

Leveraging Meta-path Contexts for Classification in Heterogeneous Information Networks

Xiang Li*, Danhao Ding[†], Ben Kao[†], Yizhou Sun[‡] and Nikos Mamoulis[§]

*School of Data Science and Engineering, East China Normal University, Shanghai, China

[†]Department of Computer Science, The University of Hong Kong, Hong Kong

[‡]Department of Computer Science, University of California, Los Angeles, USA

[§]Department of Computer Science and Engineering, University of Ioannina, Greece

Email: *xiangli@dase.ecnu.edu.cn, †{dhding2, kao}@cs.hku.hk, ‡yzsun@cs.ucla.edu, §nikos@cs.uoi.gr

Abstract—A heterogeneous information network (HIN) has as vertices objects of different types and as edges the relations between objects, which are also of various types. We study the problem of classifying objects in HINs. Most existing methods perform poorly when given scarce labeled objects as training sets, and methods that improve classification accuracy under such scenarios are often computationally expensive. To address these problems, we propose ConCH, a graph neural network model. ConCH formulates the classification problem as a multi-task learning problem that combines semi-supervised learning with self-supervised learning to learn from both labeled and unlabeled data. ConCH employs meta-paths, which are sequences of object types that capture semantic relationships between objects. ConCH co-derives object embeddings and context embeddings via graph convolution. It also uses the attention mechanism to fuse such embeddings. We conduct extensive experiments to evaluate the performance of ConCH against other 15 classification methods. Our results show that ConCH is an effective and efficient method for HIN classification.

Index Terms—heterogeneous information networks, classification, graph neural networks

I. INTRODUCTION

A *Heterogeneous Information Network* (HIN) is one whose objects are of different types and whose edges represent different types of relations between objects. Different from a *homogeneous* network, where objects (and edges) are all of one single type, HINs are more expressive in capturing the rich semantics of objects and their relationships in real-world applications. In many HINs, objects are associated with descriptive labels. For example, authors in the bibliographic network DBLP can be labeled by research areas; movies in Freebase, a human-curated knowledge base, are labeled by genres. However, object labeling is costly and it has been observed that only a small portion of objects are given labels in HINs such as Freebase and Yago [1]. Therefore, graph-based semi-supervised learning, which assigns labels to unlabeled objects in graphs based on a given set of labeled objects and the graph structure, has attracted a lot of research attention.

Conventionally, graph-based semi-supervised learning methods adopt a loss function of the form $\mathcal{L} = \mathcal{L}_{sup} + \lambda\mathcal{L}_{reg}$, where \mathcal{L}_{sup} is a loss measured with a set of labeled objects, and \mathcal{L}_{reg} is a graph-based regularization term. For example, assuming that linked objects in HINs are more likely to share the same label, some methods [2], [3] employ the graph

Laplacian regularization. However, as pointed out in [4], the edges in a graph do not necessarily imply the similarity between the connected objects. For example, one can construct a network of words using WordNet data with edges representing either synonym or antonym relations. To address the issue, *Graph Convolutional Network* (GCN) [4] has been proposed to classify objects in homogeneous networks without an explicit graph Laplacian regularization. It is, however, challenging to apply GCN on HINs for two reasons. First, GCN disregards the possibly different edge types and applies the same aggregation function for encoding all heterogeneous relations between objects [5]. Information on the edge types is therefore not used. Second, in HINs, more complex semantic relations between objects are often exhibited by multi-hop paths instead of single links. To aggregate information from important multi-hop (or *path-based*) neighbors, multiple neural-net layers have to be stacked up in GCN. However, the exponential growth in the neighborhood size as we explore path-based relations of progressively longer path lengths greatly increases the computation cost. Also, it has been reported that the performance of GCN degrades when there are numerous layers.

Recently, attempts [6]–[10] have been made to extend GCN to classify objects in HINs. In particular, some works [6], [10] leverage meta-paths. A *meta-path* is a sequence of object types that captures the semantic relation between objects in HINs. For example, if we denote the object types *Author* and *Paper* in DBLP as “A” and “P”, respectively, the meta-path *Author-Paper-Author* (APA) expresses the *co-authorship* relation. Specifically, two author objects a_1 and a_2 are *APA-related* if a path a_1-p-a_2 exists, where p represents a paper object and an edge here represents authorship. The use of meta-paths identifies a set of *significant* path-based neighbors that are semantically related to a given object. This in turn helps limit the neighborhood size by ignoring other less significant path-based neighbors and thus saves computation cost. However, existing GCN models in HINs could suffer from two main problems. On the one hand, while these methods have been shown to be effective, most studies have been performed on large training sets. When labeled objects are scarce, the performance of these methods degrades (as shown in Table I). On the other hand, some methods improve classification

accuracy by sacrificing efficiency. For example, HGT [8] proposes a Transformer [11] based aggregator to combine information from multi-type neighbors for an object. This aggregator architecture introduces more learnable parameters, which increases the training difficulty and adversely affects the model efficiency. Therefore, a research question arises: *Given a limited set of labeled objects, can we develop a simple model that is both effective and efficient for classification in HINs?*

In this paper, inspired by [5], [10], [12], we apply GCN to HINs using meta-paths and considering the finer information provided by the paths of a meta-path. For example, given a path a_1-p-a_2 , we can deduce a co-authorship relation between a_1 and a_2 . We further remark that the identity of the path, which we will call a *context* of the co-authorship relation, can be very influential in the classification task. For example, the topic (label) of paper p can provide important information of the research topics (labels) of the authors. We propose to leverage meta-path-based contexts to improve classification. We address the problems mentioned above from three perspectives. First, we leverage self-supervised learning to learn from massive unlabeled data. The learned data-driven prior knowledge could provide supplementary information for the shortage of labeled objects in the training set. Specifically, we formulate the problem as a multi-task learning problem, where the self-supervised loss acts as a regularizer to the supervised loss w.r.t. labeled objects. Second, we improve modeling efficiency by selecting the most informative meta-path-specified neighbors for an object, which effectively filters less-informative neighbors. Third, we propose to compute *context embeddings* and integrate them into the computation of object embeddings. Specifically, we propose a process that performs mutual updates of these two kinds of embeddings. Our main contributions are summarized as follows.

- We propose ConCH, which leverages meta-path-based **Contexts** to effectively and efficiently **Classify** objects in **HINs**, given scarce labeled objects as training sets.
- We formulate the classification problem in HINs as a multi-task learning problem that combines semi-supervised learning with self-supervised learning. In particular, when labeled objects are scarce, self-supervision learned from unlabeled data helps improve model performance.
- We improve model efficiency by putting forward a filtering strategy that selects the most informative path-based neighbors of an object. We aggregate information from selected neighbors to compute an object’s embedding.
- We conduct extensive experiments to evaluate ConCH against 15 other competitors. Our results show that ConCH is highly effective and also efficient.

II. RELATED WORK

We discuss four categories of related works: graph-based semi-supervised learning, network embedding, self-supervised learning and graph neural networks. For a survey of heterogeneous network representation learning, see [13].

