

Hierarchical Prototype Networks for Continual Graph Representation Learning

Xikun Zhang, Dongjin Song*, *Member, IEEE*, and Dacheng Tao*, *Fellow, IEEE*

Abstract—Despite significant advances in graph representation learning, little attention has been paid to the more practical continual learning scenario in which new categories of nodes (*e.g.*, new research areas in citation networks, or new types of products in co-purchasing networks) and their associated edges are continuously emerging, causing catastrophic forgetting on previous categories. Existing methods either ignore the rich topological information or sacrifice plasticity for stability. To this end, we present Hierarchical Prototype Networks (HPNs) which extract different levels of abstract knowledge in the form of prototypes to represent the continuously expanded graphs. Specifically, we first leverage a set of Atomic Feature Extractors (AFEs) to encode both the elemental attribute information and the topological structure of the target node. Next, we develop HPNs to adaptively select relevant AFEs and represent each node with three levels of prototypes. In this way, whenever a new category of nodes is given, only the relevant AFEs and prototypes at each level will be activated and refined, while others remain uninterrupted to maintain the performance over existing nodes. Theoretically, we first demonstrate that the memory consumption of HPNs is bounded regardless of how many tasks are encountered. Then, we prove that under mild constraints, learning new tasks will not alter the prototypes matched to previous data, thereby eliminating the forgetting problem. The theoretical results are supported by experiments on five datasets, showing that HPNs not only outperform state-of-the-art baseline techniques but also consume relatively less memory. Code and datasets are available at <https://github.com/QueueQ/HPNs>

Index Terms—Graph representation learning, continual learning, hierarchical prototype, graph neural networks.

1 INTRODUCTION

GRAPH representation learning aims to pursue a meaningful vector representation of each node so as to facilitate downstream applications such as node classification, link prediction, *etc.* Traditional methods are developed based on graph statistics [1] or hand-crafted features [2], [3]. Recently, a great amount of attention has been paid to graph neural networks (GNNs), such as graph convolutional network (GCNs) [4], GraphSAGE [5], Graph Attention Networks (GATs) [6], and their extensions [7], [8], [9], [10], [11], [12]. This is because they can jointly consider the feature and topological information of each node. Most of these approaches, however, focus on static graphs and cannot generalize to the case when new categories of nodes are emerging.

In many real world applications, different categories of nodes and their associated edges (in the form of subgraphs) are often continuously emerging in existing graphs. For instance, in a citation network [13], [14], [15], papers describing new research areas will gradually appear in the citation graph; in a co-purchasing network such as Amazon [16], new types of products will continuously be added to the graph. Given these facts, how to incorporate the feature and topological information of new nodes in a continuous and effective manner such that performance over existing nodes is uninterrupted is a critical problem to investigate, especially

when graphs are relatively large and retraining a new model over an entire graph is computationally expensive.

To address this issue, a few attempts have been made to tailor the existing continual learning techniques to graph data. For instance, adopting the memory replay based approaches, Zhou et al. [17] proposed to store a set of representative experience nodes in a buffer and replay them along with new tasks (categories) to prevent forgetting existing tasks (categories). The buffer, however, only stores node features and ignores the topological information of graphs. Inspired by the regularization based methods, Liu et al. [18] developed topology-aware weight preserving (TWP) that can preserve the topological information of existing graphs. However, its design hinders the capability of learning topology on new tasks (categories). More detailed introduction on these related works can be found in Section 2.1.

A desired learning system for continual graph representation learning is expected to continuously grasp knowledge from new categories of emerging nodes and capture their topological structures without interfering with the learned knowledge over existing graphs. Inspired by the fact that humans learn to recognize objects by forming prototypes in the brain [19], [20], [21] and can achieve extraordinary continual learning capability, we present a completely novel framework, *i.e.*, Hierarchical Prototype Networks (HPNs), to continuously extract different levels of abstract knowledge (in the form of prototypes) from graph data such that new knowledge will be accommodated while earlier experience can still be well retained. Within this framework, representation learning is simultaneously conducted to avoid catastrophic forgetting, instead of considering these two objectives separately. Although learning independent prototypes for different tasks is appealing for preventing forgetting, it may

- Mr X. Zhang and Prof D. Tao are with the School of Computer Science, in the Faculty of Engineering, at The University of Sydney, 6 Cleveland St, Darlington, NSW 2008, Australia. Email: xzha0505@uni.sydney.edu.au, dacheng.tao@sydney.edu.au.
- Dr. D. Song is with the Department of Computer Science and Engineering, University of Connecticut, Storrs, Connecticut, the United States. Email: dongjin.song@uconn.edu.

* indicates corresponding authors.

incur unbounded memory consumption with an unlimited number of new tasks. Therefore, we propose to denote each node with a composition of basic prototypes so that nodes from different tasks and their associated relational structures are represented with compositions of a limited set of basic prototypes. Take the social network as an example, if a node (person) falls into certain category, it can be decomposed into basic atomic characteristics belonging to a set of attributes (*e.g.*, gender, nationality, hobby, *etc.*) and the relationship between a pair of nodes can also be categorized into different basic types (*e.g.*, friends, coworkers, *etc.*).

Inspired by these facts, we first develop the Atomic Feature Extractors (AFE) to decompose each node into two sets of atomic embeddings, *i.e.*, atomic node embeddings which encode the node feature information and atomic structure embeddings which encode its relations to neighboring nodes within multi-hop. Next, we present Hierarchical Prototype Networks to adaptively select, compose, and store representative embeddings with three levels of prototypes, *i.e.*, atomic-level, node-level, and class-level. Given a new node, only the relevant AFEs and prototypes in each level will be activated and refined, while others are uninterrupted.

The memory efficiency and the continual learning capability of the proposed HPNs are theoretically validated. First, adopting concepts from geometry and coding theory, we show the equivalence between the prototypes of HPNs and the spherical codes, thereby deriving the upper bound of the number of prototypes. This indicates that our memory consumption is bounded regardless of the number of tasks encountered. Second, since the forgetting problem can be formalized as the prediction drift of previous data after learning new tasks, we derive the conditions under which the learning on new tasks will not alter the predictions of previous data, *i.e.* forgetting is eliminated.

The theoretical merits of the proposed HPNs are justified by experiments on five public datasets. First, the performance of HPNs not only achieves the state-of-the-art, but is also comparable to or better than the jointly trained model (set as the upper bound of continual learning models). Second, the HPNs are memory efficient. For instance, on OGB-Products dataset containing more than 2 million nodes and 47 categories of nodes, HPNs achieve around 80% accuracy with only thousands of parameters.

2 RELATED WORKS

The proposed Hierarchical Prototype Networks (HPNs) are closely related to continual learning and graph representation learning. In this section, we provide more detailed discussions by comparing HPNs with related works, especially on those directly applying existing continual learning techniques on graph data.

2.1 Continual learning

Continual learning aims to overcome the well-known catastrophic forgetting problem that a model’s performance on previous tasks decreases significantly after being trained on new tasks. Existing works for continual learning can be categorized as regularization-based methods, memory-replay based methods, and parametric isolation based methods.

In the following, we give detailed introductions on these approaches, as well as their applicability to graph data.

Regularization-based methods penalize the model objectives to maintain satisfactory performance on previous tasks [22], [23], [24], [25], [26]. For instance, Li and Hoiem [23] introduced Learning without Forgetting (LwF) which uses knowledge distillation to constrain the shift of parameters for old tasks; Kirkpartick et al. [24] proposed elastic weight consolidation (EWC) that adds a quadratic penalty to prevent the model weights from shifting too much. Recent works [25], [26] seek to constrain the gradients for new tasks in a subspace orthogonal to the updating directions that are important for previous tasks. Although the regularization-based methods can alleviate the forgetting on previous tasks, the constraints on the model weights reduce a model’s plasticity for new tasks, resulting in inferior performance compared to the other two approaches. Regularization-based methods can be directly applied to graph neural networks and are included as baselines in our experiments.

Memory-replay based methods constantly feed a model with representative data of previous tasks to prevent forgetting [27], [28], [29], [30], [31]. One example is Gradient Episodic Memory (GEM) [27] that stores representative data in episodic memory and adds a constraint to prevent the loss of the episodic memory from increasing and only allow it to decrease. Instead of storing data, Shin et al. [28] added a generative model to generate pseudo-data of previous tasks to be interleaved with new task data for rehearsal. Recent works also look for better designs of memory replay to facilitate continual learning agents [30], [31]. Memory-replay based methods are currently one of the most effective approaches for alleviating catastrophic forgetting, but these methods are not suitable for graph data, as they cannot store the topological information, which is crucial for representing graph data. In contrast, our proposed HPNs explicitly designed Atomic Feature Extractors (AFE) to encode both node and topological information. Moreover, memory-replayed methods require rehearsal of old data each time when a new task is learned. This introduces extra computation burdens.

Parametric isolation based methods adaptively introduce new parameters for new tasks to avoid the parameters of previous tasks being drastically changed [32], [33], [34], [35], [36]. For instance, progressive network [32] allocates a new sub-network for each new task and block any modification on the previously learned networks. Yoon et al. [33] proposed a more flexible model (DEN) that dynamically adds new neurons to accommodate new tasks. Recently, various innovative approaches to allocate separated parameters for different tasks have been developed [34], [35], [36]. Besides the concrete models, Knoblauch et al. [37] analyzed the required capability of an optimal continual learning agent. Although parametric isolation based methods are effective for alleviating the catastrophic forgetting problem, they also consistently increase the model complexity and memory consumption which could be a problem. Recently developed methods have used certain mechanisms to control memory consumption, such as merging parameters for similar tasks [34]. However, as task specific parameters are fitted to each individual task, these parameters are seldom reused. In the proposed HPNs, we consider the properties of graph data and decompose each node into a combination of several

prototypes. As these prototypes are largely shared by all tasks, the parameter reuse level is greatly enhanced and only limited memory needs to be consumed. Moreover, we also derive a theoretical memory upper bound for HPNs.

Overall, although existing continual learning methods can perform well on Euclidean data, they are not suitable to be directly applied to graph data, which is also shown in our experiments. Our proposed HPNs are specially designed to overcome their aforementioned limitations.

2.2 Graph representation learning

Graph representation learning aims to encode both the feature information from nodes and the topological structure of the incoming graph. Traditional methods relied on graph statistics or hand-crafted features [2], [3]. Recently, a great amount of attention has been paid to graph neural networks (GNNs), such as graph convolutional network (GCNs) [4], GraphSAGE [5], Graph Attention Networks (GATs) [6], and their extensions [7], [8], [38], [39], [40], [41], [42], [43], [44]. Instead of focusing on shallow networks, there are also works on building deep GCNs to further increase the capacity of GNNs [10], [11], [12], [45], [46]. These models are not designed for the continual learning setting and will experience catastrophic forgetting problem when learning a sequence of tasks, which is shown in our experiments.

Currently, only limited efforts have been made to pursue continual graph representation. Zhou et al. [17] integrated memory-replay to GNNs by storing experience nodes from previous tasks. However, the topological structure of graphs is ignored. Liu et al. [18] developed topology-aware weight preserving (TWP) that can preserve the topological information of previous graphs. However, preserving topology of previous graphs will hamper its capability of learning topology on new graphs. Galke et al. [47] considers a scenario where new classes of nodes may appear, which is similar to our consideration. But they focus on adapting the model to new patterns and do not explicitly maintain the performance on previous tasks. Streaming GNN [48] and Feature graph network (FGN) [49] are also related to GNNs and continual learning. However, the setting of Streaming GNN is time step incremental, while our setting is task incremental. FGN transforms node classification into graph classification by constructing feature graphs, so as to apply existing continual learning techniques. However, the topological information is not fully utilized since the information aggregation for each node does not include the neighboring nodes.

Note that continual graph representation learning is essentially different from dynamic graph representation learning [50], [51], [52], [53] and few-shot graph representation learning [54], [55], [56]. Dynamic graph representation learning focuses on learning the evolving representations of nodes over time, while in a continual learning setting we constantly deal with different graphs (different tasks) and the model is not allowed to access the previously observed graphs. Different from continual learning which aims to overcome the catastrophic forgetting problem, few-shot graph representation learning targets at fast adaption to new tasks and follows a completely different setting. Specifically, for few-shot learning, a model is first trained on several meta-training tasks and then evaluated on the meta-testing

tasks. During the evaluation, a few-shot learning model is independently trained and evaluated on each task, while a continual learning model learns sequentially on testing tasks and is eventually evaluated over all previous tasks.

