|N)$ where N is the number of negative samples. The overall time complexity of a forward pass of HCL becomes

$$\mathcal{O}\left(|\mathcal{V}| + \sum_{e_i \in \mathcal{E}} |e_i| + |\mathcal{V}| + |\mathcal{V}|N\right) = \mathcal{O}\left(|\mathcal{V}| + \sum_{e_i \in \mathcal{E}} |e_i|\right). \quad (17)$$

A forward pass of HCL requires $\mathcal{O}(|\mathcal{V}| + \sum_{e_i \in \mathcal{E}} |e_i|)$ (GPU) memory space for storing the entire input hypergraph and (intermediate) embeddings.

Now, we elaborate on the complexity of our proposed method, Basic PCL. The time complexity in Eq. (17) is applied to each partition (instead of the entire hypergraph), and as a result, the overall time complexity of a forward pass of Basic PCL becomes

$$\mathcal{O}\left(\sum_{k=1}^{|\mathcal{K}|} |\mathcal{P}_k^\mathcal{V}| + \sum_{k=1}^{|\mathcal{K}|} \sum_{e'_i \in \mathcal{P}_k^\mathcal{E}} |e'_i|\right) = \mathcal{O}\left(|\mathcal{V}| + \sum_{e_i \in \mathcal{E}} |e_i|\right), \quad (18)$$

which is equivalent to the time complexity of HCL.

The (GPU) memory requirement of Basic PCL differs from that of HCL. Since Basic PCL loads only one partition into (GPU) memory at a time, the amount of (GPU) memory required becomes

$$\mathcal{O}\left(\max_{\mathcal{P}_k \in \mathcal{P}} \left(|\mathcal{P}_k^\nu| + \sum_{e'_i \in \mathcal{P}_k^\varepsilon} |e'_i|\right)\right), \tag{19}$$

which indicates that the (GPU) memory requirement of Basic PCL is less than or equal to that of the HCL. This fact is also further supported by experiments, where HCL encounters out-of-memory issues when dealing with large-scale datasets, whereas PCL is not susceptible to such problems (see Tables 3 and 4). Regarding speed, empirical observations indicate that HCL is faster in certain real-world datasets. This is because HCL requires less time to load datasets into (GPU) memory compared to PCL, which necessitates additional time for loading and unloading each partition into (GPU) memory.

6 Experiments

In this section, we review our experiments to answer the following questions:

- *RQ1*: How accurate is PCL for the proposed pair-level tasks?
- *RQ2*: How effective is PINS?
- *RQ3*: How much GPU memory is required by PCL with and without PINS?
- *RQ4*: How effective is P-IOS?
- *RQ5*: How does the performance of PCL (+ PINS) depend on the number (or size) of partitions?

We first describe the problem settings for the proposed pair-level prediction tasks. Then, we provide details of experimental settings and datasets. Lastly, we provide answers to the above questions based on our experimental results.

6.1 Problem settings

In this subsection, we introduce how we compose train and test pairs for the two pair-level prediction tasks described in Sect. 3.

6.1.1 Hyperedge disambiguation (Task-I)

Given an original hypergraph $\mathcal{H} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$, we split hyperedges in a disjoint way. Let $\mathcal{E}_S \subseteq \mathcal{E}$ be the subset of the original hyperedge set that consists of all hyperedges of sizes greater than 10 (i.e., $\mathcal{E}_S = \{e \in \mathcal{E} : |e| > 10\}$). We randomly split every hyperedge e_i in \mathcal{E}_S into two sub-hyperedges $\{e_{i1}, e_{i2}\}$, under a constraint that the size of both sub-hyperedges should be at least 5 (i.e., $e_{i1} \cap e_{i2} = \emptyset, e_{i1} \cup e_{i2} = e_i$, and $\min(|e_{i1}|, |e_{i2}|) \geq 5$). Then, we create a new hypergraph \mathcal{H}' whose hyperedge set \mathcal{E}' is defined as:

Table 3 Task-1: On the hyperedge disambiguation task, Basic PCL (proposed) is comparable to and often better than several baseline methods, including those requiring the full data to be loaded in memory for training. We report the average and standard deviation of each metric on 10 random data splits. The best performance in each setting is highlighted in **bold**, and the second best performance is highlighted in underline. OOM indicates “out of memory”

Data type	Methods	DBLP		Trivago		OGBN-MAG		AMiner		MAG		Avg Rank
		AP	AUROC	AP	AUROC	AP	AUROC	AP	AUROC	AP	AUROC	
Only X Full graph	MLP	65.2 ± 3.4	62.2 ± 3.4	60.2 ± 1.6	60.7 ± 2.5	73.1 ± 2.6	71.9 ± 2.8	91.7 ± 0.3	92.9 ± 0.3	91.2 ± 0.6	92.6 ± 0.6	9.8
	GCN	83.6 ± 1.2	84.6 ± 1.5	61.6 ± 0.8	58.1 ± 2.4	80.1 ± 1.7	79.8 ± 2.0	OOM	OOM	OOM	OOM	9.7
	GAT	83.5 ± 1.0	84.6 ± 1.0	61.6 ± 0.5	60.0 ± 2.7	76.5 ± 2.3	75.7 ± 3.0	OOM	OOM	OOM	OOM	10.3
Full hypergraph	BGRL	84.5 ± 1.2	84.3 ± 1.7	72.0 ± 1.8	72.2 ± 2.2	83.1 ± 0.6	83.9 ± 0.5	OOM	OOM	OOM	OOM	8
	GGD	80.9 ± 1.6	80.6 ± 2.1	71.8 ± 0.9	72.5 ± 1.6	76.1 ± 1.1	76.1 ± 1.3	OOM	OOM	OOM	OOM	9.7
	HGNN	77.3 ± 2.5	80.0 ± 3.1	72.1 ± 5.1	76.6 ± 1.4	91.2 ± 0.7	92.9 ± 0.6	OOM	OOM	OOM	OOM	7.7
Partitioned graph	UniGCNII	78.1 ± 3.4	77.1 ± 3.9	71.8 ± 1.5	75.6 ± 2.3	86.9 ± 5.6	88.3 ± 6.7	OOM	OOM	OOM	OOM	8.7
	ALLSET	53.1 ± 1.2	55.8 ± 2.2	53.7 ± 1.2	56.8 ± 2.2	65.3 ± 1.4	73.4 ± 1.6	OOM	OOM	OOM	OOM	13.1
	HCL	91.0 ± 1.7	92.6 ± 2.1	73.7 ± 0.7	78.8 ± 1.1	94.2 ± 0.4	95.5 ± 0.3	OOM	OOM	OOM	OOM	4.5
Partitioned hyper-graph	GCN	84.5 ± 1.6	85.9 ± 1.6	65.9 ± 1.9	66.7 ± 2.0	82.3 ± 2.6	83.4 ± 2.8	81.2 ± 0.6	81.5 ± 0.6	OOM	OOM	7.9
	GAT	84.7 ± 1.3	86.4 ± 1.2	62.2 ± 3.1	62.5 ± 4.6	78.0 ± 1.8	79.6 ± 1.8	81.3 ± 0.5	81.0 ± 0.6	OOM	OOM	8.2
	BGRL	85.8 ± 1.7	85.6 ± 1.7	<u>82.4 ± 2.2</u>	<u>83.5 ± 2.1</u>	91.4 ± 0.5	92.8 ± 0.5	90.1 ± 0.8	90.9 ± 0.8	89.7 ± 0.7	90.8 ± 0.9	<u>3.9</u>
Partitioned hyper-graph	GGD	80.1 ± 2.4	80.8 ± 2.3	76.9 ± 1.6	78.3 ± 1.5	88.7 ± 1.0	90.2 ± 0.9	82.5 ± 1.2	83.6 ± 1.5	80.1 ± 1.3	81.2 ± 1.4	6.2
	HGNN	84.1 ± 2.0	86.2 ± 1.8	71.8 ± 0.8	75.8 ± 1.0	85.9 ± 1.1	88.1 ± 1.0	91.3 ± 0.3	92.5 ± 0.3	89.2 ± 0.7	91.6 ± 0.5	5.4
	UniGCNII	79.9 ± 1.8	80.0 ± 1.8	71.5 ± 2.6	74.2 ± 3.5	77.4 ± 0.9	75.7 ± 1.2	<u>91.8 ± 0.5</u>	92.1 ± 0.5	90.4 ± 0.4	91.9 ± 0.1	7.9
Partitioned hyper-graph	ALLSET	54.2 ± 1.6	57.6 ± 2.7	53.6 ± 1.2	56.7 ± 2.0	59.3 ± 1.6	65.6 ± 2.4	61.5 ± 1.8	68.6 ± 2.5	OOM	OOM	13.2
	Basic PCL (proposed)	<u>87.1 ± 1.9</u>	<u>88.3 ± 2.0</u>	88.2 ± 0.9	88.9 ± 0.7	<u>94.1 ± 0.4</u>	<u>95.1 ± 0.3</u>	95.5 ± 0.7	96.0 ± 0.8	96.2 ± 0.3	96.9 ± 0.3	1.4