[Graph-based semi-supervised learning]. Semi-supervised learning has been widely studied in graphs [14]–

[16]. Generally, a predictive score $f(l_j|x)$ of an object x and a label l_j is defined to indicate the likelihood that object x ’s label is l_j . The objective is to learn $f()$ for all objects and labels by minimizing a loss value that consists of two components: (1) for any labeled object x , the difference between the predictive score $f(l_j|x)$ and x ’s true label’s value, which is 1 if x ’s label is l_j ; 0 otherwise; (2) for any two linked objects x_u and x_v in the graph, the difference between their predictive scores $f(l_j|x_u)$ and $f(l_j|x_v)$ over the set of labels. The minimization of the second component is mostly achieved by explicit forms of graph Laplacian regularization. Semi-supervised learning in HINs has also been studied [2], [3], [17]. For example, GNetMine [2] is a transductive classifier that applies the idea of label predictions via predictive scores $f(l_j|x)$. As another example, Grempt [3] is a transductive regression model that employs meta-paths to deduce objects’ relations and it employs graph Laplacian regularization.

[Network embedding]. Network embedding has been widely studied to learn representation vectors of objects in networks [19]–[21]. These representation vectors can be fed into a variety of downstream tasks, such as link prediction, classification and recommendation. Early studies are for homogeneous networks [20], [22], [23] with a focus of encoding a network’s connectivity. For example, inspired by the idea of word-context pairs in sentences from word2vec [24], DeepWalk [22] introduces node-context pairs in random walk sequences. The random walk sequences are fed into the SkipGram model [24] to generate node embeddings. Other random-walk-based methods have also been proposed, such as node2vec [23]. Some works consider information other than network structure, such as object attributes [25] and object labels [26]. There are also methods that study various characteristics of network embeddings. For example, [27] puts forward a unified matrix factorization framework for some well-known network embedding methods, such as LINE [28]. Network embedding in HINs has also attracted great attention recently [29]–[31]. Some of these methods are meta-path-based. For example, metapath2vec [32] performs meta-path-based random walks to construct heterogeneous neighbors of an object. Taking object embeddings as parameters, HIN2Vec [33] constructs a binary classifier that predicts whether a given pair of objects are related by a meta-path relation. There are also specialized methods that are designed for specific tasks [34] or specific categories of HINs [35].

[Self-supervised learning]. Self-supervised learning plays an increasing role in utilizing unlabeled data. It has been widely used in computer vision [36]–[38] and natural language processing [39]–[41]. Recently, there are also works [42]–[49] that study self-supervised learning in graphs and further apply it for network representation learning. For example, deep graph infomax (DGI) [45] learns node representations by contrasting node and graph encodings. In [48], MVGRL learns node and graph representations by contrasting encodings from first-order neighbors (adjacency matrix) and a graph diffusion matrix, which are respectively taken as local and global

views of a graph structure. Further, heterogeneous deep graph infomax (HDGI) [49] extends DGI to HINs, to learn node representations by mutual information maximization.

[Graph neural networks]. Kipf and Welling [4] point out that traditional semi-supervised learning approaches on graphs rely on the assumption that edges encode similarity between objects, which may not be true. They extend the convolution operation on graphs to encode network structures, and propose the GCN model to avoid explicit forms of graph Laplacian regularization. GCN is based on spectral graph convolutional neural networks [50], [51], which decompose graph signals via graph Fourier transform and convolve on the spectral components. There are also a series of spatial graph convolutional neural network models that directly convolve information from spatially nearby neighbors of objects in graphs. For example, GraphSAGE [52] generates an object’s embedding vector by aggregating information from a fixed-size neighborhood of the object. Graph attention networks (GATs) [53] employ the attention mechanism to learn the importance of an object’s neighbors and aggregate information from these neighbors with their learned weights. However, all these models are designed for homogeneous networks.

There are also works that extend graph convolutional neural networks to HINs [5]–[10], [12], [54]–[57]. For example, to generate an object’s embedding, HetGNN [7] first derives a set of object-type-based neighbor embeddings by aggregating information from neighbors in the same type with bi-directional LSTM. After that, HetGNN adopts the vanilla attention mechanism to fuse these object-type-based neighbor embeddings. Inspired by the architecture of Transformer [11], HGT [8] proposes heterogeneous graph transformer that leverages multi-head attentions to distinguish different relation types between an object and its neighbors. Further, HGCN [9] distinguishes the various edge types and derives multiple sub-networks. In each convolutional layer, for each object type in each sub-network, HGCN aggregates information from an object’s neighbors by leveraging multiple kernels that use different convolutional filters with different aggregation strategies. Embeddings derived from various kernels in these sub-networks are fused to construct the object’s relational feature vector. Finally, relational features and original features of the object are concatenated, which are fed into a MLP to predict the object’s labels. However, with more neural-net layers stacked up in these models, the neighborhood size of an object exponentially increases. To address the problem, some methods use meta-paths to identify a subset of *multi-hop neighbors* of a given object¹. For example, for each given meta-path \mathcal{P} and an object x , HAN [6] uses a node-level attention to learn the importance of all multi-hop neighbors specified by \mathcal{P} and generates x ’s embedding by aggregating information from these neighbors. After that, a meta-path-level attention is used to learn the importance of meta-paths, and x ’s embeddings based on the various meta-paths are fused

¹ A multi-hop neighbor of an object x is one that is reached from x via a path of multiple edges.

to obtain x ’s final embedding. MAGNN [10] further improves HAN by utilizing meta-path-based contexts. Given a meta-path \mathcal{P} and an object x , MAGNN learns the importance of path instances of \mathcal{P} starting from x and generates x ’s embedding by aggregating information from these path instances with the learned weights. Finally, the embeddings learned from various meta-paths are combined to generate x ’s final embedding.

Despite their success, we observe that most existing methods perform poorly when the number of given labeled objects is limited. Also, many of them improve the model’s effectiveness by sacrificing efficiency. Different from these methods, our proposed method ConCH can achieve superior performance with high efficiency, given scarce labeled objects.

III. DEFINITIONS

In this section we formally define various concepts and the HIN classification problem.

Definition 1: Heterogeneous Information Network (HIN). Let $\mathcal{T} = \{T_1, \dots, T_m\}$ be a set of m object types. For each type T_i , let \mathcal{X}_i be the set of objects of type T_i . An HIN is a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \bigcup_{i=1}^m \mathcal{X}_i$ is a set of nodes and \mathcal{E} is a set of edges, each represents a binary relation between two objects in \mathcal{V} . An HIN $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is associated with two mappings. (1) $\phi : \mathcal{V} \rightarrow \mathcal{T}$ is an object-type mapping that maps an object in \mathcal{V} into its type, and (2) $\psi : \mathcal{E} \rightarrow \mathcal{R}$ is a link-relation mapping that maps a link in \mathcal{E} into a relation in a set of relations \mathcal{R} . A network is an HIN if $|\mathcal{T}| > 1$ or $|\mathcal{R}| > 1$; otherwise, the network is homogeneous. \square

Definition 2: Network schema. The network schema of an HIN \mathcal{G} , denoted by $T_{\mathcal{G}} = (\mathcal{T}, \mathcal{R})$, shows how objects of different types are related by the relations in \mathcal{R} . A *schematic graph* is used to represent $T_{\mathcal{G}}$ with \mathcal{T} and \mathcal{R} being the node set and the edge set, respectively. Specifically, there is an edge (T_i, T_j) in the schematic graph iff there is a relation in \mathcal{R} that relates objects of type T_i to objects of type T_j . \square




Fig. 1. An HIN and its schematic graph

As an example, Fig. 1(a) shows a movie HIN with object types: $\mathcal{T} = \{\text{Movie}(\square), \text{Actor}(\diamond), \text{Director}(\circ), \text{Producer}(\circ)\}$. The set \mathcal{R} consists of three relations, which are illustrated by the three edges in the schematic graph (Fig. 1(b)). For example, the edge A–M in Fig. 1(b) represents the relation that an (A)ctor acts in a (M)ovie, and the edge M–D denotes the relation that a (M)ovie is directed by a (D)irector.