3 HIERARCHICAL PROTOTYPE NETWORKS

In this section, we first state the problem we aim to study and the notations. Then we present Hierarchical Prototype Networks (HPNs) that consist of two core modules, *i.e.*, Atomic Feature Extractor (AFEs) and Hierarchical Prototypes (HPs), as shown in Figure 1. AFEs serve to extract a set of atomic features from the given graph, and the HPs aim to select, compose, and store the representative features in the form of different levels of prototypes. During the training stage, each node will only refine the relevant AFEs and prototypes of the model without interfering with the irrelevant parts (*i.e.*, to avoid catastrophic forgetting). In the test stage, the model will activate the relevant AFEs and prototypes to perform the inference.

3.1 Problem Statement and Notations

We study continual learning on graphs that have new categories of nodes and associated edges (in the form of subgraphs) emerging in a continuous manner. In the context of continual learning, assuming we have a sequence of p tasks $\{\mathcal{T}^s | s = 1, \dots, p\}$, in which each task \mathcal{T}^s aims to learn a satisfied representation for a new subgraph \mathcal{G}_s consisting of nodes belonging to some new categories. A desired model should maintain its performance on all previous tasks after being successively trained on the sequence of p tasks from \mathcal{T}^1 to \mathcal{T}^p .

For simplicity, we omit the subscripts in this section. Full notations will be used in the theoretical analysis. Each graph \mathcal{G} consists of a node set $\mathbb{V} = \{v_i | i = 1, \dots, N\}$ with N nodes and an edge set $\mathbb{E} = \{(v_i, v_j)\}$ denoting the connections of nodes in \mathbb{V} . Each node v_i can be represented as a feature vector $\mathbf{x}(v_i) \in \mathbb{R}^{d_v}$ that encodes node attributes, *e.g.*, gender, nationality, hobby, *etc.* The set of l -hop neighboring nodes of v_i is defined as $\mathcal{N}^l(v_i)$, with $\mathcal{N}^0(v_i) = \{v_i\}$.

3.2 Atomic Feature Extractors

As mentioned in the Introduction, to handle unlimited number of potential new categories of nodes with a limited memory consumption, we represent each node as a composition of basic prototypes selected from a limited set. To construct or refine these basic prototypes, given the input graph, we first need to obtain basic features with feature extractors. Specifically, we develop Atomic Feature Extractors (AFEs) to consider two different sets of atomic embeddings, *i.e.*, atomic node embeddings which encode the target node features and atomic structure embeddings that encode its relations to multi-hop neighbors. To ensure these generated basic prototypes are capable of encoding low-level features that can be shared across different tasks, we avoid deep structures in the AFEs and formulate them as learnable linear transformations: $\text{AFE}_{\text{node}} = \{\mathbf{A}_m \in \mathbb{R}^{d_v \times d_a} | m \in \{1, \dots, l_a\}\}$ and $\text{AFE}_{\text{struct}} = \{\mathbf{R}_n \in \mathbb{R}^{d_v \times d_r} | n \in \{1, \dots, l_r\}\}$ where \mathbf{A}_m and \mathbf{R}_n are real matrices to encode atomic node and structure information, respectively. Each matrix \mathbf{A}_m or \mathbf{R}_n




Fig. 1: The framework of HPNs. On the left, subgraphs of different tasks come in sequentially, with the current task denoted as \mathcal{T}_s . Given a node v , u_k^j denotes the j -th sampled node from k -hop neighbors. In the middle, node v and the sampled neighbors ($\mathcal{N}_{sub}(v)$) are fed into the selected AFEs ($\text{AFE}_{node}^{select}$ or $\text{AFE}_{struct}^{select}$) to get atomic embeddings ($\mathbb{E}_A^{select}(v)$), which are either matched to existing A-prototypes (\mathbb{P}_A) denoted with black dashed arrows or used to establish new A-prototypes (red dashed arrows). The selected A-prototypes $\mathbb{A}(v)$ are first mapped to node- and class-level embeddings ($\mathbb{E}_N(v)$ and $\mathbb{E}_C(v)$), which are further matched to a N- and a C-prototypes (\mathbb{P}_N and \mathbb{P}_C) to get the node- and class-level prototype representations ($\mathbb{N}(v)$ and $\mathbb{C}(v)$). Finally, $\mathbb{A}(v)$, $\mathbb{N}(v)$ and $\mathbb{C}(v)$ form the hierarchical prototype representation ($\mathbb{P}_H(v)$), which is fed into the classifier to perform node classification.

corresponds to a certain type of features, and l_a, l_r denote the number of matrices in AFE_{node} and AFE_{struct} , respectively. Given a node v , a set of atomic node embeddings is obtained by applying AFE_{node} to the feature vector $\mathbf{x}(v)$:

$$\mathbb{E}_A^{node}(v) = \{\mathbf{x}^T(v)\mathbf{A}_m | \mathbf{A}_m \in \text{AFE}_{node}\}. \quad (1)$$

To obtain atomic structure embeddings, the multi-hop neighboring nodes are considered to encode the relationship between node v and neighboring nodes within a fixed number of hops. Since the number of neighboring nodes varies from node to node, for computation efficiency in the following procedure, we uniformly sample a fixed number of vertices from 1-hop up to h -hop neighborhood. The sampled neighborhood set is denoted as, $\mathcal{N}_{sub}(v) \subseteq \bigcup_{l \in \{1, \dots, h\}} \mathcal{N}^l(v)$.

Then these selected neighboring nodes are embedded via matrices in AFE_{struct} to encode different types of interactions between the target node v and its neighbors:

$$\mathbb{E}_A^{struct}(v) = \{\mathbf{x}^T(u)\mathbf{R}_m | \mathbf{R}_m \in \text{AFE}_{struct}, u \in \mathcal{N}_{sub}\}. \quad (2)$$

Above all, with AFE_{node} and AFE_{struct} , the attribute and structure information of each node v is encoded in the two set $\mathbb{E}_A^{node}(v)$ and $\mathbb{E}_A^{struct}(v)$. For convenience of notation, we denote the complete atomic feature set of a node v as the union of the node atomic embedding set AFE_{node} and the structure embedding set AFE_{struct} , i.e.,

$$\mathbb{E}_A(v) = \mathbb{E}_A^{node}(v) \cup \mathbb{E}_A^{struct}(v). \quad (3)$$

Note that different \mathbf{A}_i and \mathbf{R}_i are designed to generate different types of atomic features. Therefore, we impose a

divergence loss on AFEs to ensure they are uncorrelated with each other and thus can map features to different subspaces:

$$\mathcal{L}_{div} = \sum_{m=1}^{l_a} \sum_{\substack{n=1 \\ n \neq m}}^{l_a} \mathbf{A}_m^T \mathbf{A}_n + \sum_{m=1}^{l_r} \sum_{\substack{n=1 \\ n \neq m}}^{l_r} \mathbf{R}_m^T \mathbf{R}_n. \quad (4)$$

3.3 Hierarchical Prototype Networks

Based on the atomic embeddings of the incoming subgraph, HPNs will distill and store the representative features in form of different levels of prototypes to memorize the current task, as shown in Figure 1. This is achieved by refining existing prototypes if the input subgraph only contains features observed before and creating new prototypes if new features are encountered. Specifically, HPNs will produce three levels of prototypes. First, the atomic-level prototypes (A-prototypes) are directly refined or created with the atomic embedding set, and they denote the prototypical basic features of nodes (e.g., gender, nationality, and social relation of people in a social network). Since the A-prototypes only describe low-level features of a node, which is similar to the output from the initial layers of a deep neural network, we further develop node-level (N-prototypes) and class-level prototypes (C-prototypes) to capture higher level information of each node. The N-prototypes are generated by associating the A-prototypes with each node to describe the node as a whole. The C-prototypes are generated from the N-prototypes to describe common features shared by a group of similar nodes. From atomic-level to class-level, the prototypes denote abstract knowledge of the subgraph

at different scales which is analog to the feature maps of convolutional neural networks at different layers.

Next, we introduce when and how HPNs will refine existing prototypes or establish new ones with the atomic embeddings. As mentioned earlier, for each task that contains certain categories of nodes, only the most relevant AFEs and prototypes will be activated while others remain uninterrupted to avoid forgetting. In other words, for each node, we will choose a subset of $\mathbb{E}_A(v)$ to refine the prototypes. Specifically, as shown in Figure 1, given a node from an incoming subgraph, The AFEs whose atomic embeddings are close enough to existing A-prototypes are chosen as the relevant ones. The closeness (distance) between embeddings and prototypes is measured with cosine similarity for two considerations. The first consideration is the numerical stability. Since we are going to develop an objective (Equation 8) to minimize the distance between the matched embeddings and prototypes, if the embeddings and the prototypes are not normalized, Equation 8 becomes maximizing the inner product between two unnormalized vectors. Unlike the cosine similarity that only adjusts the angle between two vectors and reaches the upper bound of 1 when the angle is 0, the inner product is unbounded and maximizing the inner product causes the vector lengths to increase to infinity, leading to the collapse of the model. Second, to guarantee an upper bound on the memory consumption, we restrict the prototypes on a unit sphere via normalization, so that only a limited number of prototypes will be created (Theorem 1). Due to the normalization, cosine distance is a natural choice as the distance metric.

Formally, we first obtain $\mathbb{E}_A^{\text{node}}(v)$ and $\mathbb{E}_A^{\text{struct}}(v)$ via Eq. (1) and Eq. (2), respectively. Then, we calculate the maximum cosine similarity between atomic embeddings of each AFE (\mathbf{e}_i) and the A-prototypes as:

$$\text{SimMAX}_i^{\text{id}} = \max_{\mathbf{p}} \left(\frac{\mathbf{e}_i^T \mathbf{p}}{\|\mathbf{e}_i\|_2 \|\mathbf{p}\|_2} \right), \mathbf{e}_i \in \mathbb{E}_A^{\text{id}}(v), \mathbf{p} \in \mathbb{P}_A, \quad (5)$$

where $\text{id} \in \{\text{node}, \text{struct}\}$, i ranges from 1 to l_a (or l_r), and \mathbb{P}_A is the atomic prototype set containing all A-prototypes. After that, we sort the AFEs in a descending order according to $\text{SimMAX}_i^{\text{id}}$ as $\text{AFE}_{\text{node}}^{\text{sort}} = \{\mathbf{A}_{m'} \in \mathbb{R}^{d_v \times d_a} | m' \in \{1, \dots, l_a\}\}$ and $\text{AFE}_{\text{struct}}^{\text{sort}} = \{\mathbf{R}_{n'} \in \mathbb{R}^{d_v \times d_r} | n' \in \{1, \dots, l_r\}\}$. Finally, we select the top l'_a and top l'_r ranked AFEs from these two sets as $\text{AFE}_{\text{node}}^{\text{select}}$ and $\text{AFE}_{\text{struct}}^{\text{select}}$, respectively. l'_a and l'_r are fixed hyperparameters with $l'_a \leq l_a$ and $l'_r \leq l_r$. Finally, the atomic embeddings generated by these selected AFEs are denoted as $\mathbb{E}_A^{\text{select}}(v)$, and we have $\mathbb{E}_A^{\text{select}}(v) \subset \mathbb{E}_A(v)$.