Table 4 Task-II: On the local clustering task, PCL+PINS (proposed) is comparable to and often better than several baseline methods, including those requiring the full data to be loaded in memory for training. We report the average and standard deviation of each metric on 10 random data splits. The best performance in each setting is highlighted in **bold**, and the second best performance is highlighted in underline. OOM indicates “out of memory”

Data type	Methods	DBLP		Trivago		OGBN-MAG		AMiner		MAG		Avg. Rank
		AP	AUROC	AP	AUROC	AP	AUROC	AP	AUROC	AP	AUROC	
Only X Full graph	MLP	52.2 ± 2.3	51.3 ± 3.5	50.6 ± 0.4	50.9 ± 0.7	72.1 ± 1.0	73.9 ± 0.8	70.5 ± 0.9	72.9 ± 0.8	76.2 ± 1.3	79.4 ± 0.9	7.1
	GCN	54.6 ± 4.0	54.1 ± 4.7	50.1 ± 0.2	50.2 ± 0.1	53.5 ± 0.3	53.8 ± 0.4	OOM	OOM	OOM	OOM	11.1
	GAT	55.1 ± 4.2	55.3 ± 5.9	50.1 ± 0.0	50.0 ± 0.0	OOM	OOM	OOM	OOM	OOM	OOM	11.6
Full hypergraph	BGRL	55.1 ± 5.8	52.9 ± 3.1	50.3 ± 0.2	50.4 ± 0.2	53.2 ± 0.3	53.3 ± 0.3	OOM	OOM	OOM	OOM	10.9
	GGD	58.6 ± 9.5	56.0 ± 8.6	50.3 ± 0.2	50.3 ± 0.2	53.0 ± 0.3	53.1 ± 0.3	OOM	OOM	OOM	OOM	10.1
	HGNN	56.9 ± 5.3	54.9 ± 6.3	54.8 ± 3.7	54.8 ± 3.6	78.2 ± 1.0	80.1 ± 0.7	OOM	OOM	OOM	OOM	6.3
Partitioned graph	UniGCNII	55.2 ± 4.6	52.9 ± 4.5	51.6 ± 1.0	51.1 ± 0.9	76.7 ± 1.2	79.0 ± 1.1	OOM	OOM	OOM	OOM	7.8
	ALLSET	54.2 ± 4.1	56.1 ± 5.4	51.0 ± 0.7	51.7 ± 0.7	60.0 ± 2.2	61.8 ± 2.5	OOM	OOM	OOM	OOM	8.6
	HCL	63.6 ± 6.9	61.4 ± 7.9	58.6 ± 2.6	58.8 ± 4.2	78.5 ± 1.0	80.9 ± 0.7	OOM	OOM	OOM	OOM	4.8
Partitioned Hypergraph	GCN	51.9 ± 0.4	52.2 ± 0.4	50.1 ± 0.0	50.1 ± 0.0	51.6 ± 0.4	51.7 ± 0.4	54.9 ± 0.7	54.3 ± 0.7	54.6 ± 0.4	54.3 ± 0.7	12.7
	GAT	52.4 ± 0.6	53.1 ± 0.6	50.4 ± 0.7	50.5 ± 1.0	54.1 ± 0.7	54.3 ± 0.8	55.2 ± 0.4	54.6 ± 0.4	54.6 ± 0.4	54.6 ± 0.4	10.3
	BGRL	58.8 ± 7.7	56.8 ± 5.6	61.2 ± 3.9	61.4 ± 3.9	65.6 ± 1.0	66.4 ± 0.8	65.8 ± 0.5	67.1 ± 0.6	71.0 ± 1.0	72.7 ± 0.9	3.6
Partitioned hypergraph	GGD	67.7 ± 12.7	67.2 ± 13.3	59.4 ± 3.8	59.4 ± 3.5	64.8 ± 1.0	65.5 ± 0.6	62.8 ± 0.8	64.3 ± 0.8	68.5 ± 0.9	70.5 ± 0.9	3.8
	HGNN	55.2 ± 6.5	55.2 ± 7.3	52.0 ± 1.0	52.4 ± 2.0	68.8 ± 1.3	69.9 ± 1.6	57.6 ± 1.6	57.5 ± 2.1	62.8 ± 3.0	62.5 ± 3.6	6.6
	UniGCNII	52.6 ± 1.7	52.3 ± 0.9	50.4 ± 0.3	50.1 ± 0.2	54.9 ± 0.6	54.6 ± 0.5	59.9 ± 1.2	59.6 ± 1.3	65.9 ± 2.2	66.0 ± 2.2	9.8
Partitioned hypergraph (pro-posed)	ALLSET	53.2 ± 1.6	54.6 ± 2.1	51.5 ± 0.3	52.5 ± 0.5	57.5 ± 1.9	59.2 ± 2.8	58.1 ± 1.1	60.9 ± 1.4	OOM	OOM	8.2
	PCL+PINS	63.2 ± 10.6	64.0 ± 11.5	58.6 ± 1.0	59.2 ± 1.3	77.6 ± 0.8	79.8 ± 0.6	80.2 ± 1.7	81.6 ± 0.8	83.1 ± 1.4	86.1 ± 1.5	2.1

$$\mathcal{H}' = (\mathcal{V}, \mathcal{E}'), \text{ where } \mathcal{E}' = (\mathcal{E} \setminus \mathcal{E}_S) \cup \bigcup_{e_i \in \mathcal{E}_S} \{e_{i1}, e_{i2}\}. \tag{20}$$

We denote the set of ground-truth split hyperedge pairs as $G_S = \{\{e_{i1}, e_{i2}\} : e_i \in \mathcal{E}_S\}$. We use a subset G'_S of G_S as a train positive pair set, and the others (i.e., $G_S \setminus G'_S$) as a test positive pair set.