Definition 3: Meta-path. A meta-path $\mathcal{P}: T_1 \xrightarrow{R_1} \dots \xrightarrow{R_l} T_{l+1}$ defines a composite relation $R = R_1 \circ \dots \circ R_l$ that relates objects of type T_1 to objects of type T_{l+1} . If two objects x_u and x_v are related by the composite relation R , then there

is a path, denoted by $p_{x_u \rightsquigarrow x_v}$, that connects x_u to x_v in \mathcal{G} . Moreover, the sequence of objects and links in $p_{x_u \rightsquigarrow x_v}$ matches the sequence of types T_1, \dots, T_{l+1} and relations R_1, \dots, R_l based on the object-type mapping ϕ and the link-relation mapping ψ , respectively. We say that $p_{x_u \rightsquigarrow x_v}$ is a **path instance** of \mathcal{P} , denoted by $p_{x_u \rightsquigarrow x_v} \vdash \mathcal{P}$. \square

In Fig. 1(a), the path $p_{M1 \rightsquigarrow M2} = M1 \rightarrow A1 \rightarrow M2$ is an instance of the meta-path *Movie-Actor-Movie* (abbrev. MAM), which captures the relation between two movies that feature the same actor; the path $p_{M2 \rightsquigarrow M3} = M2 \rightarrow P1 \rightarrow M3$ is an instance of the meta-path *Movie-Producer-Movie* (abbrev. MPM), which expresses the relation between two movies that are produced by the same producer. Path instances describe the details on how two objects are related by a meta-path, which can be used to define meta-path-based contexts.

Definition 4: Meta-path-based context [12]. Given two objects x_u and x_v that are related by a meta-path \mathcal{P} , the meta-path-based context $c_{uv}^{\mathcal{P}}$ is the set of path instances of \mathcal{P} connecting objects x_u and x_v , i.e., $c_{uv}^{\mathcal{P}} = \{p_{x_u \rightsquigarrow x_v} \mid p_{x_u \rightsquigarrow x_v} \vdash \mathcal{P}\}$. \square

In the following discussion, we will use *mp-context* or simply *context* as short forms of “meta-path-based context”. As an example, in Fig. 1(a), the mp-context of M1 and M2 w.r.t. the meta-path MAM is $c_{M1 M2}^{MAM} = \{M1 \rightarrow A1 \rightarrow M2, M1 \rightarrow A2 \rightarrow M2\}$.

Definition 5: Classification in HINs. Given an HIN $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, a target object type T to be classified, a label set $L = \{l_1, l_2, \dots, l_r\}$, and a set of meta-paths \mathcal{PS} , let $\mathcal{X} = \mathcal{Y} \cup \mathcal{U}$, where \mathcal{Y} is a set of labeled objects and \mathcal{U} is a set of unlabeled ones, the classification problem in HINs is to predict the labels of objects in \mathcal{U} . \square

IV. ALGORITHM

In this section we describe our ConCH algorithm. We first give an overview of ConCH, which is illustrated in Fig. 2. For an object x in an HIN, ConCH first identifies a selected set of x ’s path-based neighbors based on a given set of meta-paths that express various semantic relations (Step ①). These meta-path-based relations relate objects in the HIN with path instances, which are taken as the mp-contexts (Step ②). ConCH then constructs a bipartite graph for each given meta-path \mathcal{P} (Step ③). The bipartite graph represents objects and the mp-contexts that relate them, and it facilitates the updating of objects’ and contexts’ embeddings (Step ④). After that, the objects’ embeddings generated by various meta-paths are fused to obtain the objects’ final embeddings (Step ⑤), which are then fed into a two-layer MLP for label prediction (Step ⑥). To better utilize unlabeled data, ConCH also uses self-supervised learning. For each meta-path, in addition to the bipartite graph constructed in Step ③, ConCH further generates a “negative” bipartite graph (Step ⑦). From these two graphs, a positive sample set and a negative sample set are constructed. After that, ConCH uses a discriminator \mathcal{D} to distinguish the two sets (Step ⑧). Finally, our objective function is formulated as a multi-task learning problem by combing both the supervised

loss and the self-supervised loss. Next, we describe each component in detail.

A. Neighbor filtering

Given an object x in an HIN, our objective is to identify other objects in the network that are similar to x and use their labels to infer that of x . We call these other objects *relevant neighbors*. Note that these neighbors could be multiple hops away from x , especially if they are related to x via paths that are instances of certain given meta-paths that express important semantic relations between objects. An interesting question is how relevant neighbors can be obtained effectively and efficiently. If meta-paths are numerous (e.g., those that are obtained via automatic methods) and long, then the neighbors derived from them could cover a large scope of the network. Directly aggregating information from large numbers of neighbors would make model construction inefficient. While there are sampling methods [52] of neighbors, the sampling process itself could be time-consuming and less relevant neighbors may be sampled instead of the more relevant ones. Our first step is to filter neighbors and select the most relevant neighbors of an object.

Given a meta-path \mathcal{P} , ConCH measures the similarity between two objects x_u and x_v of the same type w.r.t. \mathcal{P} by *PathSim* [58]:

$$\text{PS}(x_u, x_v) = \frac{2 \times |\{p_{x_u \rightsquigarrow x_v} \mid p_{x_u \rightsquigarrow x_v} \vdash \mathcal{P}\}|}{|\{p_{x_u \rightsquigarrow x_u} \mid p_{x_u \rightsquigarrow x_u} \vdash \mathcal{P}\}| + |\{p_{x_v \rightsquigarrow x_v} \mid p_{x_v \rightsquigarrow x_v} \vdash \mathcal{P}\}|} \quad (1)$$

PathSim has been shown to be very effective in a variety of data mining tasks in HINs [1], [3], [18], [59]. Given an object x and a meta-path \mathcal{P} , ConCH collects a set of neighbors $N_x^{\mathcal{P}}$ that contains the top- k objects with the highest PathSim scores with x w.r.t. \mathcal{P} . Using only the top- k neighbors significantly reduces the number of relevant objects that need to be considered in the subsequent steps of the classification process. Also, it removes less relevant ones and thus helps reduce noise. This further improves the classification effectiveness and efficiency.