The $\mathbb{E}_A^{\text{select}}(v)$ is then used for refining existing prototypes or establishing new ones. To determine which embeddings contain observed features and which contain new features, a matching process is first conducted between the $\mathbb{E}_A^{\text{select}}(v)$ and \mathbb{P}_A . Formally, we measure the cosine similarity between elements in $\mathbb{E}_A^{\text{select}}(v)$ and elements in \mathbb{P}_A as

$$\text{Sim}_{E \rightarrow A}(v) = \left\{ \frac{\mathbf{e}_i^T \mathbf{p}}{\|\mathbf{e}_i\|_2 \|\mathbf{p}\|_2} \mid \mathbf{e}_i \in \mathbb{E}_A^{\text{select}}(v), \mathbf{p} \in \mathbb{P}_A \right\}. \quad (6)$$

Given $\text{Sim}_{E \rightarrow A}(v)$, the atomic embeddings with cosine similarity (not less than a certain threshold t_A) to at least one

existing A-prototype are regarded as containing observed features, since they are close to existing prototypes, *i.e.*,

$$\mathbb{E}_{\text{old}}(v) = \{ \mathbf{e}_i \mid \exists \mathbf{p} \in \mathbb{P}_A \text{ s.t. } \frac{\mathbf{e}_i^T \mathbf{p}}{\|\mathbf{e}_i\|_2 \|\mathbf{p}\|_2} \geq t_A \}. \quad (7)$$

$\mathbb{E}_{\text{old}}(v)$ is then used to refine \mathbb{P}_A . To this end, a distance loss \mathcal{L}_{dis} is computed to enhance the cosine similarity between each $\mathbf{e}_i \in \mathbb{E}_{\text{old}}(v)$ and its corresponding A-prototype $\mathbf{p}_i \in \mathbb{P}_A$, *i.e.*,

$$\mathcal{L}_{\text{dis}} = - \sum_{\mathbf{e}_i \in \mathbb{E}_{\text{old}}(v)} \frac{\mathbf{e}_i^T \mathbf{p}_i}{\|\mathbf{e}_i\|_2 \|\mathbf{p}_i\|_2} \quad (8)$$

By minimizing \mathcal{L}_{dis} , not only the existing A-prototypes in \mathbb{P}_A will get refined, the atomic embeddings will also be closer to 'standard' A-prototypes.

Next, we discuss how to deal with the atomic embeddings that are not close to any existing prototypes, *i.e.*, $\mathbb{E}_{\text{new}}(v) = \mathbb{E}_A^{\text{select}}(v) \setminus \mathbb{E}_{\text{old}}(v)$ or $\mathbb{E}_{\text{new}}(v) = \{ \mathbf{e}_i \mid \forall \mathbf{p} \in \mathbb{P}_A, \frac{\mathbf{e}_i^T \mathbf{p}}{\|\mathbf{e}_i\|_2 \|\mathbf{p}\|_2} < t_A \}$. Specifically, during the first task, \mathbb{P}_A is initialized with one random A-prototype. The matching process is still the same, and the only difference is that none of the atomic embeddings can be matched to the existing A-prototype. Therefore, all atomic embeddings may be candidates for creating new A-prototypes, as explained below.

Contrary to $\mathbb{E}_{\text{old}}(v)$, atomic embeddings in $\mathbb{E}_{\text{new}}(v)$ are regarded as new atomic features of the corresponding AFEs and should be accommodated with new prototypes. Considering that very similar embeddings may exist in $\mathbb{E}_{\text{new}}(v)$ and cause HPNs to create redundant prototypes, we first filter $\mathbb{E}_{\text{new}}(v)$ to obtain $\mathbb{E}'_{\text{new}}(v)$ which only contains the distinctive ones such that

$$\forall \mathbf{e}_i, \mathbf{e}_j \in \mathbb{E}'_{\text{new}}(v), \frac{\mathbf{e}_i^T \mathbf{e}_j}{\|\mathbf{e}_i\|_2 \|\mathbf{e}_j\|_2} < t_A. \quad (9)$$

In this way, the embeddings in $\mathbb{E}'_{\text{new}}(v)$ only contain new features, they are included into \mathbb{P}_A as new A-prototypes. Formally, the updated \mathbb{P}_A is the union of the previous A-prototypes \mathbb{P}_A and new A-prototypes $\mathbb{E}'_{\text{new}}(v)$

$$\mathbb{P}_A = \mathbb{P}_A \cup \mathbb{E}'_{\text{new}}(v). \quad (10)$$

After updating \mathbb{P}_A , $\text{Sim}_{E \rightarrow A}(v)$ will also be updated according to Eq. (6). Then, every element in $\text{Sim}_{E \rightarrow A}(v)$ would be not less than t_A , *i.e.*, each element in $\mathbb{E}_A^{\text{select}}(v)$ is matched to an A-prototype with cosine similarity larger than the threshold t_A . These matched A-prototypes are denoted as $\mathbb{A}(v)$, which represents the basic features of node v as shown in Figure 1.

As mentioned earlier, besides $\mathbb{A}(v)$, higher level representations are also beneficial for comprehensively describing a node from multiple granularities. Therefore, we further map $\mathbb{A}(v)$ to high level prototypes so as to obtain hierarchical prototype representations, Firstly, $\mathbb{A}(v)$ is mapped to an N-prototype to describe v as a whole. Specifically, the vectors in $\mathbb{A}(v)$ are concatenated as $\mathbf{a}_1 \oplus \dots \oplus \mathbf{a}'_{a+l'_r}$ ($\mathbf{a}_i \in \mathbb{A}(v)$), and then fed into a full connected layer $\text{FC}_{A \rightarrow N}(\cdot)$ to generate node-level embedding set $\mathbb{E}_N(v) = \{ \text{FC}_{A \rightarrow N}(\mathbf{a}_1 \oplus \dots \oplus \mathbf{a}'_{a+l'_r}) \}$. Here, for notation consistency with the aforementioned atomic embedding set, we also formulate the node-level embeddings as a set $\mathbb{E}_N(v)$ although it contains only

Algorithm 1: Learning Procedure for HPNs.

Input : Task sequence: $\{\mathcal{T}_1, \dots, \mathcal{T}_p\}$, HPNs
1 **for** $\mathcal{T} \leftarrow 1$ **to** p **do**
2 Get the data of the current task: \mathbb{V}, \mathbb{E} ,
 $\mathbf{X}(\mathbb{V}) = \{\mathbf{x}(v) | v \in \mathbb{V}\}$.
3 Select $\text{AFE}_{\text{node}}^{\text{select}}$ and $\text{AFE}_{\text{struct}}^{\text{select}}$.
4 Compute $\mathcal{L} = \text{HPNs}(\mathbb{V}, \mathbf{X}(\mathbb{V}), \mathbb{E})$.
5 Optimize \mathcal{L} .
Output: updated HPNs

one element. With $\mathbb{E}_N(v)$, the following procedure is to form a N-prototype representation $\mathbb{N}(v)$ for node v , which is same as the process to form $\mathbb{A}(v)$ at atomic level except that the threshold is set as t_N , instead of t_A . Finally, to enrich the node representation with the common features shared with similar nodes, we develop the C-prototypes. The generation of C-prototypes is similar to the N-prototypes. We first obtain a class-level embedding set $\mathbb{E}_C(v)$ by applying another fully connected layer $\text{FC}_{N \rightarrow C}$ to the element in $\mathbb{N}(v)$. Then, we match $\mathbb{E}_C(v)$ to existing C-prototypes or establish new ones. Different from N-prototypes, the threshold used here is t_C , which is typically larger than t_A and t_N . In this way, each C-prototype will be matched to embeddings within a large space, and therefore representing the common features shared by a wide range of similar nodes. Again, for notation convenience, we derive $\mathbb{P}_H(v)$ to represent the hierarchical prototype representations of the target node v (as shown in Figure 1), which is the union of $\mathbb{A}(v)$, $\mathbb{N}(v)$, and $\mathbb{C}(v)$:

$$\mathbb{P}_H(v) = \mathbb{A}(v) \cup \mathbb{N}(v) \cup \mathbb{C}(v). \quad (11)$$

3.4 Learning Objective

The obtained hierarchical prototypes of each node are first concatenated into a unified vector and then forward through a fully connected layer FC to obtain a c (the number of classes) dimensional feature vector, *i.e.*, $\text{FC}(\mathbf{h}_1 \oplus \dots \oplus \mathbf{h}_{l'_a+l'_r+2}), \forall \mathbf{h}_i \in \mathbb{P}_H(v)$. In this paper, we aim to perform node classification. Therefore, based on the c dimensional feature vector and the softmax function $\sigma(\cdot)$, we can estimate the label with $\hat{y}_c = \sigma(\text{FC}(\mathbf{h}_1 \oplus \dots \oplus \mathbf{h}_{l'_a+l'_r+2}))_c$ where c is the index of class. To perform node classification, with the output predictions \hat{y}_c and the target label $y_c \in \{1, 2, \dots, c\}$, the corresponding classification loss is given by

$$\mathcal{L}_{cls} = \sum_{c=1}^C -y_c \log(\hat{y}_c), \quad (12)$$

which is essentially the cross entropy loss function. As a discrepancy metric between probability distributions, cross entropy loss is generally preferred for classification tasks, because both the predictions $\{\hat{y}_c | c = 1, \dots, C\}$ and the labels are discrete probability distributions. Note that besides node classification, $\mathbb{P}_H(v)$ may also be used for other tasks based on different objective functions. In this paper, we focus on node classification and the overall loss of HPNs is:

$$\mathcal{L} = \mathcal{L}_{cls} + \alpha \cdot \mathcal{L}_{div} + \beta \cdot \mathcal{L}_{dis}, \quad (13)$$

where $\alpha \geq 0$ and $\beta \geq 0$ are hyperparameters to balance the two auxiliary losses.

3.5 Theoretical Analysis

In this subsection, we provide the theoretical upper bound on the memory consumption and analyze how the model configuration would affect HPNs' capacity in dealing with different tasks. Both theoretical results are justified and analyzed in the experiments. Only the main results are provided here, while the detailed proof and analysis are given in Section 5.

We first formulate the upper bound on the numbers of prototypes. Due to the mechanism to create new prototypes for newly emerging knowledge extracted from the graph, the memory consumption will gradually increase. However, the normalization applied to the prototypes constrains the prototype space, and thereby adds an upper bound on the memory consumption. This can be intuitively understood as the only limited number of points can be scattered on a n -dimensional hypersphere if we force the distance between each pair of points to be larger than a threshold. To formally formulate this, we first borrow the definition of spherical code from geometry and coding theory.

Definition 1 (Spherical code). *A spherical code $S(n, N, t)$, with parameters (n, N, t) is defined as the set of N points on the unit hypersphere in an n -dimension space for which the dot product of unit vectors from the origin to any two points is less than or equal to t .*

In HPNs, the prototypes at different levels are normalized into unit vectors and can be viewed as spherical codes in their own space. Taking the A-prototypes as an example, given the dimension d_a and the threshold t_A , \mathbb{P}_A is a spherical code $\mathbb{P}_A = S_A(d_a, N_A, 1 - t_A)$, where N_A is the cardinality of \mathbb{P}_A . Then the upper bound on the number of atom prototypes given d_a and t_A is equal to the maximal cardinality of $S(n, N, t)$ given n and t . Since the area of the n -dimensional unit sphere surface is limited, there exists a maximal N given a certain n and t , denoted as $\max_N S(d_a, N, 1 - t_A)$. However, finding $\max_N S(d_a, N, 1 - t_A)$ is a complex sphere packing problem, and there is not yet a general formulation of the maximal N for an arbitrary n . Therefore, given the number of two types of AFEs as l_a and l_r , we can formulate the upper bound on the numbers of different prototypes as:

Theorem 1 (Upper bounds on numbers of prototypes). *Given the notations defined in HPNs, the upper bound on the number of A-prototypes n_a can be given by*

$$n_a \leq (l_a + l_r) \max_N S(d_a, N, 1 - t_A), \quad (14)$$

and the upper bounds on the number of N-prototypes and the C-prototypes are:

$$n_N \leq \max_N S(d_n, N, 1 - t_N) \quad (15)$$

$$n_C \leq \max_N S(d_c, N, 1 - t_C) \quad (16)$$

where $S(n, N, t)$ is the spherical code defined on a n dimensional hypersphere.

Theorem 1 provides an upper bound on the memory consumption of HPNs. If the upper bound of the number of prototypes is reached, the model is still able to learn new tasks. This is because the ability to learn new tasks depends on new combinations of the AFEs. The budget (*i.e.*,

the limit) of the number of prototypes corresponds to the detail level of the basic properties extracted by AFEs. The larger budget is given to the number of prototypes, the more detailed representation can be generated for each node. In our experiments, we show that the number of parameters for most baseline methods are even higher than this upper bound.

Besides memory consumption, the more important problem for a continual learning model is the capability to maintain memory on previously learned tasks. Based on our model design, we formulate this as: whether learning new tasks affect the representations the model generates for old task data. We give explicit definitions on tasks and task distances based on set theory (in Section 5), then construct a bound to indicate what configuration would the model have to ensure this capability.