For training and testing of classifiers of this task, we also create a negative-pair set. We constrain the size of hyperedges in negative pairs since all positive pairs consist of hyperedges of size at least 5. Specifically, we use the set $\mathcal{E}'_5 = \{e \in \mathcal{E}' : |e| \geq 5\}$ of hyperedges whose size is at least 5. We compose a negative pair set with pairs of hyperedges in \mathcal{E}'_5 , ensuring that these negative pairs are not included in G_S . Specifically, the train negative pair set is the set of pairs of hyperedges $\{e_i, e_j\}$ that satisfy at least one of the two following conditions:

- **Condition 1:** $(e_i, e_j \in \bigcup_{g_s \in G'_S} g_s) \wedge (\{e_i, e_j\} \notin G'_S),$
- **Condition 2:** $(e_i \in \bigcup_{g_s \in G'_S} g_s) \wedge (e_j \notin \bigcup_{g_s \in G'_S} g_s).$

The test negative pair set is the set of pairs $\{e_i, e_j\}$ of hyperedges from \mathcal{E}'_5 that satisfy the following condition:

$$\left(e_i, e_j \notin \bigcup_{g_s \in G'_S} g_s \right) \wedge (\{e_i, e_j\} \notin G_S). \tag{21}$$

Since the number of possible negative pairs is vast, instead of considering them all, we sample some of them so that (1) the number of train negative pairs matches that of train positive pairs, and (2) the number of test negative pairs matches that of test positive pairs.

6.1.2 Local clustering (Task-II)

In our setting, the clusters are formed in a disjoint way. From the entire set of clusters $\mathcal{C} = \{C_1, C_2, \dots, C_{|\mathcal{C}|}\}$, a subset $\mathcal{C} \simeq \subsetneq \mathcal{C}$ is given for training. The train positive pair set is defined as the set of node pairs that belong together to a **known** ground-truth cluster, formally,

$$\left\{ \{v_i, v_j\} \in \binom{\mathcal{V}}{2} : v_i, v_j \in C', C' \in \mathcal{C} \simeq \right\}. \tag{22}$$

The test positive pair set is defined as the set of node pairs that belong together to an **unknown** ground-truth cluster, formally

$$\left\{ \{v_i, v_j\} \in \binom{\mathcal{V}}{2} : v_i, v_j \in C', C' \in (\mathcal{C} \setminus \mathcal{C} \simeq) \right\}. \tag{23}$$

Instead of using all possible positive pairs within each cluster, we sample a certain portion of node pairs from each cluster.

We define the train node set as $V_{TR} = \bigcup_{C' \in \mathcal{C}_{\approx}} C'$ and the test node set as $V_{TE} = \bigcup_{C' \in (\mathcal{C} \setminus \mathcal{C}_{\approx})} C'$. Then, the train negative pair set is defined as the node pairs from V_{TR} that are not included in the train positive pair set. Similarly, we define the test negative pair set as the node pairs from V_{TE} , except for the pairs that belong to the test positive pairs. As in the previous task, instead of using all negative pairs, we sample some of them so that (1) the number of train negative pairs matches that of train positive pairs, and (2) the number of test negative pairs matches that of test positive pairs. Throughout the experiments, in each dataset, we use the labels of nodes as the ground-truth clusters of the nodes.

6.2 Experimental settings

In this section, we provide an overview of the experimental setup for our study.

6.2.1 Baseline methods

We compare the proposed methods against 16 baseline methods in the two proposed tasks (Task-I and Task-II), which include Multi-Layer Perceptron (MLP) (Rumelhart et al. 1986), graph encoders trained via supervised learning or self-supervised learning⁹, and hypergraph encoders trained via supervised learning and (full-batch) self-supervised learning. These baseline methods serve as a means to assess the effectiveness of several ideas behind the proposed methods for Task-I and Task-II: the ideas are (1) hypergraph data modeling (2) self-supervised learning strategy, and (3) partitioning.

For graph encoders trained via supervised learning, we use GCN (Kipf and Welling 2017) and GAT (Veličković et al. 2018), and for self-supervised learning methods, we use GCN trained by BGRL (Addanki et al. 2021) or GGD (Zheng et al. 2022b). For hypergraph encoders trained via supervised learning, we use HGNN (Feng et al. 2019), UniGCNII (Huang and Yang 2021), and AllSet (Chien et al. 2021). All the above methods are trained and evaluated on both entire and partitioned (hyper)graphs. Lastly, we use HGNN trained via self-supervised learning on the entire hypergraph, which we call Hypergraph Contrastive Learning (HCL), as an additional baseline method.

6.2.2 Implementations

We implement our methods and HCL using PyTorch 1.11.0 (Paszke et al. 2019) and PyTorch Geometric 2.0.4 (Fey and Lenssen 2019), and for all other baseline methods, we use their official code.

⁹ Since graph encoders require a graph topology as an input, we convert original hypergraphs into graphs by Clique Expansion, described in Appendix A.2.

6.2.3 Machine specification

All experiments are conducted on a machine with NVIDIA RTX 8000 D6 GPUs (48GB memory) and two Intel Xeon Silver 4214R Processors.

6.2.4 Datasets

We conduct experiments on five real-world hypergraph datasets: two existing small-scale datasets (DBLP (Rossi and Ahmed 2015) and Trivago (Chodrow et al. 2021)), one medium-scale hypergraph dataset transformed from a heterogeneous graph (OGBN-MAG (Wang et al. 2020)), and the two proposed large-scale hypergraph datasets (AMINER and MAG). Some statistics of these datasets are reported in Table 2. Details regarding nodes, hyperedges, and labels of these datasets are described in Appendix A.1.

6.2.5 Experimental set-ups

We use the average precision score (AP) and the area under the ROC curve score (AUROC) as two quantitative evaluation metrics since the proposed tasks are sort of binary classification tasks. For each dataset and for each task, we evaluate each model using 10 data splits. Specifically, we measure average AP and average AUROC together with their standard deviations over 10 splits.

For Task-I, we split the ground-truth hyperedge pairs into 10 and 90% and use them as the positive train pairs and positive test pairs, respectively. In addition, we use half of the positive train pairs for validation and use the remaining half for training. By following the procedures described in Sect. 6.1, we create the same number of negative train/validation/test pairs.

Similarly for Task-II, we split the ground-truth clusters (i.e., classes) into 10 and 90% and use them as known and unknown ground-truth clusters, respectively. Here also, we use half of the known clusters for validation and use the remaining half for training. Since the DBLP dataset has only six clusters, we use four of them (specifically, two for training and two for validation) as known ground-truth clusters, which results in a 66/34% split for Task-II. Based on the clusters, by following the procedures described in Sect. 6.1, we create positive and negative train/validation/test pairs. Specifically, for each cluster, we randomly sample the following number of positive pairs from all possible positive pairs that can be created from the corresponding cluster (and use the same number of negative pairs) in each dataset: 10000 in DBLP and Trivago and 50000 in OGBN-MAG, AMINER, and MAG.

For all partition-based methods, the number of partitions (i.e., $|\mathcal{P}|$) is set to 4 for DBLP, 32 for Trivago, 128 for OGBN-MAG, and 256 for AMINER and MAG, unless otherwise stated. The effect of the number of partitions is also explored in Sect. 6.3.5. For all cases, we utilize the Adam optimizer (Kingma and Ba 2015)

with a fixed weight decay scalar of 10^{-6} . The learning rate for each case is tuned as a hyperparameter. The specific hyperparameters for the proposed methods and all baseline methods are described in Appendix A.3. For PCL, unless otherwise stated, PINS is used for Task-II but not for Task-I, and P-IOS is not used for both tasks. This is because PINS does not lead to accuracy gains in most cases for Task-I, and P-IOS limits the scalability of PCL. PINS, however, is helpful for prediction accuracy, as shown in Sect. 6.3.4.

6.3 Experimental results

6.3.1 RQ1. Overall performance on downstream tasks

As shown in Tables 3 and 4, PCL achieves overall the best AP and AUROC scores among all the considered methods (i.e., best in terms of average rank) on both tasks. There are two notable observations, which we describe below.