B. Context feature construction

Recall that given two objects x_1 and x_l , and a meta-path \mathcal{P} , the mp-context $c_{1l}^{\mathcal{P}}$ is a set of path instances of \mathcal{P} between x_1 and x_l . Learning a context embedding from scratch by treating it as learnable parameter is very costly. To reduce the number of parameters, we construct a feature vector for each context. Specifically, we apply a conventional HIN embedding method, such as *metapath2vec* [32], to obtain initial object embeddings. Then, we generate a path instance’s embedding by aggregating the embedding vectors of all the objects in the path instance. Given a path instance $p_{x_1 \rightsquigarrow x_l} = x_1 x_2 \dots x_l \vdash \mathcal{P}$, its embedding is computed by

$$o_{p_{x_1 \rightsquigarrow x_l}} = \text{MEAN}(\{e_i\}_{i=1}^l), \quad (2)$$

where e_i is the initial embedding vector of $x_i, \forall 1 \leq i \leq l$. An *initial* context embedding is further obtained by aggregating path instances’ embeddings:

$$o_{c_{1l}^{\mathcal{P}}} = \text{MEAN}(\{o_{p_{x_1 \rightsquigarrow x_l}}\} \mid p_{x_1 \rightsquigarrow x_l} \in c_{1l}^{\mathcal{P}}), \quad (3)$$




Fig. 2. The overall framework of the ConCH model. For details of each step, see Sec. IV.

which serves as the context $c_{li}^{\mathcal{P}}$'s feature vector.

C. Meta-path-based bipartite graphs

In the previous steps, we obtained a selected set of path-based neighbors of objects and mp-contexts. We next update their embeddings through a mutual update process. Note that a meta-path-based relation between two objects is influenced by the meta-path instances (i.e., a context) that connect the two objects; At the same time, a context is influenced by the objects that compose the meta-path instances. Hence, we update object embeddings and context embeddings together. To facilitate such an update, we first construct, for each meta-path \mathcal{P} , a bipartite graph $G_{\mathcal{P}} = (\mathcal{X}, V_{\mathcal{C}}^{\mathcal{P}}, E_{OC}^{\mathcal{P}})$, where \mathcal{X} is the set of objects to be classified, $V_{\mathcal{C}}^{\mathcal{P}}$ is the set of meta-path contexts derived from \mathcal{P} , and $E_{OC}^{\mathcal{P}}$ is the set of edges such that e_{ij} connects object x_i in \mathcal{X} and context c_j in $V_{\mathcal{C}}^{\mathcal{P}}$ if the path instances in c_j start from or end at x_i . Note that our neighbor filtering scheme (see Section IV-A) derives top- k neighbors for each object. Therefore, the degree of any object x in the bipartite graph is at most k .

Fig. 2 shows example bipartite graphs derived from a movie HIN. For the meta-path Movie-Actor-Movie (MAM), M1 and M3 are connected by the path instance $M1 \rightarrow A1 \rightarrow M3$, and M2 and M4 are connected by $M2 \rightarrow A2 \rightarrow M4$. The bipartite graph for MAM is shown in Fig. 2(a), in which context $C1 = \{M1 \rightarrow A1 \rightarrow M3\}$ and $C2 = \{M2 \rightarrow A2 \rightarrow M4\}$.

We update object embeddings (denoted by $h_x^{\mathcal{P}}$) and context embeddings (denoted by $h_{c_j}^{\mathcal{P}}$) from each other. For each context c_j that is linked to two objects x_u and x_v in $G_{\mathcal{P}}$, we update its embedding vector $h_{c_j}^{\mathcal{P}}$ in the $(t+1)$ -st timestep by aggregating information from both x_u and x_v :

$$h_{c_j}^{\mathcal{P},(t+1)} = \text{ReLU} \left(W_1^{(t)} \cdot h_{x_u}^{\mathcal{P},(t)} + W_1^{(t)} \cdot h_{x_v}^{\mathcal{P},(t)} + W_2^{(t)} \cdot h_{c_j}^{\mathcal{P},(t)} \right), \quad (4)$$

where $W_1^{(t)}$ and $W_2^{(t)}$ are two linear transformation matrices to be learned. For any object $x_i \in \mathcal{X}$, it connects to at most

k contexts. From these contexts, we update the embedding vector of x_i in the $(t+1)$ -st timestep by:

$$h_{x_i}^{\mathcal{P},(t+1)} = \text{ReLU} \left(W_3^{(t)} \cdot \sum_j h_{c_j}^{\mathcal{P},(t)} + W_4^{(t)} \cdot h_{x_i}^{\mathcal{P},(t)} \right), \quad (5)$$

where $W_3^{(t)}$ and $W_4^{(t)}$ are another two weight matrices. Here, we use the *sum* aggregator. It is efficient and since we have retained only the top- k neighbors that are the most informative, the aggregation should also be effective.

D. Semantic aggregation

Our next step is to fuse together the object embeddings derived from different meta-paths. Since different meta-paths represent different semantic relations, they may contribute in different extents in the classification task. We use the vanilla attention mechanism to learn the importance of meta-paths. Given a meta-path \mathcal{P} and the embedding vector $h_{x_i}^{\mathcal{P}}$ of an object x_i , we feed $h_{x_i}^{\mathcal{P}}$ into a two-layer MLP and get:

$$\tilde{w}_{x_i}^{\mathcal{P}} = a^T \cdot \xi(W_5 \cdot \tanh(W_6 \cdot h_{x_i}^{\mathcal{P}})), \quad (6)$$

where a is the weight vector, ξ is the leaky_ReLU function, and W_5 and W_6 are two linear transformation matrices. We regularize $\tilde{w}_{x_i}^{\mathcal{P}}$ by softmax function and compute the attention weight $w_{x_i}^{\mathcal{P}}$ of the meta-path \mathcal{P} by

$$w_{x_i}^{\mathcal{P}} = \text{softmax}(\tilde{w}_{x_i}^{\mathcal{P}}) = \frac{\exp(\tilde{w}_{x_i}^{\mathcal{P}})}{\sum_q \exp(\tilde{w}_{x_i}^{\mathcal{P}_q})}. \quad (7)$$

Finally, we aggregate the embedding vectors across meta-paths and generate the final embedding vector z_i of x_i by:

$$z_i = \text{ReLU} \left(\sum_q w_{x_i}^{\mathcal{P}_q} \cdot h_{x_i}^{\mathcal{P}_q} \right). \quad (8)$$

E. The ConCH model

After the semantic aggregation, a standard way to guide parameter learning is to use label information in the training set. Specifically, we first feed z_i in Eq. 8 into a two-layer

MLP to output a label vector \tilde{z}_i of length r (r is the number of labels):

$$\tilde{z}_i = W_7(\text{ReLU}(W_8 \cdot z_i)), \quad (9)$$

where W_7 and W_8 are learnable weight matrices. The label corresponding to the largest-value entry in \tilde{z}_i is taken as the predicted label of x_i . Then we minimize the Cross-Entropy loss function:

$$\mathcal{L}_{sup} = - \sum_{x_i \in \mathcal{Y}} y_i \ln(\tilde{z}_i), \quad (10)$$

where \mathcal{Y} is the training set of labeled objects and y_i is the one-hot encoded true label vector of object x_i . While the supervised loss in Eq. 10 is useful, the model performance could be adversely affected when the number of labeled objects is limited. Further, self-supervised learning has shown great potential in utilizing unlabeled data. Therefore, we incorporate self-supervised learning into the learning process to leverage both labeled and unlabeled objects.