Theorem 2 (Task distance preserving). *For HPNs trained on consecutive tasks \mathcal{T}^p and \mathcal{T}^{p+1} . If $l_a d_a + l_r d_r \geq (l_r + 1) d_v$ and \mathbf{W} is column full rank, then as long as $t_A < \lambda_{\min}(l_r + 1) \text{dist}(\mathbb{V}_p, \mathbb{V}_{p+1})$, learning on \mathcal{T}^{p+1} will not modify representations HPNs generate for data from \mathcal{T}^p , i.e., catastrophic forgetting is avoided.*

The proof of Theorem 2 requires a number of definitions, lemmas, and corollaries, and the details are provided in Section 5. The theoretical results also provide insights into model implementations. In Theorem 2, λ_i is eigenvalues of the $\mathbf{W}^T \mathbf{W}$, where \mathbf{W} denotes the matrix constructed via AFEs (details in Section 5). d_v , d_a and d_r are dimensions of data and two kinds of atomic embeddings. The bound in this theorem is not tight, since the tight bound would be dependent on the specific dataset properties. But this informs us that either the number of AFEs or the dimension of the prototypes has to be large enough to ensure that data from two tasks can be well separated in the representation space. Then, according to Theorem 1, the upper bound of the memory consumption is dependent on $S(d_a, N, t_A)$, $S(d_n, N, t_N)$, and $S(d_c, N, t_C)$. As $S(n, N, t)$ grows fast with n , we prefer larger number of AFEs with smaller prototype dimensions. We also empirically verify this in Section 4.6. Besides, the upper bound proposed in Theorem 1 is explicitly computed and compared to experimental results.

4 EXPERIMENTS

In the experiments, we answer the following six questions: (1) Whether HPNs can outperform state-of-the-art approaches? (2) How does each component of HPNs contribute to its performance? (3) Whether HPNs can memorize previous tasks after learning each new task? (4) Are HPNs sensitive to the hyperparameters? (5) Whether the theoretical results can be empirically verified? (6) Whether the learned prototypes can be interpreted via visualization? All experiments are repeated on 5 different datasets. Due to space limit, only main results are presented in this paper, while the details over some datasets are included in Section 1.1 of the supplementary materials.

4.1 Datasets

To assess the effectiveness of the proposed HPNs, we consider 5 public datasets which include 3 citation networks

TABLE 1: The detailed statistics of 5 datasets used in our experiments.

Dataset	Cora	Citeseer	Actor	OGB-Arxiv	OGB-Products
# nodes	2,708	3,327	7,600	169,343	2,449,029
# edges	5,429	4,732	33,544	1,166,243	61,859,140
# features	1,433	3,703	931	128	100
# classes	7	6	4	40	47
# tasks	3	3	2	20	23

(Cora [13], Citeseer [13], OGB-Arxiv [14], [15]), 1 actor co-occurrence network (Actor) [57], and 1 product co-purchasing networks (OGB-Products [16]). The statistics of these 5 datasets are summarized in Table 1. Key properties of the datasets are presented here, and more details are included in Section 1.1 of the supplementary materials.

4.1.1 Citation networks

Cora contains 2708 documents, 5429 links denoting the citations among the documents. For training, 140 documents are selected with 20 examples for each class. The validation set contains 500 documents and the test set contains 1000 examples. In our continual learning setting, the first 6 classes are selected and grouped into 3 tasks (2 classes per task) in the original order. Citeseer contains 3312 documents and 4732 links. 20 documents per class are selected for training, 500 documents are selected for validation, and 1000 documents are selected as the test set. For continual learning setting, the documents from 6 classes are grouped into 3 tasks with 2 classes per task in the original order. The Cora and Citeseer datasets can be downloaded via Cora&Citeseer.

The OGB-Arxiv dataset is collected in the Open Graph Benchmark OGB. It is a directed citation network between all Computer Science (CS) arXiv papers indexed by MAG [14]. Totally it contains 169,343 nodes and 1,166,243 edges. The dataset contains 40 classes. As the dataset is imbalanced and the numbers of examples in different classes differs significantly, directly grouping the classes into 2-class groups like the Cora and Citeseer will cause the tasks to be imbalanced. Therefore, we reordered the classes in an descending order according to the number of examples contained in each class, and then group the classes according to the new order. In this way, the number of examples contained in different classes of each task are arranged to be as balanced as possible. Specifically, the class indices of each task are: $\{(35, 12), (15, 21), (28, 30), (16, 24), (10, 34), (8, 4), (5, 2), (27, 26), (36, 19), (23, 31), (9, 37), (13, 3), (20, 39), (22, 6), (38, 33), (25, 11), (18, 1), (14, 7), (0, 17), (29, 32)\}$, where each tuple denotes a task consisting of two classes.

4.1.2 Actor co-occurrence network

The actor co-occurrence network is a subgraph of the film-director-actor-writer network [58]. Each node in this dataset corresponds to an author, and the edges between the nodes are co-occurrence on the same Wikipedia pages. The whole dataset contains 7600 nodes and 33544 edges. Each node is accompanied with a feature vector of 931 dimensions. For this dataset, we also constructed 2 tasks with 2 classes per task. The link to this dataset is Actor. The balanced splitting of the classes is $\{(0, 1), (2, 3)\}$.

TABLE 2: Performance comparisons between HPNs and baselines on 5 different datasets.

C.L.T.	Base	Cora		Citeseer		Actor		OGB-Arxiv		OGB-Products	
		AM/%	FM/%	AM/%	FM/%	AM/%	FM /%	AM/%	FM /%	AM/%	FM /%
None	GCN	63.5±1.9	-42.3±0.4	64.5±3.9	-7.7±1.6	43.6±3.6	-9.1±2.9	56.8±4.3	-19.8±3.2	45.2±5.6	-27.8±7.1
	GAT	71.9±3.8	-33.1±2.3	66.8±0.9	-19.6±0.3	53.1±2.7	-4.3±1.6	54.3±3.5	-21.7±4.6	44.9±6.9	-30.3±5.2
	GIN	68.3±2.3	-35.4±3.4	57.7±2.3	-36.4±0.3	45.5±2.3	-8.9±2.8	53.2±6.5	-23.5±8.1	43.1±7.4	-31.4±8.8
EWC [24]	GCN	63.1±1.2	-42.7±1.6	54.4±4.2	-30.3±0.9	44.3±1.1	-7.1±1.4	72.1±2.4	-9.1±1.9	66.7±0.5	-8.4±0.4
	GAT	72.2±1.5	-32.2±1.6	65.7±2.5	-19.7±2.3	54.2±2.5	-2.5±1.5	73.2±1.1	-10.8±2.1	67.9±1.0	-9.6±1.3
	GIN	69.6±2.6	-28.5±2.8	57.9±3.4	-36.3±2.4	47.6±2.1	-7.2±1.6	74.1±1.7	-8.3±2.0	67.3±2.3	-13.6±1.5
LwF [23]	GCN	76.1±1.4	-21.3±2.4	67.0±0.2	-8.3±2.7	49.7±2.5	-3.6±1.4	69.9±3.9	-12.1±2.8	66.3±2.5	-11.8±3.4
	GAT	70.8±2.8	-34.6±4.1	66.1±4.1	-18.9±1.5	52.8±2.7	-6.2±2.2	68.9±4.4	-13.6±3.3	65.1±4.1	-13.2±2.9
	GIN	74.1±2.7	-23.3±0.8	63.1±1.9	-16.5±2.2	49.7±2.6	-4.1±1.5	71.4±4.8	-15.9±5.6	65.9±4.0	-10.7±3.1
GEM [27]	GCN	75.7±3.0	-6.5±4.4	41.8±2.6	-31.9±1.4	52.7±3.1	+3.9±2.9	75.4±1.7	-13.6±0.5	71.3±1.7	-10.5±0.9
	GAT	69.8±3.0	-26.1±2.6	71.3±2.2	+9.0±1.5	54.3±3.6	-2.0±0.9	76.6±0.7	-11.3±0.4	70.4±0.8	-10.9±1.6
	GIN	80.2±3.3	-2.0±4.2	49.7±0.5	-24.5±0.9	45.2±2.8	-11.1±1.5	77.3±2.1	-11.2±1.6	76.5±3.3	-7.2±2.5
MAS [59]	GCN	65.5±1.9	-21.4±3.7	59.5±3.1	-0.1±2.4	50.7±2.4	-1.5±0.8	69.8±0.4	-18.8±0.9	62.0±1.1	-17.9±1.9
	GAT	84.7±0.7	-5.6±2.0	69.1±1.1	-4.8±3.3	53.7±2.6	-1.6±0.8	70.6±1.3	-16.7±1.6	64.4±2.3	-14.5±3.2
	GIN	76.7±2.6	-4.0±3.6	65.2±3.9	+0.0±2.0	51.6±2.6	-0.6±1.3	65.3±2.9	-17.0±2.3	61.4±3.8	-20.9±2.9
ERGN. [60]	GCN	63.5±2.4	-42.3±0.7	54.2±3.9	-30.3±1.9	52.4±3.3	+0.6±1.4	63.3±1.7	-18.1±0.9	60.7±2.8	-26.6±3.3
	GAT	71.1±2.5	-34.3±1.0	65.5±0.3	-20.4±3.9	51.4±2.2	-7.2±3.2	63.5±2.4	-19.5±1.9	61.3±1.7	-25.1±0.8
	GIN	68.3±0.4	-35.4±0.4	57.7±3.1	-36.4±1.3	42.7±3.9	-13.0±2.1	69.2±1.8	-11.8±1.4	61.8±4.7	-23.4±7.9
TWP [18]	GCN	68.9±0.9	-5.7±1.5	60.5±3.8	-0.3±4.4	50.6±2.0	-4.8±1.1	75.6±0.3	-10.4±0.5	69.9±0.4	-9.0±1.1
	GAT	81.3±3.2	-14.4±1.5	69.8±1.5	-8.9±2.6	54.0±1.8	-2.1±1.9	75.8±0.5	-5.9±0.3	69.3±2.3	-8.9±1.5
	GIN	73.7±3.2	-3.9±2.6	68.9±0.7	-2.4±1.9	49.9±1.9	-3.6±2.0	76.6±1.8	-11.3±1.1	69.9±1.4	-10.3±2.7
Join.	GCN	93.7± 0.5	-	78.9±0.4	-	57.0±0.9	-	77.2±0.8	-	72.9±1.2	-
	GAT	93.9± 0.9	-	79.3±0.8	-	57.1±0.9	-	81.8±0.3	-	73.7±2.4	-
	GIN	93.2± 1.2	-	78.7±0.9	-	56.9±0.6	-	82.3±1.9	-	77.9±2.1	-
HPNs		93.7±1.5	+0.6±1.0	79.0±0.9	-0.6±0.7	56.8±1.4	-0.9±0.9	85.8± 0.7	+0.6±0.9	80.1±0.8	+2.9±1.0

4.1.3 Product co-purchasing network

OGB-Products is collected in the Open Graph Benchmark OGB, representing an Amazon product co-purchasing network link. It contains 2,449,029 nodes and 61,859,140 edges. Nodes represent products sold in Amazon, and edges indicate that the connected products are purchased together. In our experiments, we select 46 classes and omit the last class containing only 1 example. Similar to OGB-Arxiv, we reorder the classes and the class indices of each tasks are: {(4, 7), (6, 3), (12, 2), (0, 8), (1, 13), (16, 21), (9, 10), (18, 24), (17, 5), (11, 42), (15, 20), (19, 23), (14, 25), (28, 29), (43, 22), (36, 44), (26, 37), (32, 31), (30, 27), (34, 38), (41, 35), (39, 33), (45, 40)}.