First, in the two proposed large-scale hypergraphs AMINER and MAG, PCL shows the best performance with a significant gap from the second best model (spec., 3.7% higher AP on Task-I and 9.7% higher AP on Task-II in AMINER dataset). This result demonstrates the effectiveness of PCL in pair-level tasks on large-scale hypergraph datasets. Note that the graph representation methods operating on clique-expanded graphs show poor performance on both tasks in large-scale datasets. This result highlights the effectiveness of the hypergraph modeling of large-scale group interactions when tackling pair-level tasks.

Second, although PCL is trained on partitioned hypergraphs, surprisingly, it shows performance comparable to or even better than that of HCL, which uses entire hypergraphs, without partitioning, for contrastive learning. That is, the topological information loss due to partitioning is not severe enough to harm the overall performance of models, and sometimes it is even helpful. We suspect that partitioning may increase the *hardness* of negative samples that are used during contrastive learning, which may lead to performance improvement. Specifically, when selecting negative samples for contrastive learning, there is inherent randomness in the choice of which negative samples. If we choose negative samples from the entire hypergraph (as in HCL), it is likely that distant nodes are chosen. In this case, representations of different nodes can easily be dissimilar, since it is likely that distant nodes have different neighbors. However, if we select negative samples within a partition (as in PCL), it is relatively more likely that nodes sharing many neighbors are chosen¹⁰, and this makes the encoder hard to maximize the dissimilarity between representations of such nodes. As a result, the hardness of negative samples increases in PCL, and learning to distinguish hard negative samples from positive samples potentially enhances the quality of the trained encoder's output representation (Robinson et al. 2021; Hassani and Khasahmadi 2020).

¹⁰ This is because partitioning algorithms generally assign such nodes in the same partition.

Table 5 Effectiveness of PINS on Task-II. PCL+PINS outperforms PCL w/o PINS in most datasets

Data	Metric	PCL w/o PINS	PCL+PINS
DBLP	AP	62.4 ± 11.5	63.2 ± 10.6
	AUROC	61.1 ± 9.9	64.0 ± 11.5
Trivago	AP	58.6 ± 1.0	58.6 ± 1.0
	AUROC	58.7 ± 1.0	59.2 ± 1.3
OGBN-MAG	AP	77.3 ± 0.7	77.6 ± 0.8
	AUROC	79.6 ± 0.3	79.8 ± 0.6
AMiner	AP	78.6 ± 1.3	80.2 ± 1.7
	AUROC	81.3 ± 1.2	81.6 ± 0.8
MAG	AP	83.3 ± 1.1	83.1 ± 1.4
	AUROC	86.3 ± 0.7	86.1 ± 1.5

6.3.2 RQ2. Effectiveness of PINS

We demonstrate the effectiveness of PINS on Task-II by comparing the performances of PCL w/o PINS (Basic PCL) and PCL+PINS. As shown in Table 5, PCL+PINS outperforms PCL w/o PINS in four out of five datasets. The performance gain by PINS is up to +2.9% (AUROC on DBLP), while the performance degradation is up to -0.2% (AUROC on MAG). Thus, we conclude that PINS is effective on Task-II. However, in our preliminary study, PINS does not increase the performance on Task-I.

6.3.3 RQ3. Efficiency of PINS

Despite the effectiveness of PINS on Task-II, one may concern with additional computational and memory costs caused by PINS. Regarding the concern, we compare the average running time per epoch of PCL+PINS and PCL w/o PINS (Basic PCL), and we also measure their average GPU memory usage during the contrastive learning process¹¹.

As reported in Table 7, the average running time of PCL+PINS is 20, 32, and 23% longer than that of PCL w/o PINS on OGBN-MAG, AMiner, and MAG, respectively¹². Moreover, it uses 0.1, 0.01, and 0.01% more average GPU memory, compared to PCL w/o PINS, for the above three datasets. Therefore, the additional cost due to PINS is not substantial.

¹¹ At each mini-batch (partition) of contrastive learning, we record the GPU memory usage after completing the gradient computation (spec., execute `loss.backward()` and check the current GPU memory allocation using `torch.cuda.memory_allocated(device)`). After training an encoder in every mini-batch, we calculate the average GPU memory usage for the current epoch by averaging the usage across all partitions. Finally, we compute the average GPU memory usage across all epochs.

¹² The total contrastive training epochs are 50.

Table 6 Effectiveness of P-IOS. PCL+P-IOS outperforms PCL w/o P-IOS in all cases

Task	Data	Metric	PCL w/o P-IOS	PCL+P-IOS
Task-I	DBLP	AP	87.1 ± 1.9	90.8 ± 1.8
		AUROC	88.3 ± 2.0	92.9 ± 1.2
	Trivago	AP	88.2 ± 0.9	89.4 ± 1.2
		AUROC	88.9 ± 0.7	89.5 ± 0.9
	*OGBN-MAG	AP	94.1 ± 0.4	94.8 ± 0.4
		AUROC	95.1 ± 0.2	96.2 ± 0.2
Task-II	DBLP	AP	62.4 ± 11.5	64.7 ± 11.5
		AUROC	61.1 ± 9.9	62.3 ± 12.3
	Trivago	AP	58.6 ± 1.0	59.2 ± 1.0
		AUROC	58.7 ± 1.0	59.3 ± 1.0
	OGBN-MAG	AP	77.3 ± 0.7	78.8 ± 0.9
		AUROC	79.6 ± 0.3	80.9 ± 0.7

Table 7 Cost of PINS. The additional cost due to PINS is not large in terms of time and marginal in terms of memory requirements

	Method	OGBN-MAG	AMiner	MAG
Average Running Time Per Epoch	PCL w/o PINS	20.102	89.505	103.450
(Sec)	PCL+PINS	24.025	118.338	133.739
Average GPU Memory Usage	PCL w/o PINS	2.344	10.573	23.942
(GB)	PCL+PINS	2.347	10.575	23.945

6.3.4 RQ4. Effectiveness of P-IOS

To demonstrate the effectiveness of P-IOS, we compare PCL w/o P-IOS (Basic PCL) with PCL+P-IOS on both Task-I and Task-II in the DBLP, Trivago, and OGBN-MAG datasets. As shown in Table 6, the prediction performance is improved with P-IOS in all the settings. This result is intuitive since P-IOS mitigates information loss caused by partitioning via restoring each hyperedge. However, PCL+P-IOS is not applicable to large-scale hypergraphs since with P-IOS, resulting partitions are too large, causing an out-of-memory error.

6.3.5 RQ5. Tendency of PCL+PINS with Respect to the Number (or Size) of Partitions

At last, we investigate how the performance of PCL+PINS changes with respect to the number (or size) of partitions on both tasks. Note that the number of partitions and the size of each partition (i.e., number of nodes and number of hyperedges) is in inverse proportion (see Table 8). To this end, for each partition size, we show the AP score distribution of PCL+PINS over 10 data splits of each dataset. As shown in Fig. 5, there is no clear tendency between the number of partitions (or the size of

Table 8 Detailed statistics regarding partitions. $|\mathcal{P}|$ indicates the number of partitions, $\overline{|\mathcal{P}_i^v|}$ indicates the average number of nodes in each partition, $\overline{|\mathcal{P}_i^e|}$ indicates the average number of hyperedges in each partition, and $sd(|\mathcal{P}_i^v|)$ and $sd(|\mathcal{P}_i^e|)$ indicate the standard deviations of the numbers of nodes and hyperedges, respectively, in each partition

	OGBN-MAG			AMiner			MAG		
	Large	Medium	Small	Large	Medium	Small	Large	Medium	Small
$ \mathcal{P} $	32	64	128	256	512	1024	256	512	1024
$\overline{ \mathcal{P}_i^v }$	23012.2	11506.1	5753.0	51806.9	25903.5	12951.7	106720.2	53360.1	26680.1
$\overline{ \mathcal{P}_i^e }$	45454.3	24265.5	12659.5	122358.5	62491.5	31772.9	205835.5	108065.1	56134.2
$sd(\mathcal{P}_i^v)$	72.8	25.7	20.1	190.9	78.3	41.0	334.4	174.3	51.6
$sd(\mathcal{P}_i^e)$	15408.3	7765.6	5984.1	45433.8	25306.5	11955.4	37684.6	21684.5	12252.0




Fig. 5 Performance tendency of PCL+PINS with respect to the size of partitions. The boxplot describes the test AP score distribution over 10 data splits. For each dataset, boxplots are arranged in order of Large, Medium, and Small sizes. No clear tendency exists between the number/size of the partition and the performance of PCL+PINS

partitions) and the performance of PCL+PINS, and especially in AMINER, there is no clear difference in the distributions depending on the size of partitions.