After the final embedding vector of an object is derived (in Eq. 8), we compute a summary vector s as the global representation encoder, which retains the information from all the objects of the type to be classified. We simply take the averaging operator and get:

$$s = \text{MEAN}(\{z_i\}_{i=1}^n), \quad (11)$$

where n is the number of objects to be classified. Then we maximize the mutual information between node embeddings and the global representation, which encourages each single object to have access to the information of others. Following [45], we use a noise-contrastive type objective function with a standard binary cross entropy loss:

$$\mathcal{L}_{ss} = \frac{1}{N+M} \left[\sum_{i=1}^N \mathbb{E}_{Pos} [\log \mathcal{D}(z_i, s)] + \sum_{j=1}^M \mathbb{E}_{Neg} [\log(1 - \mathcal{D}(\hat{z}_j, s))] \right], \quad (12)$$

where \mathcal{D} is a discriminator to distinguish the positive sample set $Pos = \{(z_i, s)\}_{i=1}^N$ from the negative sample set $Neg = \{(\hat{z}_j, s)\}_{j=1}^M$. The sample (z_i, s) is positive when object x_i belongs to the original graph; the sample (\hat{z}_j, s) is negative when object x_j is the generated fake one. $\mathcal{D}(z_i, s)$ represents the probability score assigned to the input node-summary pair, and it is defined as

$$\mathcal{D}(z_i, s) = \sigma(z_i^T W_{\mathcal{D}} s), \quad (13)$$

where σ is the sigmoid function and $W_{\mathcal{D}}$ is the learnable weight matrix. To generate negative samples, inspired by [49], for each meta-path-based bipartite graph, we fix the adjacency matrix unchanged and randomly shuffle rows of the initial object feature matrix, which generates a “negative” bipartite graph. For example, Fig. 2(a) and Fig. 2(c) show the original bipartite graph (we call it “positive” for comparison) and the “negative” one derived by the meta-path MAM, respectively. Feature vectors of movies in the negative bipartite graph are shuffled from that in the positive graph. For the constructed negative bipartite graphs, we repeat steps ④ and ⑤ (in Sec IV-C and IV-D) to get \hat{z}_j for each object x_j and further construct Neg .

Finally, to learn from both labeled and unlabeled data, we formulate the problem as a multi-task learning problem by combining \mathcal{L}_{sup} and \mathcal{L}_{ss} to get a loss function:

$$\mathcal{L} = \mathcal{L}_{sup} + \lambda \mathcal{L}_{ss}, \quad (14)$$

where λ controls the relative importance of the two terms. Here, the self-supervised loss acts as a regularizer that learns from unlabeled data. The learned data-driven prior knowledge could provide supplementary information to the training set when the number of labeled objects is limited. This further enhances the model’s capability. The loss function can be optimized by stochastic gradient descent. To prevent overfitting, we further regularize all the weight matrices W_j ’s mentioned in Eqs. 4 - 13.

The major computation cost of ConCH comes from the updates of object and context embeddings. Let n be the number of objects to be classified. In each bipartite graph, there are at most kn contexts. Since the degree of each object is at most k and that of a context is at most 2, the total time complexity of embedding updates is $O(3knd_1d_2)$, where d_1 and d_2 are respectively the maximum numbers of rows and columns in the transformation matrices of the MLPs. Further, for each meta-path, we generate a “positive” bipartite graph and a “negative” one. Therefore, the total time complexity of embedding updates is $O(6knd_1d_2|\mathcal{PS}|)$, where $|\mathcal{PS}|$ is the number of meta-paths. Since the computation for each bipartite graph is independent from the others’, ConCH can be easily parallelized. Finally, we summarize ConCH in Alg. 1.

Algorithm 1 ConCH

Input: An HIN $\mathcal{G} = (\mathcal{V}, \mathcal{E})$; target object type T and object set \mathcal{X}_T ; a pre-defined meta-path set \mathcal{PS} ; object feature matrix X ; the number of layers L ;

Output: Final object embeddings $\{z_i\}_{i=1}^{|\mathcal{X}_T|}$

- 1: **for** \mathcal{P} in \mathcal{PS} **do**
 - 2: Calculate \mathcal{P} -induced *PathSim* scores based on Eq. 1
 - 3: Filter \mathcal{P} -based neighbors for $x_i, \forall x_i \in \mathcal{X}_T$
 - 4: Construct initial context features $o_{c_{\mathcal{P}}}$ based on Eq. 3
 - 5: Construct the bipartite graph $G_{\mathcal{P}}$
 - 6: **end for**
 - 7: \triangleright Lines 8-22 correspond to one epoch
 - 8: **for** \mathcal{P} in \mathcal{PS} **do**
 - 9: $\hat{X} = \text{shuffle}(X)$
 - 10: \triangleright Let f_i and \hat{f}_i denote the feature vectors of object x_i in X and \hat{X} , respectively
 - 11: $h_{x_i}^{\mathcal{P},(0)} \leftarrow f_i, \hat{h}_{x_i}^{\mathcal{P},(0)} \leftarrow \hat{f}_i, \forall x_i \in \mathcal{X}_T$
 - 12: $h_{c_j}^{\mathcal{P},(0)} \leftarrow o_{c_{\mathcal{P}}}, \hat{h}_{c_j}^{\mathcal{P},(0)} \leftarrow o_{c_{\mathcal{P}}}, \forall c_j \in V_C^{\mathcal{P}}$
 - 13: **for** $t \leftarrow 0$ to $L - 1$ **do**
 - 14: Update $h_{c_j}^{\mathcal{P},(t+1)}$ and $\hat{h}_{c_j}^{\mathcal{P},(t+1)}$ by Eq. 4, $\forall c_j \in V_C^{\mathcal{P}}$
 - 15: Update $h_{x_i}^{\mathcal{P},(t+1)}$ and $\hat{h}_{x_i}^{\mathcal{P},(t+1)}$ by Eq. 5, $\forall x_i \in \mathcal{X}_T$
 - 16: **end for**
 - 17: **end for**
 - 18: Compute z_i and \hat{z}_i by Eqs. 6 - 8, $\forall x_i \in \mathcal{X}_T$
 - 19: Construct the global vector s based on Eq. 11
 - 20: Construct the positive sample set $Pos = \{(z_i, s)\}_{i=1}^N$ and the negative sample set $Neg = \{(\hat{z}_j, s)\}_{j=1}^M$
 - 21: Construct $\mathcal{L}_{sup}, \mathcal{L}_{ss}$ and \mathcal{L}
 - 22: Optimize \mathcal{L} to update weight matrices W_j ’s
 - 23: **return** $\{z_i\}_{i=1}^{|\mathcal{X}_T|}$
-

V. EXPERIMENTAL EVALUATION

We evaluate ConCH’s effectiveness and efficiency. We compare ConCH with 15 other methods by their *Micro-F1* and *Macro-F1* scores. We also show results on meta-path weight learning and parameter analysis.