4.2 Experimental Setup and Evaluation Metrics

To perform continual graph representation learning with new categories of nodes continuously emerging, like the traditional continual learning works, we split each dataset into multiple tasks and train the model on them sequentially. Each new task brings a subgraph with new categories of nodes and associated edges, *e.g.*, task 1 contains classes 1 and 2, task 2 contains new classes 3 and 4, *etc.* Each model is trained on a sequence of tasks, and the performance will be evaluated on all previous tasks. Specifically, we adopt accuracy mean (AM) and forgetting mean (FM) as metrics for evaluation. After learning on all tasks, the AM and FM are computed as the average accuracy and the average accuracy decrease on all previous tasks. Given a sequence of T tasks, we denote the model's accuracy on the q -th task after learning the p -th task as $M_{p,q}$, then $\{M_{T,q}|q = 1, \dots, T\}$ are the model' accuracy on each learnt task after learning the entire sequence. Accordingly, the mathematical formulations

of AM and FM are:

$$AM = \frac{\sum_{q=1}^T M_{T,q}}{T}, \quad (17)$$

$$FM = \frac{\sum_{q=1}^{T-1} (M_{T,q} - M_{q,q})}{T-1}. \quad (18)$$

Negative FM indicates the existence of forgetting, zero FM denotes no forgetting and positive FM denotes positive knowledge transfer between tasks. For HPNs, we set $d_a = d_n = d_c = 16$, $l_a = l_r = 22$, and $h = 2$. The threshold t_A , t_N , and t_C are selected by cross validation on $\{0.01, 0.05, 0.1, 0.15, 0.2, 0.25, 0.3, 0.35, 0.4\}$. The experiments on the important hyperparameters are provided in Section 4.6. All experiments are repeated 5 times with Nvidia Titan Xp GPU. The average performance and standard deviations are reported. More details are provided in Section 1.2 of the supplementary materials and our code.

4.3 Comparisons with Baseline Methods

We compare HPNs with various baseline methods. Experience Replay based GNN (ERGN) [17] and Topology-aware Weight Preserving (TWP) [18] are developed for continual graph representation learning. The others approaches, including Elastic Weight Consolidation (EWC) [24], Learning without Forgetting (LwF) [23], Gradient Episodic Memory (GEM) [27], and Memory Aware Synapses (MAS) [59]) are popular continual learning methods for Euclidean data. All the baselines are implemented based on three popular backbone models, *i.e.*, Graph Convolutional Networks (GCNs) [4], Graph Attentional Networks (GATs) [6], and Graph Isomorphism Network (GIN) [7].

Note that Joint training (Join.) in Table 2 does not represent continual learning. It allows a model to access

		GAT			HPNs		
Te.	Tr.	\mathcal{T}_1	$\mathcal{T}_{1,2}$	$\mathcal{T}_{1,2,3}$	\mathcal{T}_1	$\mathcal{T}_{1,2}$	$\mathcal{T}_{1,2,3}$
		\mathcal{T}_1		94.12%	47.96%	71.49%	95.02%
\mathcal{T}_2			93.30%	49.68%		92.66%	92.22%
\mathcal{T}_3				94.44%			96.83%

TABLE 3: Accuracy (%) changes of GAT and HPNs.

		GAT+GEM			GAT+MAS		
Te.	Tr.	\mathcal{T}_1	$\mathcal{T}_{1,2}$	$\mathcal{T}_{1,2,3}$	\mathcal{T}_1	$\mathcal{T}_{1,2}$	$\mathcal{T}_{1,2,3}$
		\mathcal{T}_1		93.21%	91.40%	66.97%	94.12%
\mathcal{T}_2			81.21%	81.43%		77.97%	70.84%
\mathcal{T}_3				63.89%			93.25%

TABLE 4: Accuracy (%) changes of GAT+GEM and GAT+MAS.

data of all tasks at any time and thus is often used as an upper bound for continual learning. [61]. Therefore, FM is not applicable to Joint training approach.

In Table 2, we observe that regularization based approaches, *e.g.*, EWC and TWP, generally obtain lower forgetting, but the accuracy (AM) is limited by the constraints. However, the forgetting problem of regularization based methods will become increasingly severe when the number of tasks is relatively large, as shown in Section 4.5. Memory replay based methods such as GEM achieve better performance without using any constraint. However, the memory consumption is higher (Section 4.7). HPNs significantly outperform all baselines without inheriting their limitations. Compared to regularization based methods, HPNs do not impose constraints to limit the model’s expressiveness, therefore the performance is much better. Compared to memory replay based methods, HPNs do not only perform better but also are memory efficient as shown in Section 4.7. Joint training (Join.) achieves comparable performance to HPNs on small datasets but is significantly worse on large OGB datasets. This is because joint training (Join.) is a multi-task setting, inter-task interference may cause negative transfer, which is not obvious on small datasets with only a few tasks but becomes prominent on large datasets with tens of tasks. In HPNs, different tasks can choose different combinations of the parameters and thus task interference is dramatically alleviated.

Among all the models in Table 2, we could observe several kinds of forgetting behavior among the baselines. First, some of them are with low AM and severe forgetting (negative FM with large |FM|) like the pure GNNs in the first 3 rows of Table 2 in the paper. These models may perform well on individual tasks, but the AM is brought down by

TABLE 5: Ablation study on prototypes of different levels of prototypes over Cora.

Conf.	A-p.	N-p.	C-p.	AM%	FM%
1	✓			89.2±1.3	-0.1±0.5
2	✓	✓		91.7±1.1	-0.2±0.8
3	✓	✓	✓	93.7±1.5	+0.6±1.0

TABLE 6: Ablation study on different loss terms over Cora.

Conf.	\mathcal{L}_{cls}	\mathcal{L}_{div}	\mathcal{L}_{dis}	AM%	FM%
1	✓			92.4±1.3	+0.8±0.7
2	✓	✓		92.9±1.1	+0.3±1.0
3	✓		✓	92.8±0.9	+0.0±1.2
4	✓	✓	✓	93.7±1.5	+0.6±1.0

catastrophic forgetting. To explain this in details, we expand the results of GAT without any continual learning techniques in Table 3, in which Tr. denotes on which tasks has the model been trained and Te. denotes on which task the model is tested.

The first row of Table 3 shows the model performance on task 1 after it has been trained on the task 1, on the first two tasks, and on the first three tasks. We see that GAT performs well on each individual task it has just learnt. But after being trained on more tasks, the performance on previous tasks drops dramatically, making the AM to be relatively low. In contrast, the performance change on previous tasks of HPNs is also shown in Table 3, and we could see that HPNs can well maintain the performance on previous tasks throughout the training on all tasks.

Second, there are also some models without severe forgetting but the still low AM. Typically these are the models which preserve the performance on previous tasks with certain constraints on the models. These constraints can indeed alleviate forgetting, but it also limits the model’s flexibility in learning new tasks, and thus the performance on new tasks will degrade. For example, comparing Table 3 with Table 4, GAT achieves accuracy higher than 93% on individual tasks, but for GAT+MAS, although the forgetting on task 1 is alleviated, the performance on task 2 drops significantly. GAT+GEM also suffers from the similar problem.

4.4 Ablation Study

We conduct ablation studies on different levels of prototypes and different combinations of three loss terms. In Table 5, we show the performance of HPNs when A-, N-, and C-Prototypes are gradually added (Cora dataset). We notice both AM and FM of HPNs increase when higher level prototypes are considered. This suggests that high level prototypes can enhance the model’s performance and robustness against forgetting. The effect of different combinations of loss terms are shown in Table 6. The first three rows show that adding \mathcal{L}_{div} or \mathcal{L}_{dis} with \mathcal{L}_{cls} may slightly improve the performance. By jointly considering these three terms, the performance (AM) can be further improved. This is because \mathcal{L}_{div} pushes different AFEs away from each other and \mathcal{L}_{dis} makes the prototypes of each AFE be more close to its output. Jointly considering \mathcal{L}_{div} and \mathcal{L}_{dis} with \mathcal{L}_{cls} can make the prototype space better separated as shown in Section 4.8.

4.5 Learning Dynamics

For continual learning, it is important to memorize previous tasks after learning each new task. To measure this, instead of directly measuring the average accuracy on previous tasks which may mix up the accuracy change caused by forgetting and task differences, we develop a new metric, *i.e.*, average




Fig. 2: Left: dynamics of ARS for continual learning tasks on OGB-Arxiv. Middle: impact of t_A on the number of prototypes in HPNs over Cora. Right: dynamics of memory consumption of HPNs on OGB-Products.




Fig. 3: (a) and (b) are AM and FM of HPNs with different number of AFEs and prototype dimensions on OGB-Arxiv. (c) and (d) are AM and FM change with when t_A varies on Cora.




Fig. 4: The left two figures show the results of fixing $\alpha = 1.0$ and tuning β from 0.0001 to 1000. The right two figures show the results of fixing $\beta = 1.0$ and tuning α . Logarithmic horizontal axis is adopted, and the results are obtained on both Cora and Citeseer datasets.

retaining score (ARS), to address this problem. Specifically, after learning on a task \mathcal{T}^i , the ratio between the model’s accuracy on a previous task \mathcal{T}^{i-m} and its accuracy on \mathcal{T}^{i-m} after it had been just learned on \mathcal{T}^{i-m} is defined as the retaining ratio. Then the ARS is the average retaining ratio of all previous tasks after learning a new task.

Figure 2(left) shows the ARS change of HPNs and two baselines. GAT represents the models without continual learning techniques. TWP+GAT is the best baseline in terms of forgetting. GAT forgets quickly, while TWP significantly alleviates the forgetting problem for GAT. But as more tasks come in, the forgetting of TWP+GAT increases. As different tasks require different parameters, TWP+GAT (regularization based) is seeking a trade-off between old and new tasks. With more new tasks, TWP+GAT tends to gradually adapt to new tasks and forget old ones. On contrary, HPNs maintain the ARS very well. This is because HPNs learn prototypes to denote the common basic features and learning new tasks does not hurt the parameters for old tasks. New tasks can be handled with new combinations of the existing basic

prototypes. If necessary, new prototypes can be established for more expressiveness.

4.6 Parameter Sensitivity

As discussed in Section 3.5, the number of AFEs and the prototype dimensions are key factors in determining the continual learning capability and memory consumption. Here, we conduct experiments with different number of AFEs and prototype dimensions to justify the theoretical results. We keep the dimensions of different prototypes equal and the number of two types of AFEs equal for simplicity.

As shown in Figure 3(a) and (b), larger dimensions and the number of AFEs yield better AM and FM, which is consistent with Theorem 2. Besides, AM is mostly determined by the number of AFEs since HPNs compose prototypes with different AFEs to represent each target node. The number of possible combinations determines its expressiveness. Considering the above results and the bound (Theorem 1) for the number of prototypes, using large number of AFEs

TABLE 7: Final parameter amount for models trained on OGB-Products

	None	EWC	LwF	GEM	MAS	ERGNN	TWP	Joint	HPNs
GCN	2,336	46,720	4,672	2,202,336	2,336	6,738	9,344	2,336	4,908
GAT	20,032	400,640	40,064	2,220,032	20,032	24,432	80,128	20,032	
GIN	2,352	47,040	4,704	2,202,352	2,352	6,752	9,408	2,352	




Fig. 5: Visualization of hierarchical prototype representations of nodes in the test set of Cora.

and small dimension can ensure both high performance and low memory usage, as verified in Section 4.7.

We also evaluate the effectiveness of HPNs when prototype thresholds vary from 0.01 to 0.4. In Figure 3(c) and (d), we observe that the performance (AM and FM) of HPNs are generally stable when t_A varies and slightly better when t_A is between 0.2 and 0.3. This is because when t_A is too small or too large, we will have too many or too less prototypes (consistent with Theorem 1) as shown in Figure 2(middle), which may cause the problem of overfitting or underfitting.

Finally, we explore the influence of the hyperparameters of the auxiliary losses (α and β in Equation 13) on the performance. As shown in Figure 4, we fix $\alpha = 1.0$ or $\beta = 1.0$ in turn, and tune the other one to check the parameter sensitivity.

From Figure 4, we can observe that the performance is relatively robust when α and β vary from 0.0001 to 1000. Through monitoring the training process, we suggest the reason is that \mathcal{L}_{div} and \mathcal{L}_{dis} decrease faster than \mathcal{L}_{cls} , thus the total loss is not very sensitive to the coefficients on them. Based on this observation, we choose to simply set $\alpha = \beta = 1$ in implementations.

In this subsection, we systematically studied the dependency of the model on the hyperparameters. According to the results, HPNs are robust against most of the hyperparameters. The number of AFEs will cause significant difference but the relation between the performance and the number of AFEs is simple. For a given dataset, the performance first increases with the number of AFEs, and becomes stable when the number AFEs reaches a certain threshold (*i.e.* large enough).

4.7 Memory Consumption

We compare memory consumption of different methods, as well as an explicitly theoretical memory upper bound, with the baselines on OGB-Products (the largest dataset). We also show the actual memory consumption of HPNs in the process of continual learning. We measure the total number of parameters in our model including not only the prototypes but also the parameters of the other modules like the AFEs. The number of prototypes is converted into the number of parameters. For example, an 8-dimensional

prototype is counted as 8 parameters, since the prototypes and the other parameters use the same data type with the same precision. Similarly, for other models in Table 7, we also show the total number of parameters. For the constant memory models like EWC, LwF, etc., only the trainable parameters are counted. For the memory based methods like GEM, the space consumption of the memory buffer is also counted.