7 Conclusion

In this work, in order to bridge the gap between previous studies and real-world applications of hypergraph learning, we make three contributions that are summarized as follows:

- In terms of tasks, we provide two novel pair-level hypergraph-learning tasks (hyperedge disambiguation and local clustering) that can be used for formulating various real-world problems.
- In terms of datasets, we propose two large-scale hypergraph datasets (AMINER and MAG) that enable the evaluation of hypergraph neural networks at scale.
- In terms of training methods, we suggest PCL, a scalable hypergraph learning method. PCL is based on hypergraph partitioning and contrastive learning, equipped with two additional techniques (PINS and P-IOs) for reducing infor-

mation loss caused by partitioning. We experimentally verify the superiority of PCL over 16 baseline methods on the proposed pairwise prediction tasks and the effectiveness of PCL+PINS and PCL+P-IOs.

For reproducibility, we make the source code and datasets used in the paper available at <https://github.com/kswoo97/pcl>.

Appendix: Additional experimental settings

Details of datasets

DBLP is a co-authorship hypergraph where nodes and hyperedges correspond to publications and authors, respectively. Each publication’s class is labeled according to its field of study. Trivago is a hotel-web search hypergraph where each node indicates each hotel and each hyperedge corresponds to a user. If a user (hyperedge) has visited the website of a particular hotel (node), the corresponding node is added to the respective user hyperedge. Furthermore, each hotel’s class is labeled based on the country in which it is located. OGBN-MAG is originally a heterogeneous graph that contains comprehensive academic information including venue, author, publication, and affiliation information. We transform this heterogeneous graph into a hypergraph as described in Sect. 4, while a label of each node (publication) indicates a published venue of the corresponding publication.

Details of graph-based baseline methods

Since graph representation models (Kipf and Welling 2017; Veličković et al. 2018) require ordinary graph structure as an input, we transform original hypergraph datasets into ordinary graph datasets by using *clique expansion*, where each hyperedge is replaced with a clique in the resulting graph. Formally, the clique expansion is a transformation of a given hyperedge set \mathcal{E} to a clique expanded edge set $\mathcal{E}_G = \bigcup_{e \in \mathcal{E}} \binom{e}{2}$.

Algorithm 2 Clique expansion with sampling

Input: Hyperedge set \mathcal{E} , sampling threshold k

Output: Clique-expansion graph E

```

1  $E \leftarrow \emptyset$ 
2 foreach  $e \in \mathcal{E}$  do
3   if  $|e| \leq k$  then
4      $E \leftarrow E \cup \binom{e}{2}$ 
5   else
6     foreach  $v \in e$  do
7        $\mathcal{V}' \leftarrow$  uniformly sample  $k$  number of nodes from  $e \setminus v$ 
8        $E \leftarrow E \cup \{\{v, v'\} : v' \in \mathcal{V}'\}$ 
9 return  $E$ 

```

Specifically, for full-graph datasets of DBLP and Trivago, we directly obtain \mathcal{E}_G from \mathcal{E} , the entire hyperedge set. For full-graph datasets of OGBN-MAG, the size of the resulting clique expanded edges is too large to be loaded into the main memory. To reduce its scale, we additionally employ sampling, as described in Algorithm 2. Specifically, for each hyperedge e' whose size is greater than k and for each constituent node $v \in e'$, we uniformly sample k other nodes from e' (line 7) and create k edges joining v and each of the k sampled nodes. Here, we set $k = 10$ for the OGBN-MAG dataset. We fail to create full-graph datasets of AMiner and MAG since clique expansion runs out of memory even with small k around 3, and thus we cannot perform experiments on them.

For partitioned-graph datasets of DBLP, Trivago, and OGBN-MAG, we apply clique expansion to the hyperedge set in each partition and use the resulting clique-expanded edge set as that of the corresponding partition. For partitioned-graph datasets of AMiner and MAG, due to the scalability issue, we apply the sampling strategy described in Algorithm 2 to each partition \mathcal{P}_i of \mathcal{P} (i.e., the input is $\mathcal{P}_i^{\mathcal{E}}$ instead of \mathcal{E}) and treat the resulting edge set as the edge set of the corresponding partition. Here, we set k to 10.

Details of hyperparameter settings

We now provide detailed hyperparameter settings of representation models and training methods. The number of layers and hidden dimension of all representation models are fixed to 2 and 128, respectively.

For representation models that are trained via supervised learning methods, we train each model for 100 epochs. We tune a learning rate of each model within $\{0.01, 0.001, 0.0001\}$. For every 10 epochs, we measure the validation AP score and save the model parameters. Then, we designate the checkpoint with the highest validation AP score as the final model parameters.

For representation models that are trained via all versions of PCL, we tune the number of self-supervised learning epochs within $\{25, 50\}$, while we set a broader search space, specifically $\{20, 40, 60, 80, 100\}$, for that of other self-supervised learning methods. We tune the learning rate of the self-supervised learning within $\{0.001, 0.0001\}$ for all self-supervised learning methods. In addition, for methods that require augmentation steps, we tune the extent of node feature augmentation p_v within $\{0.3, 0.4\}$, and the extent of topological augmentation p_e within $\{0.3, 0.4\}$. Furthermore, for methods that require negative samples for contrastive learning, we tune the number of negative samples N within $\{1, 2\}$. The temperature parameter τ for all self-supervised learning methods, and the scalar λ that controls the strength of inter-partition loss in PCL+PINS are both fixed to 0.5. Lastly, we train downstream task classifiers of all self-supervised learning methods with a learning rate of 0.001. We train the classifiers for 100 epochs, and for every 10 epochs, we measure the validation AP score and save the classifier parameters. Then, we designate the checkpoint with the highest validation AP score as the final classifier parameters.

Funding This work was supported by Samsung Electronics Co., Ltd., National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. NRF-2020R1C1C1008296), and Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2022-0-00157, Robust, Fair, Extensible Data-Centric Continual Learning) (No. 2019-0-00075, Artificial Intelligence Graduate School Program (KAIST)).

Code and Data Availability The source code used in this paper and the large-scale hypergraph datasets that we build are publicly available at <https://github.com/kswoo97/pcl> for reproducibility.

Declarations

Conflict of interest The authors have no relevant financial or non-financial interests to disclose.