A. Datasets

We use three datasets: DBLP, Yelp, and Freebase. DBLP² is a bibliographic network of academic publications. Yelp³ contains information of businesses, such as user data and reviews. Freebase⁴ is a knowledge base of facts recorded with the RDF model. The three datasets are representative HINs. We define a classification task for each dataset as follows:

- **DBLP:** We extracted a subset from the *dblp-4area* dataset [60], which contains 4,057 authors (A), 14,376 papers (P) and 20 conferences (C). Links include A-P (an author publishes a paper) and P-C (a paper is published at a conference). Authors are classified into one of four research areas: database (DB), data mining (DM), machine learning (ML), and information retrieval (IR). Following [61], we compute, for each author, a 300-dimension attribute vector by averaging the word embeddings⁵ of keyword terms of his/her published papers. We consider the meta-path set {APA, APAPA, APCPA}. Our task is to classify authors into their research areas. The ground truth labels of authors are given in *dblp-4area* data, which labels each author by his/her primary research area.

- **Yelp-Restaurant:** The dataset is about restaurant businesses in Yelp. It contains 2,614 businesses (B), 33,360 reviews (R), 1,286 users (U) and 82 food relevant keywords (K). Links include B-R (a business receives a review), U-R (a user writes a review) and K-R (a keyword is included in a review). We classify restaurants by the types of food they provide: *Fast Food*, *Sushi Bars*, or *American New Food*. Each restaurant is associated with two categorical attributes: reservation (whether reservation is required) and service (whether waiter service is provided). We use the meta-paths {BRURB, BRKRB}. The ground truth labels of restaurants are given in Yelp.

- **Freebase-Movie:** We constructed a dataset on movies from Freebase that contains 3,492 movies (M), 33,401 actors (A), 2,502 directors (D) and 4,459 producers (P). Links include M-A (movie and its actor), M-D (movie and its director) and M-P (movie and its producer). Our task is to classify movies by genre, which includes *Action*, *Comedy* and *Drama*. We encode a one-hot feature vector for each movie. In our experiments, we consider meta-paths {MAM, MDM, MPM}. We take the given genres of movies as the ground truth.

B. Algorithms for comparison

We first compare ConCH with 11 other methods, which can be grouped into two categories. Readers are referred to Section II for more details.

²<https://dblp.uni-trier.de/>

³https://www.yelp.com/academic_dataset/

⁴<https://www.freebase.com/>

⁵<http://nlp.stanford.edu/data/glove.840B.300d.zip>

(Network-embedding-based methods): **node2vec** and **meta-path2vec** (or **mp2vec** for short) are two network embedding methods that generate objects’ embedding vectors. These vectors are fed into a logistic regression classifier to predict objects’ labels. Since **node2vec** is designed for homogeneous networks, we apply it to an HIN by ignoring the heterogeneity of the network. For **metapath2vec**, we test all the given meta-paths and report the best result.

(Graph-neural-network-based methods): **GCN** is a model that extends the convolution operation on graphs. **GAT** further integrates the attention mechanism in the convolutional layer. **MVGRL** is a self-supervised learning method. All the methods are designed for homogeneous networks. We apply them by converting an HIN to a homogeneous network using meta-paths and report the best result. **HAN**, **HetGNN**, **MAGNN**, **HGT**, **HDGI** and **HGCN** are state-of-the-art heterogeneous neural network models. Note that **HAN**, **MAGNN** and **HDGI** utilize meta-paths while others not. In particular, **HDGI** uses **HAN** as the base neural network model. For **MVGRL**, **HetGNN** and **HDGI**, since they are unsupervised, the learned object embeddings are fed into a logistic regression classifier to predict objects’ labels.

C. Experimental setup

We implement ConCH using PyTorch. The model is initialized by Glorot initialization [62] and trained by Adam [63]. For ConCH and all its variants, we set the learning rate to 0.001, the dropout ratio to 0.5 and the penalty weight on the ℓ_2 -norm regularizer to 0.0005 on all three datasets. We set k in neighbor filtering to 5, 10, 10 and the number of layers L to 2, 1, 1 in DBLP, Yelp and Freebase, respectively. For the self-supervised regularization weight λ , we fine tune it from {0.001, 0.01, 0.1, 1}. We use early stopping with a patience of 100, i.e., we stop training if the validation accuracy does not decrease for 100 consecutive epochs. To run ConCH, we first run **metapath2vec** with the default parameter settings in the original codes provided by the authors to construct context features. We set the initial context embedding dimensionality to 128 for DBLP and Yelp, and 32 for Freebase. We perform neighbor filtering and context feature extraction as preprocessing steps. For fair comparison, we set the output embedding dimensionality for all the methods to be 128. We run all the experiments on a server with 128G memory and a single NVIDIA 2080Ti GPU. We also feed all the methods the same training/validation/test set splits. For all the baseline methods, we use the original codes released by their authors. For **MVGRL**, **HetGNN** and **HDGI**, since they are unsupervised, we train the models by the train loss early stopping with 100 epochs. For other graph neural network models, we use the same early stopping criteria as ConCH for fairness. For these models, we fine tune the model hyper-parameters by grid search. These include the number of model layers {1, 2, 3} and the learning rate {0.01, 0.001, 0.0001}. For **HGCN**, as suggested in the original paper, we further perform the hyper-parameter tuning for the number of inception layers {1, 2, 3, 4}, the convolutional kernel size {1, 2, 3, 4}

and the order of label propagation $\{0, 1, 2\}$. For all the meta-path-based methods, we employ the same meta-path sets. For the random-walk-based embedding methods, we use the default parameters reported in their original papers. For each method, we run experiments 10 times and report the average results. We provide our code and data here: <https://github.com/dingdanhao110/Conch>.

D. Performance results

Table I summarizes the performance results. We evaluate the methods on three datasets under two evaluation metrics. Moreover, to compare the methods when labeled objects are scarce, we vary the training set size from very small (2% and 5%) to moderate (10% and 20%). There are thus in total $3 \times 2 \times 4 = 24$ “contests” in the comparison study. Each row in the table corresponds to one contest. For each contest, we highlight the winner’s score in bold. From the table, we make the following observations:

(1) As the training set size decreases, the performance of many baseline methods drops sharply. For example, on DBLP, the Micro-F1 score of HDGI is 0.9233 with 20% labeled objects, which is close to the winner’s score. However, with only 2% labeled objects, the score of HDGI degrades to 0.7733 while the best score (ConCH) is 0.9429.

(2) For the two network-embedding-based methods, mp2vec performs better on DBLP and Yelp, while node2vec takes the lead on Freebase. Although mp2vec is an embedding method for HINs that uses meta-paths, it can take only one meta-path as input, which affects its performance.

(3) ConCH outperforms graph neural network models GCN and GAT. This is because both methods are designed for homogeneous networks without considering objects’ and links’ heterogeneity. While MVGRL leverages self-supervision, it ignores heterogeneity and is outperformed by ConCH.

(4) Although HAN is designed for HINs and it uses meta-paths, HAN omits mp-contexts, which lowers its effectiveness. For example, its Macro-F1 score on Yelp with 20% labeled data is 0.6637, which is significantly lower than that of ConCH. This shows the importance of mp-contexts. Moreover, since HDGI takes HAN as the base neural network model, its performance is also adversely affected due to the ignorance of mp-contexts. MAGNN considers mp-contexts and it performs well on DBLP. However, as we will show in Sec. V-G, MAGNN is computationally expensive. Further, MAGNN needs to preprocess and maintain information for each meta-path instance, which requires a lot of storage. For example, the meta-paths processing in the Yelp task generates a large number of path instances between objects, which causes out-of-memory errors. As a result, in our experiments, MAGNN fails to run on Yelp.