In Table 7, even on the dataset with millions of nodes and 23 tasks, HPNs can accommodate all tasks with a small amount of parameters. Besides, the dynamic change of parameter amount is shown in Figure 2(right). The red dashed line denotes the theoretical upper bound (6,163), and the computation details are included in supplemental materials. In Figure 2(right), we notice the actual memory usage of HPNs is much lower than the upper bound. Moreover, even the upper bound is among the lowest for memory consumption compared to baselines. The model we use here is the same as the one in Section 4.3

4.8 Visualization

To concretely show the prototype representations generated by HPNs, we apply t-SNE [62] to visualize the node representations of the Cora dataset (test set) after learning each task, as shown in Figure 5. The three figures display the representations generated after each of the three tasks arrives sequentially. Each task contains two classes denoted with different colors, *i.e.* (red, blue), (green, salmon), and (purple, orange). In Figure 5, with new tasks coming in sequentially from left to right, the representations are consistently well separated, which is beneficial for downstream tasks.

5 PROOF DETAILS

5.1 Overview

The main theoretical results are briefly introduced in the paper. In this section, we provide detailed explanations and proofs for the theoretical results.

5.2 Memory consumption upper bound

The proof and detailed analysis on the memory consumption upper bound has already been given in the paper. In this

subsection, we give the computation of a specific case of the general theoretical results.

Although the general formulation of the upper bound is not available, we can specially compute $\max_N S(d_a, N, 1 - t_A)$ for certain n_s , and verify it with experiments. For example, when $n = 2$, the distribution becomes distributing points on a circle with unit radius. Then, $\max_N S(d_a, N, 1 - t_A)$ can be obtained by evenly distributing the points on the circle with an interval of t_A . Finally, the explicit value of $\max_N S(d_a, N, 1 - t_A)$ can be formulated as:

$$\max_N S(d_n, N, 1 - t_N) = \frac{2\pi}{\arccos(1 - t_A)}, \quad (19)$$

then we have:

$$n_A \leq (l_a + l_r) \frac{2\pi}{\arccos(1 - t_A)}. \quad (20)$$

And the upper bound of the number of N- and C- prototypes can be formulated similarly. The above results are used in Section 4.7 in the paper.

5.3 Task distance preserving

In this subsection, we give proofs and detailed analysis on the Theorem 2 in the paper. At below, Lemma 2, Lemma 3, Lemma 4, and Corollary 1 are from existing knowledge ranging from geometry to linear algebra. The other parts are of our own contributions.

In continual learning, the key challenge is to overcome the catastrophic forgetting, which refers to the performance degradation on previous tasks after training the model on new tasks. Based on our model design, we formulate this as: whether learning new tasks affects the representations the model generates for old task data. First, we give definitions on the tasks and task distances:

Definition 2 (Task set). *The p -th task in a sequence is denoted as \mathcal{T}^p and contains a subgraph \mathcal{G}_p consisting of nodes belonging to some new categories. We denote the associated node set and adjacency matrix as \mathbb{V}_p and A_p . Each $v_p^i \in \mathbb{V}_p$ has a feature vector $\mathbf{x}(v_p^i)$ and a label $y(v_p^i)$.*

Then, the reason for catastrophic forgetting is that different tasks in a sequence are drawn from heterogeneous distributions, making the model sequentially trained on different tasks unable to maintain satisfying performances on previous tasks. Therefore, given the definition of the tasks (Definition 2), we then give a formal definition to quantify the difference between two tasks.

Definition 3 (Task distance). *We define the distance between two tasks as the set distance between the node sets of these two tasks, i.e.*

$$\mathbf{dist}(\mathbb{V}_p, \mathbb{V}_q) = \inf \|\mathbf{x}(v_p^i) - \mathbf{x}(v_q^j)\|, \forall v_p^i \in \mathbb{V}_p, v_q^j \in \mathbb{V}_q.$$

Lemma 1. *The distance between any two tasks is non-negative, i.e. $\forall i, j \in \{1, \dots, M^T\}$, $\mathbf{dist}(\mathbb{V}_p, \mathbb{V}_q) \geq 0$, where M^T is the number of tasks contained in the sequence.*

The real-world data could be complex and sometimes may even contain noises that are impossible for any model to learn, which needs extra considerations when justifying the effectiveness of the model. Formally, we give the definition of the contradictory data.

Definition 4 (Contradictory data). *$\forall v_p^i \in \mathbb{V}_p, p = 1, \dots, M^T$, if $\exists v_q^j \in \mathbb{V}_q, j = 1, \dots, M^T$, $\text{st.} \forall l \in \mathbb{N}^*$, $\forall u \in \mathcal{N}^l(v_p^i)$ and $\forall v \in \mathcal{N}^l(v_q^j)$, $\mathbf{x}(u) = \mathbf{x}(v)$ but $y(v_p^i) \neq y(v_q^j)$, then we say $(v_p^i, y(v_p^i))$ and $(v_q^j, y(v_q^j))$ are contradictory data, as it is contradictory for any model to give different predictions for one node based on the same node features and graph structures. (\mathbb{N}^* denotes the set of non-negative integers)*

Remark 1. *Contradictory data is ignored or simply regarded as outliers in previous works, but in this work, we explicitly analyze its affect for the comprehensiveness of our theory. contradictory data has different situations. If v_p^i and v_q^j are from different tasks, then $y(v_p^i) \neq y(v_q^j)$ is plausible. Because they may be describing a same thing from different aspects. For example, an article from the citation network may be both categorized as 'physics related' and 'computer science related'. In this situation, it would be easy to add an task indicator to the feature of the node, then the feature of v_p^i and v_q^j are no longer equal and are not contradictory data anymore.*

But within one task, contradictory data are most likely to be wrongly labeled, e.g. it does not make sense if an article is both 'related to physics' and 'not related to physics'.

Besides the distance between tasks, the distance between the embeddings obtained by the AFEs will also be a crucial concept in the proof.

Definition 5 (Embedding distance). *Each input node v_p^i is given a set of atomic embeddings $\mathbb{E}_A(v_p^i) = \mathbb{E}_A^{\text{node}}(v_p^i) \cup \mathbb{E}_A^{\text{struct}}(v_p^i)$, where $\mathbb{E}_A^{\text{node}}(v_p^i) = \{\mathbf{a}_i^j | j \in \{1, \dots, l_a\}\}_p$ containing the atomic node embeddings of v_p^i and $\mathbb{E}_A^{\text{struct}}(v_p^i) = \{\mathbf{r}_i^j | k \in \{1, \dots, l_r\}\}_p$ containing the atomic structure embeddings. $\mathbf{a}_i^j \in \mathbb{R}^{d_a}$ and $\mathbf{r}_i^j \in \mathbb{R}^{d_r}$. To define the distance between representations of two nodes, we concatenate the atomic embeddings of each node into a single vector in a higher dimensional space, i.e. each node v_p^i corresponds to a latent vector $\mathbf{z}_p^i = [\mathbf{a}_i^1; \dots; \mathbf{a}_i^{l_a}; \mathbf{r}_i^1; \dots; \mathbf{r}_i^{l_r}] \in \mathbb{R}^{l_a \times d_a + l_r \times d_r}$. Then we define the distance between representations of two nodes v_p^i and v_q^j as the Euclidean distance between their corresponding latent vector \mathbf{z}_p^i and \mathbf{z}_q^j , i.e. $\mathbf{dist}(\mathbf{z}_p^i, \mathbf{z}_q^j) = \|\mathbf{z}_p^i - \mathbf{z}_q^j\|_2$.*

Then we will give some explanations on the linear algebra related theories.

Lemma 2 (Bounds for real quadratic forms). *Given a real symmetric matrix \mathbf{A} , and an arbitrary real vector variable \mathbf{x} , we have*

$$\lambda_{\min} \leq \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \leq \lambda_{\max}, \quad (21)$$

where λ_{\min} and λ_{\max} are the minimum and maximum eigenvalues of matrix \mathbf{A} .

Lemma 3 (Real symmetric matrix). *For a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{A}^T \mathbf{A} \in \mathbb{R}^{n \times n}$ is a real symmetric matrix, $\text{rank}(\mathbf{A}^T \mathbf{A}) = \text{rank}(\mathbf{A})$, and the non-zero eigenvalues of $\mathbf{A}^T \mathbf{A}$ are squares of the non-zero singular values of \mathbf{A} .*

Lemma 4 (Rank and number of non-zero singular values). *For a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, the number of non-zero singular values equals the rank of \mathbf{A} , i.e. $\text{rank}(\mathbf{A})$*

Corollary 1. *For a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$. Without loss of generality, we assume $n \leq m$. If \mathbf{A} is column full rank, i.e. $\text{rank}(\mathbf{A}) = n$,*

then \mathbf{A} has n non-zero singular values. Besides, $\text{rank}(\mathbf{A}^T \mathbf{A}) = n$, and $\mathbf{A}^T \mathbf{A}$ has n non-zero singular values.

Given the explanations above, we then derive the bound for the change of the distance among data, which will be further used for analyzing the separation of data from different tasks.

Lemma 5 (Embedding distance bound). *Given two nodes $v_p^i \in \mathbb{V}_p$ and $v_q^j \in \mathbb{V}_q$ with vertex feature $\mathbf{x}(v_p^i), \mathbf{x}(v_q^j) \in \mathbb{R}^{d_v}$, their multi-hop neighboring node sets are denoted as $\bigcup_{l \in \mathbb{N}^*} \mathcal{N}^l(v_p^i)$ and $\bigcup_{l \in \mathbb{N}^*} \mathcal{N}^l(v_q^j)$. The AFEs for generating atomic embeddings are*

$\text{AFE}_{\text{node}} = \{\mathbf{A}_i \in \mathbb{R}^{d_a \times d_v} | i \in \{1, \dots, l_a\} \text{ and } \text{AFE}_{\text{struct}} = \{\mathbf{R}_j \in \mathbb{R}^{d_r \times d_v} | j \in \{1, \dots, l_r\}\}$, corresponding to matrices for atomic node embeddings and atomic structure embeddings, respectively. Then, the square distance

$$\text{dist}^2(\mathbf{z}_p^i - \mathbf{z}_q^j) = \|\mathbf{z}_p^i - \mathbf{z}_q^j\|_2^2 \geq \lambda_{\min}(\|\mathbf{x}(v_p^i) - \mathbf{x}(v_q^j)\|_2^2 + \sum_{k=1}^{l_r} \|\mathbf{x}(u_k) - \mathbf{x}(\nu_k)\|_2^2),$$

if $l_a \times d_a + l_r \times d_r \geq d_v$, where u_k are nodes sampled from $\bigcup_{l \in \mathbb{N}^*} \mathcal{N}^l(v_p^i)$, ν_k are nodes sampled from $\bigcup_{l \in \mathbb{N}^*} \mathcal{N}^l(v_q^j)$, λ_i

are the eigenvalues of $\mathbf{W}^T \mathbf{W}$, and $\mathbf{W} \in \mathbb{R}^{(l_r+1)d_v \times (l_a d_a + l_r d_r)}$ is constructed with the matrices in AFE_{node} and $\text{AFE}_{\text{struct}}$. Specifically, \mathbf{W} is a block matrix constructed as follows:

1. $\mathbf{W}_{1:l_a d_a, 1:d_v}$ are filled by the concatenation of $\{\mathbf{A}_i | i = 1, \dots, l_a\}$, i.e. $[\mathbf{A}_1; \dots; \mathbf{A}_{l_a}] \in \mathbb{R}^{l_a d_a \times d_v}$.

2. For $\mathbf{W}_{l_a d_a + 1:l_a d_a + l_r d_r, 1:(l_r+1)d_v}$, the construction is first filling $\mathbf{W}_{l_a d_a + (k-1)d_r : l_a d_a + kd_r, kd_v : (k+1)d_v}$ with \mathbf{R}_k , $k = 1, \dots, l_r$.