References

- Addanki R, Battaglia P, Budden D, et al (2021) Large-scale graph representation learning with very deep GNNs and self-supervision. [arXiv:2107.09422](https://arxiv.org/abs/2107.09422), <https://doi.org/10.48550/arXiv.2107.09422>
- Ahmed I, Galoppo T, Hu X et al (2021) Graph regularized autoencoder and its application in unsupervised anomaly detection. *IEEE Trans Pattern Anal Mach Intell (TPAMI)* 44(8):4110–4124. <https://doi.org/10.1109/TPAMI.2021.3066111>
- Arya D, Gupta DK, Rudinac S, et al (2020) HyperSAGE: Generalizing inductive representation learning on hypergraphs. [arXiv:2010.04558](https://arxiv.org/abs/2010.04558), <https://doi.org/10.48550/arXiv.2010.04558>
- Bai S, Zhang F, Torr PH (2021) Hypergraph convolution and hypergraph attention. *Pattern Recogn* 110(107):637. <https://doi.org/10.1016/j.patcog.2020.107637>
- Benson AR, Abebe R, Schaub MT et al (2018) Simplicial closure and higher-order link prediction. *Proceed Nat Acad Sci* 115(48):E11221–E11230. <https://doi.org/10.1073/pnas.1800683115>
- Caldwell AE, Kahng AB, Markov IL (2000) Improved algorithms for hypergraph bipartitioning. In: *Proceedings of the 2000 Asia and South Pacific design automation conference (ASP-DAC)*, pp 661–666. <https://doi.org/10.1109/ASPAC.2000.835182>
- Caron E, van Eck NJ (2014) Large scale author name disambiguation using rule-based scoring and clustering. In: *Proceedings of the 19th international conference on science and technology indicators (STI)*, pp 79–86. https://doi.org/10.1007/978-981-32-9298-7_12
- Çatalyürek ÜV, Aykanat C (2011) PaToH (partitioning tool for hypergraphs). In: *Encyclopedia of parallel computing*. Springer, pp 1479–1487. https://doi.org/10.1007/978-0-387-09766-4_93
- Chen J, Ma T, Xiao C (2018a) FastGCN: Fast learning with graph convolutional networks via importance sampling. In: *International conference on learning representations (ICLR)*, [https://doi.org/10.48550/arXiv.1801.10247](https://arxiv.org/abs/1801.10247)
- Chen J, Zhu J, Song L (2018b) Stochastic training of graph convolutional networks with variance reduction. In: *International conference on machine learning (ICML)*, pp 942–950, <https://doi.org/10.48550/arXiv.1710.10568>
- Chen T, Kornblith S, Norouzi M, et al. (2020) A simple framework for contrastive learning of visual representations. In: *International conference on machine learning (ICML)*, pp 1597–1607, <https://doi.org/10.48550/arXiv.2002.05709>
- Chiang WL, Liu X, Si S, et al (2019) Cluster-GCN: An efficient algorithm for training deep and large graph convolutional networks. In: *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining (KDD)*, pp 257–266, <https://doi.org/10.1145/3292500.3330925>
- Chien E, Pan C, Peng J, et al (2021) You are AllSet: A multiset function framework for hypergraph neural networks. In: *International conference on learning representations (ICLR)*, [https://doi.org/10.48550/arXiv.2106.13264](https://arxiv.org/abs/2106.13264)
- Chodrow PS, Veldt N, Benson AR (2021) Generative hypergraph clustering: from blockmodels to modularity. *Sci Adv* 28:eabh1303. <https://doi.org/10.1126/sciadv.abh1303>
- Contisciani M, Battiston F, De Bacco C (2022) Inference of hyperedges and overlapping communities in hypergraphs. *Nat Commun* 13(1):7229. <https://doi.org/10.1038/s41467-022-34714-7>

- Deng K, Xing L, Zheng L et al (2019) A user identification algorithm based on user behavior analysis in social networks. *IEEE Access* 7:47114–47123. <https://doi.org/10.1109/ACCESS.2019.2909089>
- Do MT, Yoon Se, Hooi B, et al (2020) Structural patterns and generative models of real-world hypergraphs. In: Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining (KDD), pp 176–186. <https://doi.org/10.1145/3394486.3403060>
- Dong Y, Sawin W, Bengio Y (2020) HNHN: hypergraph networks with hyperedge neurons. *arXiv:2006.12278*, <https://doi.org/10.48550/arXiv.2006.12278>
- Feng Y, You H, Zhang Z, et al (2019) Hypergraph neural networks. In: Proceedings of the AAAI conference on artificial intelligence (AAAI), pp 3558–3565. <https://doi.org/10.1609/aaai.v33i01.33013558>
- Fey M, Lenssen JE (2019) Fast graph representation learning with pytorch geometric. *arXiv:1903.02428*, <https://doi.org/10.48550/arXiv.1903.02428>
- Gao T, Yao X, Chen D (2021) SimCSE: Simple contrastive learning of sentence embeddings. In: Proceedings of the 2021 conference on empirical methods in natural language processing (EMNLP), pp 6894–6910. <https://doi.org/10.18653/v1/2021.emnlp-main.552>
- Girvan M, Newman ME (2002) Community structure in social and biological networks. *Proc Natl Acad Sci* 99(12):7821–7826. <https://doi.org/10.1073/pnas.122653799>
- Grover A, Leskovec J (2016) Node2vec: scalable feature learning for networks. In: Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining (KDD), pp 855–864. <https://doi.org/10.1145/2939672.2939754>
- Grunig G, Durmus N, Zhang Y et al (2022) Molecular clustering analysis of blood biomarkers in world trade center exposed community members with persistent lower respiratory symptoms. *Int J Environ Res Public Health* 19(13):8102. <https://doi.org/10.3390/ijerph19138102>
- Guo M, Yi T, Zhu Y, et al (2021) JiTUNE: Just-in-time hyperparameter tuning for network embedding algorithms. *arXiv:2101.06427*, <https://doi.org/10.48550/arXiv.2101.06427>
- Hamilton WL, Ying R, Leskovec J (2017) Inductive representation learning on large graphs. In: Advances in neural information processing systems (NeurIPS), <https://doi.org/10.48550/arXiv.1706.02216>
- Hassani K, Khasahmadi AH (2020) Contrastive multi-view representation learning on graphs. In: International conference on machine learning (ICML), pp 4116–4126. <https://doi.org/10.48550/arXiv.2006.05582>
- He K, Fan H, Wu Y, et al (2020) Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR), pp 9729–9738. <https://doi.org/10.48550/arXiv.1911.05722>
- Hein M, Setzer S, Jost L, et al (2013) The total variation on hypergraphs-learning on hypergraphs revisited. In: Advances in neural information processing systems (NeurIPS), <https://doi.org/10.48550/arXiv.1312.5179>
- Huang J, Yang J (2021) UniGNN: A unified framework for graph and hypergraph neural networks. In: Proceedings of the Thirtieth international joint conference on artificial intelligence (IJCAI), pp 2563–2569. <https://doi.org/10.24963/ijcai.2021/353>
- Huang W, Zhang T, Rong Y, et al (2018) Adaptive sampling towards fast graph representation learning. In: Advances in neural information processing systems (NeurIPS), <https://doi.org/10.48550/arXiv.1809.05343>
- Hwang H, Lee S, Park C, et al (2022) AHP: learning to negative sample for hyperedge prediction. In: Proceedings of the 45th International ACM SIGIR conference on research and development in information retrieval (SIGIR), pp 2237–2242. <https://doi.org/10.1145/3477495.3531836>
- Jecmen S, Yoon M, Conitzer V, et al (2023) A dataset on malicious paper bidding in peer review. In: Proceedings of the ACM web conference 2023 (WWW), pp 3816–3826. <https://doi.org/10.1145/3543507.3583424>
- Karypis G, Aggarwal R, Kumar V, et al (1997) Multilevel hypergraph partitioning: Application in VLSI domain. In: Proceedings of the 34th annual Design Automation Conference (DAC), pp 526–529. <https://doi.org/10.1145/266021.266273>
- Kim S, Choe M, Yoo J, et al (2022) Reciprocity in directed hypergraphs: measures, findings, and generators. In: IEEE International conference on data mining (ICDM), <https://doi.org/10.1109/ICDM54844.2022.00122>
- Kingma DP, Ba J (2015) Adam: a method for stochastic optimization. In: International conference on learning representations (ICLR), <https://doi.org/10.48550/arXiv.1412.6980>