(5) HetGNN and HGT can achieve overall good performance over all three datasets. However, they are also easily affected when the number of labeled objects is limited. For example, with 2% labeled objects, the Macro-F1 scores of HetGNN and HGT on Yelp are 0.8486 and 0.7836, respectively, while the winner’s is 0.8860. HGCN constructs relational

features for an object by aggregating its neighbors’ label information. However, the constructed relational features and the original features of the object could be in different feature space, which restricts the model’s effectiveness.

(6) ConCH achieves the best performance in all 24 cases. ConCH uses a simple neighbor filtering scheme to remove the less relevant neighbors of an object. Moreover, ConCH, based on meta-paths, leverages mp-contexts to boost the classification performance. We further observe that, given scarce labeled objects, ConCH achieves superior performance compared with other methods. For example, with 2% labeled objects, the Micro-F1 score of ConCH on Yelp is 0.8841 while that of the first runner-up is only 0.8558. This is because the self-supervision learned from unlabeled data provides additional information that helps classify objects.

E. Ablation study

We conduct an ablation study on ConCH to understand the characteristics of its main components. One variant updates objects’ embeddings by directly aggregating information from its multi-hop neighbors specified by meta-paths without considering contexts derived from the path instances. This helps us understand the importance of including mp-contexts in object classification. We call this variant **ConCH_nc** (no contexts). Another variant randomly selects k meta-path-based neighbors from a given object as *relevant neighbors*. This is in contrast to ConCH, which selects relevant neighbors by picking the top- k most similar path-based neighbors based on PathSim scores (see Section IV-A). We call this variant **ConCH_rd** (random), which helps us evaluate the effectiveness of our neighbor-filtering strategy. To show the importance of the self-supervised regularization, we train the model with \mathcal{L}_{sup} only and call this variant **ConCH_su** (supervised). Moreover, while we formulate our problem as a multi-task learning objective, we can also train the model by a *pre-training & fine-tuning* strategy. Specifically, we first train the model by using \mathcal{L}_{ss} only and then take the learned parameters as the initialization to the model that uses \mathcal{L}_{sup} only. This variant can be used to show the advantage of the multi-task learning framework and we call it **ConCH_ft** (fine-tuning). Finally, we consider a variant of ConCH that assigns meta-paths equal weights without the attention mechanism. We call this variant **ConCH_ew** (equal weights).

Similar to Sec. V-D, we compare ConCH with these variants for four training set sizes on three datasets w.r.t. Macro-F1 and Micro-F1 measures. The results are given in Figs. 3 - 5. From these figures, we observe:

(1) ConCH beats ConCH_nc in all 24 cases. In particular, ConCH significantly outperforms ConCH_nc on Yelp and Freebase. This shows that mp-context information is particularly important for HINs that are with heterogeneity in object types and link types. When using meta-paths, the inclusion of mp-contexts is essential for effective classification.

(2) ConCH achieves better performance than ConCH_rd. Since ConCH_rd randomly selects k path-based neighbors for an object, the performance gaps between ConCH and

TABLE I
THE CLASSIFICATION RESULTS (%) OVER THE METHODS

Datasets	Metrics	Training	node2vec	mp2vec	GCN	GAT	MVGRL	HAN	HetGNN	MAGNN	HGT	HDGI	HGCN	ConCH
DBLP	Macro-F1	2%	90.60	91.16	91.15	86.73	91.43	80.65	89.92	92.35	89.97	67.71	82.36	93.86
		5%	92.29	92.46	91.86	90.59	91.66	87.74	91.99	93.01	92.00	85.68	86.69	94.22
		10%	92.08	92.48	92.20	91.19	91.99	90.98	92.50	93.68	92.32	90.37	88.36	93.77
		20%	92.08	92.41	92.46	91.23	92.36	92.51	92.86	94.17	93.14	91.60	86.81	94.29
	Micro-F1	2%	91.54	91.88	91.94	88.27	92.01	82.30	90.92	92.99	90.87	77.33	85.46	94.29
		5%	92.82	93.00	92.56	91.20	92.26	88.38	92.59	93.58	92.62	88.10	88.65	94.64
		10%	92.62	93.04	92.88	91.77	92.57	91.49	93.01	94.13	92.90	91.35	90.09	94.22
		20%	92.65	92.95	93.10	91.78	92.89	92.91	93.36	94.56	93.65	92.33	88.67	94.70
Yelp	Macro-F1	2%	75.54	84.95	84.80	70.34	53.77	58.12	84.86	-	78.36	51.63	59.07	88.60
		5%	86.61	89.44	85.96	80.50	53.82	61.55	88.48	-	88.64	56.83	62.94	90.11
		10%	88.65	90.15	86.34	80.78	54.09	61.42	89.45	-	90.00	63.36	73.86	91.31
		20%	89.62	90.57	86.39	81.95	55.91	66.37	89.85	-	90.32	72.59	76.61	92.11
	Micro-F1	2%	78.28	85.58	83.17	76.67	72.42	70.89	85.25	-	80.78	65.53	70.03	88.41
		5%	85.43	88.31	84.09	81.30	73.15	72.71	87.51	-	87.66	68.80	72.40	89.69
		10%	87.50	89.10	84.53	81.64	73.22	73.86	88.56	-	89.10	72.64	78.58	90.78
		20%	88.45	89.56	84.58	82.16	73.33	74.95	88.96	-	89.48	77.51	79.73	91.56
Freebase	Macro-F1	2%	52.81	52.08	52.35	47.72	50.77	51.39	56.32	51.54	53.20	52.69	47.73	56.46
		5%	54.27	52.68	55.28	48.46	50.79	54.57	59.74	57.30	57.00	57.00	51.65	61.07
		10%	55.38	54.26	59.08	50.00	51.14	57.96	62.15	59.50	60.39	59.93	56.69	63.35
		20%	57.35	57.50	60.90	51.09	51.15	60.11	62.92	62.23	61.76	61.39	60.35	64.75
	Micro-F1	2%	59.71	59.01	64.89	64.69	62.75	65.59	63.39	63.31	64.06	63.86	55.43	65.82
		5%	59.67	57.83	66.00	65.53	64.44	65.83	65.40	65.65	66.09	65.41	58.79	67.55
		10%	60.19	58.91	67.65	65.86	65.52	66.79	67.41	66.93	67.69	66.82	63.03	69.21
		20%	61.95	61.84	68.49	66.38	65.92	67.32	68.05	68.41	68.85	67.55	66.75	70.04

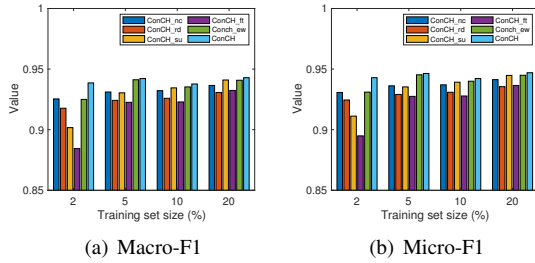


Fig. 3. The ablation study results on DBLP




Fig. 4. The ablation study results on Yelp

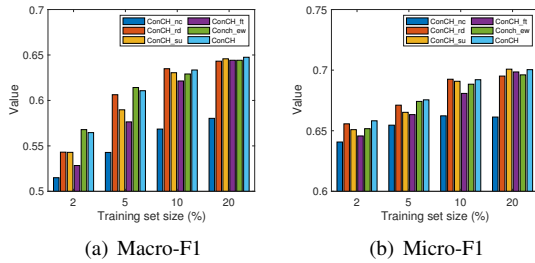


Fig. 5. The ablation study results on Freebase

ConCH_rd show that ConCH’s top- k neighbor-filtering strategy is very effective in selecting more influential path-based neighbors to improve classification accuracy.