3. For other parts, fill with zeros.

Proof. Given vertex v_p^i , we concatenate its feature vector with the l_r neighbors sampled from $\bigcup_{l \in \mathbb{N}^*} \mathcal{N}^l(v_p^i)$, i.e. $\mathbf{x}'_{p,i} = [\mathbf{x}(v_p^i); \mathbf{x}(u_1); \dots; \mathbf{x}(u_{l_r})] \in \mathbb{R}^{(l_r+1)d_v \times 1}$, $u_j \in \bigcup_{l \in \mathbb{N}^*} \mathcal{N}^l(v_p^i)$.

Then with the constructed block matrix \mathbf{W} , we could formulate the generation of \mathbf{z}_p^i as: $\mathbf{z}_p^i = \mathbf{W} \mathbf{x}'_{p,i}$.

Similarly, we can formulate \mathbf{z}_q^j for another vertex v_q^j . And their distance can be formulated as:

$$\text{dist}(\mathbf{z}_p^i, \mathbf{z}_q^j) = \|\mathbf{z}_p^i - \mathbf{z}_q^j\|_2 = \sqrt{(\mathbf{z}_p^i - \mathbf{z}_q^j)^T (\mathbf{z}_p^i - \mathbf{z}_q^j)},$$

where $(\mathbf{z}_p^i - \mathbf{z}_q^j)^T (\mathbf{z}_p^i - \mathbf{z}_q^j)$ can be further expanded as:

$$\begin{aligned} (\mathbf{z}_p^i - \mathbf{z}_q^j)^T (\mathbf{z}_p^i - \mathbf{z}_q^j) &= (\mathbf{W} \mathbf{x}'_{p,i} - \mathbf{W} \mathbf{x}'_{q,j})^T (\mathbf{W} \mathbf{x}'_{p,i} - \mathbf{W} \mathbf{x}'_{q,j}) \\ &= (\mathbf{W} (\mathbf{x}'_{p,i} - \mathbf{x}'_{q,j}))^T (\mathbf{W} (\mathbf{x}'_{p,i} - \mathbf{x}'_{q,j})) \\ &= (\mathbf{x}'_{p,i} - \mathbf{x}'_{q,j})^T \mathbf{W}^T \mathbf{W} (\mathbf{x}'_{p,i} - \mathbf{x}'_{q,j}) \end{aligned}$$

According to lemma 3, $\mathbf{W}^T \mathbf{W}$ is a real symmetric matrix, with lemma 2, we have:

$$\frac{(\mathbf{x}'_{p,i} - \mathbf{x}'_{q,j})^T \mathbf{W}^T \mathbf{W} (\mathbf{x}'_{p,i} - \mathbf{x}'_{q,j})}{(\mathbf{x}'_{p,i} - \mathbf{x}'_{q,j})^T (\mathbf{x}'_{p,i} - \mathbf{x}'_{q,j})} \geq \lambda_{\min}$$

According to lemma 3, with $l_a d_a + l_r d_r \geq (l_r + 1)d_v$ and the constraint of column full rank on \mathbf{W} , $\mathbf{W}^T \mathbf{W} \in$

$\mathbb{R}^{(l_r+1) \times (l_r+1)}$ has $l_r + 1$ positive eigenvalues, thus $\lambda_{\min} > 0$. Then we decompose $(\mathbf{x}'_{p,i} - \mathbf{x}'_{q,j})^T (\mathbf{x}'_{p,i} - \mathbf{x}'_{q,j})$ as:

$$\begin{aligned} & (\mathbf{x}'_{p,i} - \mathbf{x}'_{q,j})^T (\mathbf{x}'_{p,i} - \mathbf{x}'_{q,j}) \\ &= \sum_{k=0}^{(l_r+1)d_v-1} ((\mathbf{x}'_{p,i})_k - (\mathbf{x}'_{q,j})_k)^2 \\ &= \sum_{k=1}^{(l_r+1)d_v} ((\mathbf{x}'_{p,i})_k - (\mathbf{x}'_{q,j})_k)^2 \\ &= \sum_{k=1}^{d_v} ((\mathbf{x}'_{p,i})_k - (\mathbf{x}'_{q,j})_k)^2 + \sum_{k=1}^{l_r} \sum_{m=kd_v+1}^{(k+1)d_v} ((\mathbf{x}'_{p,i})_k - (\mathbf{x}'_{q,j})_k)^2 \\ &= \|\mathbf{x}(v_p^i) - \mathbf{x}(v_q^j)\|_2^2 + \sum_{m=1}^{l_r} \|\mathbf{x}(u_m) - \mathbf{x}(\nu_m)\|_2^2 \end{aligned}$$

$$\therefore \text{dist}^2(\mathbf{z}_p^i - \mathbf{z}_q^j) = \|\mathbf{z}_p^i - \mathbf{z}_q^j\|_2^2 \geq \lambda_{\min}(\|\mathbf{x}(v_p^i) - \mathbf{x}(v_q^j)\|_2^2 + \sum_{k=1}^{l_r} \|\mathbf{x}(u_k) - \mathbf{x}(\nu_k)\|_2^2) \quad \square$$

The key point in these theories is that for any task sequence with certain distance among the tasks, there exists a configuration that ensures HPNs to be capable of preserving the task distance after projecting the data into the hidden space, so that only the prototypes associated with the current task are refined and the prototypes corresponding to the other tasks are preserved. Specifically, theorem on zero-forgetting can be formulated as follows:

Theorem 3 (Task distance preserving). *For HPNs trained on consecutive tasks \mathcal{T}^p and \mathcal{T}^{p+1} . If $l_a d_a + l_r d_r \geq (l_r + 1)d_v$ and \mathbf{W} is column full rank, then as long as $t_A < \lambda_{\min}(l_r + 1) \text{dist}(\mathbb{V}_p, \mathbb{V}_{p+1})$, learning on \mathcal{T}^{p+1} will not modify representations HPNs generate for data from \mathcal{T}^p , i.e. catastrophic forgetting is avoided.*

In Theorem 3, λ_i is eigenvalues of the $\mathbf{W}^T \mathbf{W}$, where \mathbf{W} is the matrix mentioned before constructed via AFEs. d_v , d_a and d_r are dimensions of data, atomic node embeddings, and atomic structure embeddings.

Proof. Following the proofs above, suppose two nodes v_p^i and v_q^j are embedded into \mathbf{z}_p^i and \mathbf{z}_q^j with the embedding module. Then the distance between \mathbf{z}_p^i and \mathbf{z}_q^j could be formulated as:

$$\text{dist}(\mathbf{z}_p^i, \mathbf{z}_q^j) = \|\mathbf{z}_p^i - \mathbf{z}_q^j\|_2 = \sqrt{(\mathbf{z}_p^i - \mathbf{z}_q^j)^T (\mathbf{z}_p^i - \mathbf{z}_q^j)}$$

According to lemma 5, we have $\text{dist}^2(\mathbf{z}_p^i, \mathbf{z}_q^j) = \|\mathbf{z}_p^i - \mathbf{z}_q^j\|_2^2 \geq \lambda_{\min}(\|\mathbf{x}(v_p^i) - \mathbf{x}(v_q^j)\|_2^2 + \sum_{k=1}^{l_r} \|\mathbf{x}(u_k) - \mathbf{x}(\nu_k)\|_2^2)$.

$\therefore v_p^i \in \mathbb{V}_p, v_q^j \in \mathbb{V}_q$, $\therefore \|\mathbf{x}(v_p^i) - \mathbf{x}(v_q^j)\|_2^2 \geq \text{dist}^2(\mathbb{V}_p, \mathbb{V}_q)$.

Similarly, $\|\mathbf{x}(u_k) - \mathbf{x}(\nu_k)\|_2^2 \geq \text{dist}^2(\mathbb{V}_p, \mathbb{V}_q)$, for $\forall k$.

$\therefore \|\mathbf{z}_p^i, \mathbf{z}_q^j\|_2^2 \geq \lambda_{\min}(l_r + 1) \text{dist}^2(\mathbb{V}_p, \mathbb{V}_q)$

$\therefore \text{dist}(\mathbf{z}_p^i, \mathbf{z}_q^j) = \|\mathbf{z}_p^i - \mathbf{z}_q^j\|_2 \geq \sqrt{\lambda_{\min}(l_r + 1) \text{dist}^2(\mathbb{V}_p, \mathbb{V}_q)}$

\therefore If $t_A < \sqrt{\lambda_{\min}(l_r + 1) \text{dist}^2(\mathbb{V}_p, \mathbb{V}_q)}$, the embeddings of two nodes from two different tasks will not be assigned to same A-prototypes.

Above all, if the conditions in Theorem 3 are satisfied, learning on new tasks will not modify the prototypes for previous tasks. Besides, the data from previous tasks will be exactly matched to the correct prototypes after training the model on new tasks. In practice, the conditions may not be easy to be satisfied all the time. However, as mentioned in the

paper, the bound given in Theorem 3 is not tight, thus fully satisfying the conditions may not be necessary. Therefore, in the experimental section in the paper, we practically show how the important factors included in these conditions influence the performance (Section 3.6 in the paper). The results demonstrate that the more we satisfy the conditions, the better performance we will obtain, and certain factors (number of AFEs) influence more than the others. \square

Remark 2. When $\text{dist}(\mathbb{V}_p, \mathbb{V}_q) = 0$, i.e. there exists a non-empty set $\mathbb{V}_\cap = \mathbb{V}_p \cap \mathbb{V}_q$, st. $\text{dist}(\mathbb{V}_p \setminus \mathbb{V}_\cap, \mathbb{V}_q \setminus \mathbb{V}_\cap) > 0$, then Theorem 3 holds. As for the \mathbb{V}_\cap containing examples exactly same in \mathbb{V}_p and \mathbb{V}_q , there are two situations:

1. $\forall v \in \mathbb{V}_\cap, y_p(v) = y_q(v)$, where $y_p(\cdot)$ and $y_q(\cdot)$ denote the associated labels in task p and q

2. $\exists v \in \mathbb{V}_\cap, y_p(v) \neq y_q(v)$

For situation 1, \mathbb{V}_\cap will not cause forgetting on the previous task, as these shared data are exactly same and will optimize the model to same direction. For situation 2, if no task indicator is provided, then these data are contradictory data, if task indicator is provided, then the indicator could be merged into the feature vector of the node, i.e. $\mathbf{x}(v_p)$, then v_p will not belong to \mathbb{V}_\cap .

6 CONCLUSION

In this paper, we proposed Hierarchical Prototype Networks (HPNs), to continuously extract different levels of abstract knowledge (in the form of prototypes) from streams of tasks on graph representation learning. The continual learning performance of HPNs is both theoretically analyzed and experimentally justified from different perspectives on multiple public datasets with up to millions of nodes and tens of millions of edges. Besides, the memory consumption of HPNs is also theoretically analyzed and experimentally verified.

HPNs provide a general framework for continual graph representation learning and can be further extended to tackle more challenging tasks. In the future, we will extend HPNs to heterogeneous graphs by adapting the AFEs and prototypes to capture multiple types of nodes and edges. Moreover, we will also equip the prototypes with connections to denote the relationship among the abstract knowledge and investigate more challenging graph reasoning tasks.

7 ACKNOWLEDGMENT

Dacheng Tao and Xikun Zhang are supported by Australian Research Council Projects FL-170100117, DP-180103424, IH-180100002, and IC-190100031. Dongjin Song is supported by UConn Scholarship Facilitation Fund (SFF) Award, NSF CNS-2154191, and USDA NIFA 2022-67022-36603.