- Kipf TN, Welling M (2017) Semi-supervised classification with graph convolutional networks. In: International conference on learning representations (ICLR), <https://doi.org/10.48550/arXiv.1609.02907>
- Klamt S, Haus UU, Theis F (2009) Hypergraphs and cellular networks. *PLoS Comput Biol* 5(5):e1000385. <https://doi.org/10.1371/journal.pcbi.1000385>
- Ko J, Kook Y, Shin K (2022) Growth patterns and models of real-world hypergraphs. *Knowl Inf Syst* 64(11):2883–2920. <https://doi.org/10.1007/s10115-022-01739-9>
- Konstantinova EV, Skorobogatov VA (2001) Application of hypergraph theory in chemistry. *Discret Math* 235(1–3):365–383. [https://doi.org/10.1016/S0012-365X\(00\)00290-9](https://doi.org/10.1016/S0012-365X(00)00290-9)
- Koren Y, Bell R, Volinsky C (2009) Matrix factorization techniques for recommender systems. *Computer* 42(8):30–37. <https://doi.org/10.1109/MC.2009.263>
- Lee D, Shin K (2023) I'm me, we're us, and i'm us: Tri-directional contrastive learning on hypergraphs. In: Proceedings of the AAAI conference on artificial intelligence (AAAI), <https://doi.org/10.48550/arXiv.2206.04739>
- Lee G, Choe M, Shin K (2021) How do hyperedges overlap in real-world hypergraphs?-patterns, measures, and generators. In: Proceedings of the web conference 2021 (WWW), pp 3396–3407, <https://doi.org/10.1145/3442381.3450010>
- Lee J, Lee Y, Kim J, et al (2019) Set transformer: a framework for attention-based permutation-invariant neural networks. In: International conference on machine learning (ICML), pp 3744–3753, <https://doi.org/10.48550/arXiv.1810.00825>
- Li P, Milenkovic O (2018) Submodular hypergraphs: p-laplacians, cheeger inequalities and spectral clustering. In: International conference on machine learning (ICML), pp 3014–3023, <https://doi.org/10.48550/arXiv.1803.03833>
- Li Z, Huang C, Xia L, et al (2022) Spatial-temporal hypergraph self-supervised learning for crime prediction. In: IEEE 38th international conference on data engineering (ICDE), pp 2984–2996, <https://doi.org/10.1109/ICDE53745.2022.00269>
- Liu Z, Ma Y, Ouyang Y, et al (2021) Contrastive learning for recommender system. [arXiv:2101.01317](https://arxiv.org/abs/2101.01317), <https://doi.org/10.48550/arXiv.2101.01317>
- Luo Q, Yu D, Cai Z, et al (2021) Hypercore maintenance in dynamic hypergraphs. In: IEEE 37th international conference on data engineering (ICDE), pp 2051–2056, <https://doi.org/10.1109/ICDE51399.2021.00199>
- Luo X, Ju W, Qu M, et al (2022) DualGraph: Improving semi-supervised graph classification via dual contrastive learning. In: IEEE 38th international conference on data engineering (ICDE), pp 699–712, <https://doi.org/10.1109/ICDE53745.2022.00057>
- Malatras A, Geneiatakis D, Vakalis I (2017) On the efficiency of user identification: a system-based approach. *Int J Inf Secur* 16(6):653–671. <https://doi.org/10.1007/s10207-016-0340-2>
- Maleki S, Saless D, Wall DP, et al (2022) HyperNetVec: Fast and scalable hierarchical embedding for hypergraphs. In: *Network Science (NetSci)*, Springer, pp 169–183, https://doi.org/10.1007/978-3-030-97240-0_13
- Mayer C, Mayer R, Bhowmik S, et al (2018) HYPE: Massive hypergraph partitioning with neighborhood expansion. In: IEEE International conference on big data (Big Data), pp 458–467, <https://doi.org/10.1109/BigData.2018.8621968>
- Milojević S (2013) Accuracy of simple, initials-based methods for author name disambiguation. *J Informet* 7(4):767–773. <https://doi.org/10.1016/j.joi.2013.06.006>
- Muttakin MN, Hossain MI, Rahman MS (2022) Overlapping community detection using dynamic dilated aggregation in deep residual GCN. [arXiv:2210.11174](https://arxiv.org/abs/2210.11174), <https://doi.org/10.48550/arXiv.2210.11174>
- Nair V, Hinton GE (2010) Rectified linear units improve restricted boltzmann machines. In: Proceedings of the 27th international conference on machine learning (ICML-10), pp 807–814
- Oord Avd, Li Y, Vinyals O (2018) Representation learning with contrastive predictive coding. [arXiv:1807.03748](https://arxiv.org/abs/1807.03748), <https://doi.org/10.48550/arXiv.1807.03748>
- Paszke A, Gross S, Massa F, et al (2019) PyTorch: An imperative style, high-performance deep learning library. In: Advances in Neural information processing systems (NeurIPS), <https://doi.org/10.48550/arXiv.1912.01703>
- Pennington J, Socher R, Manning CD (2014) GloVe: Global vectors for word representation. In: Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP), pp 1532–1543, <https://doi.org/10.3115/v1/D14-1162>
- Qu C, Tao M, Yuan R (2018) A hypergraph-based blockchain model and application in internet of things-enabled smart homes. *Sensors* 18(9):2784. <https://doi.org/10.3390/s18092784>

- Robinson J, Chuang CY, Sra S, et al (2021) Contrastive learning with hard negative samples. In: International conference on learning representations (ICLR), <https://doi.org/10.48550/arXiv.2010.04592>
- Rossetti G, Cazabet R (2018) Community discovery in dynamic networks: a survey. *ACM Comput Surv* 51(2):1–37. <https://doi.org/10.1145/3172867>
- Rossi E, Frasca F, Chamberlain B, et al (2020) SIGN: scalable inception graph neural networks. *arXiv:2004.11198*, <https://doi.org/10.48550/arXiv.2004.11198>
- Rossi R, Ahmed N (2015) The network data repository with interactive graph analytics and visualization. In: Proceedings of the AAAI conference on artificial intelligence (AAAI), <https://doi.org/10.1609/aaai.v29i1.9277>
- Ruan B, Gan J, Wu H, et al (2021) Dynamic structural clustering on graphs. In: Proceedings of the 2021 international conference on management of data (SIGMOD), pp 1491–1503. <https://doi.org/10.1145/3448016.3452828>
- Rumelhart DE, Hinton GE, Williams RJ (1986) Learning representations by back-propagating errors. *Nature* 323(6088):533–536. <https://doi.org/10.1038/323533a0>
- Sanyal DK, Bhowmick PK, Das PP (2021) A review of author name disambiguation techniques for the PubMed bibliographic database. *J Inf Sci* 47(2):227–254. <https://doi.org/10.1177/0165551519888605>
- Schlag S, Heuer T, Gottesbüren L et al (2023) High-quality hypergraph partitioning. *ACM J Exp Algorithms* 27:1–39. <https://doi.org/10.1145/3529090>
- Shchur O, Günnemann S (2019) Overlapping community detection with graph neural networks. *arXiv:1909.12201*, <https://doi.org/10.48550/arXiv.1909.12201>
- Sinha A, Shen Z, Song Y, et al (2015) An overview of microsoft academic service (MAS) and applications. In: Proceedings of the 24th international conference on world wide web (WWW), pp 243–246. <https://doi.org/10.1145/2740908.2742839>
- Tang J, Zhang J, Yao L, et al (2008) ArnetMiner: Extraction and mining of academic social networks. In: Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining (KDD), pp 990–998. <https://doi.org/10.1145/1401890.1402008>
- Torres L, Blevins AS, Bassett D et al (2021) The why, how, and when of representations for complex systems. *SIAM Rev* 63(3):435–485. <https://doi.org/10.1137/20M1355896>
- Tsitsulin P, Palowitch J, Perozzi B, et al (2020) Graph clustering with graph neural networks. *arXiv:2006.16904*, <https://doi.org/10.48550/arXiv.2006.16904>
- Tynes M, Gao W, Burrill DJ et al (2021) Pairwise difference regression: a machine learning meta-algorithm for improved prediction and uncertainty quantification in chemical search. *J Chem Inf Model* 61(8):3846–3857. <https://doi.org/10.1021/acs.jcim.1c00670>
- Veličković P, Fedus W, Hamilton WL, et al (2018) Deep graph infomax. In: International conference on learning representations (ICLR), <https://doi.org/10.48550/arXiv.1809.10341>
- Wang K, Shen Z, Huang C et al (2020) Microsoft academic graph: when experts are not enough. *Quant Sci Stud* 1(1):396–413. https://doi.org/10.1162/qss_a_00021
- Wang Z, Zheng L, Li Y, et al (2019) Linkage based face clustering via graph convolution network. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR), pp 1117–1125. <https://doi.org/10.1109/CVPR.2019.00121>
- Wu F, Souza A, Zhang T, et al (2019) Simplifying graph convolutional networks. In: International conference on machine learning (ICML), pp 6861–6871. <https://doi.org/10.48550/arXiv.1902.07153>
- Xia X, Yin H, Yu J, et al (2021) Self-supervised hypergraph convolutional networks for session-based recommendation. In: Proceedings of the AAAI conference on artificial intelligence (AAAI), pp 4503–4511. <https://doi.org/10.1609/aaai.v35i5.16578>
- Xie J, Kelley S, Szymanski BK (2013) Overlapping community detection in networks: the state-of-the-art and comparative study. *ACM Comput Surv* 45(4):1–35. <https://doi.org/10.1145/2501654.2501657>
- Xie X, Sun F, Liu Z, et al (2022) Contrastive learning for sequential recommendation. In: IEEE 38th International conference on data engineering (ICDE), pp 1259–1273. <https://doi.org/10.1109/ICDE53745.2022.00099>
- Yadati N, Nimishakavi M, Yadav P, et al (2019) HyperGCN: a new method of training graph convolutional networks on hypergraphs. In: Advances in neural information processing systems (NeurIPS), pp 1509–1520. <https://doi.org/10.48550/arXiv.1809.02589>
- Yadati N, Nitin V, Nimishakavi M, et al (2020) NHP: neural hypergraph link prediction. In: Proceedings of the 29th ACM international conference on information & knowledge management (CIKM), pp 1705–1714. <https://doi.org/10.1145/3340531.3411870>

- Yang J, Leskovec J (2011) Patterns of temporal variation in online media. In: Proceedings of the fourth ACM international conference on Web search and data mining (WSDM), pp 177–186, <https://doi.org/10.1145/1935826.1935863>
- Yang J, Leskovec J (2013) Overlapping community detection at scale: a nonnegative matrix factorization approach. In: Proceedings of the sixth ACM international conference on Web search and data mining (WSDM), pp 587–596, <https://doi.org/10.1145/2433396.2433471>
- Yin N, Feng F, Luo Z, et al (2022) Dynamic hypergraph convolutional network. In: IEEE 38th international conference on data engineering (ICDE), pp 1621–1634, <https://doi.org/10.1109/ICDE53745.2022.00167>
- You Y, Chen T, Sui Y, et al (2020) Graph contrastive learning with augmentations. In: Advances in neural information processing systems (NeurIPS), pp 5812–5823, <https://doi.org/10.48550/arXiv.2010.13902>
- Yu J, Yin H, Li J, et al (2021) Self-supervised multi-channel hypergraph convolutional network for social recommendation. In: Proceedings of the web conference 2021 (WWW), pp 413–424, <https://doi.org/10.1145/3442381.3449844>
- Zaheer M, Kottur S, Ravanbakhsh S, et al (2017) Deep sets. In: Advances in neural information processing systems (NeurIPS), <https://doi.org/10.48550/arXiv.1703.06114>
- Zeng H, Zhou H, Srivastava A, et al (2019) GraphSAINT: Graph sampling based inductive learning method. In: International conference on learning representations (ICLR), <https://doi.org/10.48550/arXiv.1907.04931>
- Zhang D, Huang X, Liu Z, et al (2020a) AGL: A scalable system for industrial-purpose graph machine learning. Proc VLDB Endow (PVLDB) 13(12): 3125–3137. <https://doi.org/10.14778/3415478.3415539>
- Zhang F, Liu X, Tang J, et al (2019) OAG: Toward linking large-scale heterogeneous entity graphs. In: Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining (KDD), pp 2585–2595, <https://doi.org/10.1145/3292500.3330785>
- Zhang J, Gao M, Yu J, et al (2021) Double-scale self-supervised hypergraph learning for group recommendation. In: Proceedings of the 30th ACM international conference on information & knowledge management (CIKM), pp 2557–2567, <https://doi.org/10.1145/3459637.3482426>
- Zhang J, Li F, Xiao X, et al (2022) Hypergraph convolutional networks via equivalency between hypergraphs and undirected graphs. [arXiv:2203.16939](https://arxiv.org/abs/2203.16939), <https://doi.org/10.48550/arXiv.2203.16939>
- Zhang R, Zou Y, Ma J (2020b) Hyper-SAGNN: a self-attention based graph neural network for hypergraphs. In: International conference on learning representations (ICLR), <https://doi.org/10.48550/arXiv.1911.02613>
- Zhang S, Tong H (2016) FINAL: Fast attributed network alignment. In: Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data (KDD), pp 1345–1354, <https://doi.org/10.1145/2939672.2939766>
- Zheng C, Chen H, Cheng Y, et al. (2022a) ByteGNN: efficient graph neural network training at large scale. Proc VLDB Endow (PVLDB) 15(6):1228–1242. <https://doi.org/10.14778/3514061.3514069>
- Zheng D, Ma C, Wang M, et al (2020) DistDGL: Distributed graph neural network training for billion-scale graphs. In: IEEE/ACM 10th workshop on irregular applications: architectures and algorithms (IA3), pp 36–44, <https://doi.org/10.1109/IA351965.2020.00011>
- Zheng Y, Pan S, Lee VC, et al (2022b) Rethinking and scaling up graph contrastive learning: an extremely efficient approach with group discrimination. In: Advances in neural information processing systems (NeurIPS), pp 10809–10820, <https://doi.org/10.48550/arXiv.2206.01535>
- Zhu Y, Xu Y, Yu F, et al (2020) Deep graph contrastive representation learning. [arXiv:2006.04131](https://arxiv.org/abs/2006.04131), <https://doi.org/10.48550/arXiv.2006.04131>
- Zhu Y, Xu Y, Yu F, et al (2021) Graph contrastive learning with adaptive augmentation. In: Proceedings of the web conference 2021 (WWW), pp 2069–2080, <https://doi.org/10.1145/3442381.3449802>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

Authors and Affiliations

Sunwoo Kim¹ · Dongjin Lee² · Yul Kim³ · Junggho Park³ · Taeho Hwang³ ·
Kijung Shin^{1,2} 

✉ Kijung Shin
kijungs@kaist.ac.kr

Sunwoo Kim
kswoo97@kaist.ac.kr

Dongjin Lee
dongjin.lee@kaist.ac.kr

Yul Kim
yul5.kim@samsung.com

Junggho Park
j0106.park@samsung.com

Taeho Hwang
taeho.hwang@samsung.com

¹ Kim Jaechul Graduate School of AI, KAIST, Seoul, South Korea

² School of Electrical Engineering, KAIST, Daejeon, South Korea

³ Samsung Research, Seoul, South Korea