REFERENCES

- [1] B. Perozzi, R. Al-Rfou, and S. Skiena, "Deepwalk: Online learning of social representations," in *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2014, pp. 701–710.
- [2] S. Bhagat, G. Cormode, and S. Muthukrishnan, "Node classification in social networks," in *Social Network Data Analytics*. Springer, 2011, pp. 115–148.
- [3] D. Liben-Nowell and J. Kleinberg, "The link-prediction problem for social networks," *Journal of the American Society for Information Science and Technology*, vol. 58, no. 7, pp. 1019–1031, 2007.
- [4] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," *arXiv preprint arXiv:1609.02907*, 2016.
- [5] W. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," in *Advances in Neural Information Processing Systems*, 2017, pp. 1024–1034.
- [6] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio, "Graph attention networks," *arXiv preprint arXiv:1710.10903*, 2017.
- [7] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, "How powerful are graph neural networks?" *arXiv preprint arXiv:1810.00826*, 2018.
- [8] J. Chen, T. Ma, and C. Xiao, "Fastgcn: fast learning with graph convolutional networks via importance sampling," *arXiv preprint arXiv:1801.10247*, 2018.
- [9] D. Zou, Z. Hu, Y. Wang, S. Jiang, Y. Sun, and Q. Gu, "Layer-dependent importance sampling for training deep and large graph convolutional networks," in *Advances in Neural Information Processing Systems*, 2019, pp. 11 247–11 256.
- [10] G. Li, M. Muller, A. Thabet, and B. Ghanem, "Deepgcn: Can gcn go as deep as cnns?" in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 9267–9276.
- [11] M. Chen, Z. Wei, Z. Huang, B. Ding, and Y. Li, "Simple and deep graph convolutional networks," in *International Conference on Machine Learning*. PMLR, 2020, pp. 1725–1735.
- [12] Z. Wang, Z. Li, R. Wang, F. Nie, and X. Li, "Large graph clustering with simultaneous spectral embedding and discretization," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [13] P. Sen, G. Namata, M. Bilgic, L. Getoor, B. Galligher, and T. Eliassi-Rad, "Collective classification in network data," *AI Magazine*, vol. 29, no. 3, pp. 93–93, 2008.
- [14] K. Wang, Z. Shen, C. Huang, C.-H. Wu, Y. Dong, and A. Kanakia, "Microsoft academic graph: When experts are not enough," *Quantitative Science Studies*, vol. 1, no. 1, pp. 396–413, 2020.
- [15] T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," *arXiv preprint arXiv:1310.4546*, 2013.
- [16] K. Bhatia, K. Dahiya, H. Jain, P. Kar, A. Mittal, Y. Prabhu, and M. Varma, "The extreme classification repository: Multi-label datasets and code," 2016. [Online]. Available: <http://manikvarma.org/downloads/XC/XMLRepository.html>
- [17] F. Zhou and C. Cao, "Overcoming catastrophic forgetting in graph neural networks with experience replay," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 5, 2021, pp. 4714–4722.
- [18] H. Liu, Y. Yang, and X. Wang, "Overcoming catastrophic forgetting in graph neural networks," *arXiv preprint arXiv:2012.06002*, 2020.
- [19] C. R. Bowman, T. Iwashita, and D. Zeithamova, "Tracking prototype and exemplar representations in the brain across learning," *Elife*, vol. 9, p. e59360, 2020.
- [20] T. Dandan, Z. Haixue, L. Wenfu, Y. Wenjing, Q. Jiang, and Z. Qinglin, "Brain activity in using heuristic prototype to solve insightful problems," *Behavioural Brain Research*, vol. 253, pp. 139–144, 2013.
- [21] D. Zeithamova, W. T. Maddox, and D. M. Schyns, "Dissociable prototype learning systems: Evidence from brain imaging and behavior," *Journal of Neuroscience*, vol. 28, no. 49, pp. 13 194–13 201, 2008.
- [22] H. Jung, J. Ju, M. Jung, and J. Kim, "Less-forgetting learning in deep neural networks," *arXiv preprint arXiv:1607.00122*, 2016.
- [23] Z. Li and D. Hoiem, "Learning without forgetting," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 12, pp. 2935–2947, 2017.
- [24] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska et al., "Overcoming catastrophic forgetting in neural networks," *Proceedings of the National Academy of Sciences*, vol. 114, no. 13, pp. 3521–3526, 2017.
- [25] M. Farajtabar, N. Azizan, A. Mott, and A. Li, "Orthogonal gradient descent for continual learning," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 3762–3773.
- [26] G. Saha and K. Roy, "Gradient projection memory for continual learning," in *International Conference on Learning Representation*, 2021.
- [27] D. Lopez-Paz and M. Ranzato, "Gradient episodic memory for continual learning," in *Advances in Neural Information Processing Systems*, 2017, pp. 6467–6476.
- [28] H. Shin, J. K. Lee, J. Kim, and J. Kim, "Continual learning with deep generative replay," in *Advances in Neural Information Processing Systems*, 2017, pp. 2990–2999.
- [29] R. Aljundi, M. Lin, B. Goujaud, and Y. Bengio, "Gradient based sample selection for online continual learning," in *Advances in Neural Information Processing Systems*, 2019, pp. 11 816–11 825.

- [30] L. Caccia, E. Belilovsky, M. Caccia, and J. Pineau, "Online learned continual compression with adaptive quantization modules," in *International Conference on Machine Learning*. PMLR, 2020, pp. 1240–1250.
- [31] A. Chrysakis and M.-F. Moens, "Online continual learning from imbalanced data," in *International Conference on Machine Learning*. PMLR, 2020, pp. 1952–1961.
- [32] A. A. Rusu, N. C. Rabinowitz, G. Desjardins, H. Soyer, J. Kirkpatrick, K. Kavukcuoglu, R. Pascanu, and R. Hadsell, "Progressive neural networks," *arXiv preprint arXiv:1606.04671*, 2016.
- [33] J. Yoon, E. Yang, J. Lee, and S. J. Hwang, "Lifelong learning with dynamically expandable networks," *arXiv preprint arXiv:1708.01547*, 2017.
- [34] J. Yoon, S. Kim, E. Yang, and S. J. Hwang, "Scalable and order-robust continual learning with additive parameter decomposition," in *International Conference on Learning Representation*, 2020.
- [35] M. Wortsman, V. Ramanujan, R. Liu, A. Kembhavi, M. Rastegari, J. Yosinski, and A. Farhadi, "Supermasks in superposition," *arXiv preprint arXiv:2006.14769*, 2020.
- [36] Y. Wu, Y. Chen, L. Wang, Y. Ye, Z. Liu, Y. Guo, and Y. Fu, "Large scale incremental learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 374–382.
- [37] J. Knoblauch, H. Husain, and T. Diethe, "Optimal continual learning has perfect memory and is np-hard," in *International Conference on Machine Learning*. PMLR, 2020, pp. 5327–5337.
- [38] H. Bahonar, A. Mirzaei, S. Sadri, and R. C. Wilson, "Graph embedding using frequency filtering," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 2, pp. 473–484, 2019.
- [39] X. Zhang, C. Xu, X. Tian, and D. Tao, "Graph edge convolutional neural networks for skeleton-based action recognition," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 31, no. 8, pp. 3047–3060, 2019.
- [40] X. Zhang, C. Xu, and D. Tao, "Context aware graph convolution for skeleton-based action recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 14 333–14 342.
- [41] H. Lei, N. Akhtar, and A. Mian, "Spherical kernel for efficient graph convolution on 3d point clouds," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [42] H. Gao, Y. Liu, and S. Ji, "Topology-aware graph pooling networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [43] Y. Xu, Q. Zhang, J. Zhang, and D. Tao, "Vita: Vision transformer advanced by exploring intrinsic inductive bias," *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [44] Q. Zhang, Y. Xu, J. Zhang, and D. Tao, "Vitaev2: Vision transformer advanced by exploring inductive bias for image recognition and beyond," *arXiv preprint arXiv:2202.10108*, 2022.
- [45] Y. Rong, W. Huang, T. Xu, and J. Huang, "Dropedge: Towards deep graph convolutional networks on node classification," in *International Conference on Learning Representations*, 2019.
- [46] X. Zhang, C. Xu, and D. Tao, "On dropping clusters to regularize graph convolutional neural networks," in *European Conference on Computer Vision*, 2020.
- [47] L. Galke, I. Vagliano, and A. Scherp, "Incremental training of graph neural networks on temporal graphs under distribution shift," *arXiv preprint arXiv:2006.14422*, 2020.
- [48] J. Wang, G. Song, Y. Wu, and L. Wang, "Streaming graph neural networks via continual learning," in *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, 2020, pp. 1515–1524.
- [49] C. Wang, Y. Qiu, and S. Scherer, "Lifelong graph learning," *arXiv preprint arXiv:2009.00647*, 2020.
- [50] W. Yu, W. Cheng, C. C. Aggarwal, K. Zhang, H. Chen, and W. Wang, "Netwalk: A flexible deep embedding approach for anomaly detection in dynamic networks," in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2018, pp. 2672–2681.
- [51] G. H. Nguyen, J. B. Lee, R. A. Rossi, N. K. Ahmed, E. Koh, and S. Kim, "Continuous-time dynamic network embeddings," in *Companion Proceedings of the The Web Conference 2018*, 2018, pp. 969–976.
- [52] L. Zhou, Y. Yang, X. Ren, F. Wu, and Y. Zhuang, "Dynamic network embedding by modeling triadic closure process," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018.
- [53] Y. Ma, Z. Guo, Z. Ren, J. Tang, and D. Yin, "Streaming graph neural networks," in *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2020, pp. 719–728.
- [54] F. Zhou, C. Cao, K. Zhang, G. Trajcevski, T. Zhong, and J. Geng, "Meta-gnn: On few-shot node classification in graph meta-learning," in *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 2019, pp. 2357–2360.
- [55] Z. Guo, C. Zhang, W. Yu, J. Herr, O. Wiest, M. Jiang, and N. V. Chawla, "Few-shot graph learning for molecular property prediction," in *Proceedings of the Web Conference 2021*, 2021, pp. 2559–2567.
- [56] H. Yao, C. Zhang, Y. Wei, M. Jiang, S. Wang, J. Huang, N. Chawla, and Z. Li, "Graph few-shot learning via knowledge transfer," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 04, 2020, pp. 6656–6663.
- [57] H. Pei, B. Wei, K. C.-C. Chang, Y. Lei, and B. Yang, "Geom-gcn: Geometric graph convolutional networks," *arXiv preprint arXiv:2002.05287*, 2020.
- [58] J. Tang, J. Sun, C. Wang, and Z. Yang, "Social influence analysis in large-scale networks," in *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge discovery and data mining*, 2009, pp. 807–816.
- [59] R. Aljundi, F. Babiloni, M. Elhoseiny, M. Rohrbach, and T. Tuytelaars, "Memory aware synapses: Learning what (not) to forget," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 139–154.
- [60] F. Zhou, C. Cao, T. Zhong, K. Zhang, G. Trajcevski, and J. Geng, "Continual graph learning," *arXiv preprint arXiv:2003.09908*, 2020.
- [61] G. M. Van de Ven and A. S. Tolias, "Three scenarios for continual learning," *arXiv preprint arXiv:1904.07734*, 2019.
- [62] L. Van der Maaten and G. Hinton, "Visualizing data using t-sne." *Journal of Machine Learning Research*, vol. 9, no. 11, 2008.



Xikun Zhang is Ph.D. student at the School of Computer Science in The University of Sydney. He obtained a M.phil from the School of Computer Science in The University of Sydney, and a Bachelor of Applied Physics from University of Science and Technology of China. His research interests mainly include applying deep learning to tackle graph related problems.



Dongjin Song (M'19) is an assistant professor in the Department of Computer Science & Engineering at the University of Connecticut (UConn). His research interests include data science, artificial intelligence, deep learning, time series analysis, graph representation learning, and related applications. Papers describing his research have been published at top-tier data science and artificial intelligence conferences, such as ICML, ICLR, KDD, ICDM, SDM, AAAI, IJCAI, CVPR, ICCV, MM, etc. He won the UConn REP award in

2021, UConn SFF award in 2022, and the best scientific project award in the D4D Challenge in 2013.



Dacheng Tao (Fellow, IEEE) is currently a Professor of Computer Science and an ARC Laureate Fellow in the School of Computer Science and the Faculty of Engineering at The University of Sydney. He mainly applies statistics and mathematics to artificial intelligence and data science. His research is detailed in one monograph and over 200 publications in prestigious journals and proceedings at prominent conferences such as IEEE TPAMI, IJCV, JMLR, NeurIPS, ICML, ICLR, CVPR, ICCV, ECCV, AAAI, IJCAI, ICDM and

ACM SIGKDD, with several best paper awards, such as the Best Theory/Algorithm Paper Runner Up Award at IEEE ICDM'07, the Distinguished Paper Award at 2018 IJCAI, the 2014 ICDM 10-year Highest-Impact Paper Award, and the 2017 IEEE Signal Processing Society Best Paper Award. He received the 2015 and 2020 Australian Eureka Prize and the 2018 IEEE ICDM Research Contributions Award. He is a fellow of the Australian Academy of Science, TWAS, AAAS, ACM and IEEE.