(3) Given 20% labeled objects, ConCH_su achieves comparable performance with ConCH. However, as the training set size decreases, the performance gap between ConCH and ConCH_su gets larger. This further shows the importance of leveraging self-supervised learning to cover the shortage of labeled objects.

(4) ConCH clearly outperforms ConCH_ft on three datasets. ConCH_ft, which switches the objective function from \mathcal{L}_{ss} in pre-training to \mathcal{L}_{sup} in fine-tuning, solves two optimization problems. However, ConCH unifies the two loss functions into one objective that is much easier to be optimized and can better leverage both labeled and unlabeled objects. Our results are similar to those reported in [42].

(5) ConCH performs better than ConCH_ew in 22 cases. This further shows the importance of the attention mechanism in learning meta-path weights.

F. Attention weight learning

ConCH learns the importance of meta-paths by the attention mechanism. To show the effectiveness of attention weight learning, we compute the average learned weights of meta-paths for all three datasets. We show the results (Fig. 6) with 20% labeled objects as training sets for illustration.

Fig. 6(a) shows the learned meta-paths’ weights for the DBLP dataset. Recall that the classification task is to classify authors by their research areas. From the figure, we see that the weight of the meta-path APCPA (authors that publish papers in the same venue) is almost 1, while that of meta-paths APA (co-authorship) and APAPA (authors that share a common co-author) are very close to 0. An almost-0 weight for APA may seem counter-intuitive considering that co-authorship is a strong signal that two authors work in the same area. The




Fig. 6. Learned weights of meta-paths

reason why APA is given such a low weight is that it is a sparse relation. An author typically co-authors with only a handful of others in the community. Further, an author is related to a large number of other authors by APCPA (conference co-attendee relation). Moreover, authors related by APA are also related by APCPA, hence, the former is subsumed by the latter. In this example, we see that ConCH correctly selects the relevant meta-path APCPA over APA and APAPA.

Fig. 6(b) shows the learned meta-paths’ weights for the Yelp dataset. Recall that the task is to classify restaurants by their food categories. We see that the meta-path BRKRB (restaurants whose reviews contain the same food keyword) is given a much larger weight than BRURB (restaurants that are visited by the same customer). This is reasonable because food keywords in reviews directly indicate food category. On the other hand, a customer could visit restaurants that serve different categories of food.

The Freebase task is to classify movies by genre. Fig. 6(c) shows that all three meta-paths are useful. The weights of meta-paths MAM (movies share the same actor) and MDM (movies filmed by the same director) are a bit larger than that of MPM (movies produced by the same producer). From our discussion, we see that ConCH’s meta-path weight learning strategy is highly effective.

G. Efficiency study

In this section we study ConCH’s efficiency. For fairness, we compare the training time for ConCH, HAN, MAGNN, HGT and HGCN, as they are semi-supervised classification methods in HINs. For all these methods, we use the same training/validation set and run the experiments for 300 epochs. We take 20% labeled objects as training set for illustration. We repeat all the experiments three times and show the average training time of these methods w.r.t. the Micro-F1 score on the validation set in Fig. 7. Note that we cannot run MAGNN on Yelp. To recap the major differences of these methods: ConCH, HAN and MAGNN are based on meta-paths. However, ConCH selects the k most informative path-based neighbors for an object based on a filtering scheme. HAN uses all meta-path-based neighbors (i.e., no filtering) and it learns the neighbors’ relative importance by the vanilla attention mechanism. HAN does not use mp-contexts either. Different from ConCH, which uses mp-contexts to construct high-level context objects with features, MAGNN utilizes mp-contexts in a fine-grained manner by independently considering all the meta-path instances. It learns the importance of meta-path instances and further

aggregates information from these path instances to generate object embeddings. HGT employs a Transformer based aggregator to aggregate information from multi-type neighbors for an object, while HGCN leverages multiple kernels that use different convolutional filters.

From Figs. 7(a)-(c), we can see that ConCH consistently converges fast to the best results over all three datasets. MAGNN and HGT can also reach high performance scores, however, they need much longer time to converge. In particular, for MAGNN, when meta-paths are of longer lengths, the number of path instances could be numerous, which adversely affects model efficiency. For example, ConCH is about $50\times$ faster than MAGNN on DBLP; ConCH achieves almost $40\times$ speedup than HGT on Yelp. Although HAN and HGCN are faster than MAGNN and HGT, their performances are poor. These results show that our approach ConCH is very efficient and also highly effective. Finally, the runtime of ConCH depends on k , which is the number of selected relevant neighbors for an object. Fig. 7(d) shows how ConCH’s runtime of one training epoch varies with k . From the figure, we see that the runtime increases fairly linearly with k . Thus, ConCH is scalable w.r.t. k .

To further study ConCH’s efficiency, we conduct experiments on a large dblp-4area subgraph extracted from the AMiner citation network⁶. We first extract papers whose fields of study in the original dataset contain at least one of four areas: database (DB), data mining (DM), machine learning (ML), and information retrieval (IR). For simplicity, we only consider conference papers. We extract authors and conferences related to these papers. The resulting dataset contains 416,554 papers (P), 537,435 authors (A) and 2,649 conferences (C). Links include A-P (an author publishes a paper) and P-C (a paper is published at a conference). Our task is to classify papers into one of the four research areas. For each paper, following [61], we compute a 300-dimension attribute vector by averaging the word embeddings of keywords it contains. We consider the meta-path set {PAP, PCP}. The ground truth labels of papers are derived from the original dataset, where we label each paper by the research area with the largest field of study weight.

Figure 8 shows the efficiency study on the AMiner dataset. All these methods are semi-supervised classification methods in HINs. Compared with others, ConCH converges fast to the best performance. Table II further summarizes the classification performance over all the methods on AMiner. ConCH clearly outperforms all the baseline methods, which again shows its effectiveness. Due to the large dataset size, both MVGRL and MAGNN cause out-of-memory errors. Note that MVGRL requires both adjacency matrix and diffusion matrix as input to provide both local and global views of a graph structure, respectively. However, the diffusion matrix computed by all the approaches suggested in the original paper is too dense, which causes out-of-memory errors. This further shows that ConCH is efficient and also very effective.

⁶<https://originalstatic.aminer.cn/misc/dblp.v12.7z>




Fig. 7. The results of efficiency study. Specifically, (a)-(c) compare the training time vs. the Micro-F1 score on three datasets; (d) shows the runtime of one training epoch in ConCH vs. k on three datasets

TABLE II
THE CLASSIFICATION RESULTS (%) OVER ALL THE METHODS ON AMINER

Datasets	Metrics	Training	node2vec	mp2vec	GCN	GAT	MVGRL	HAN	HetGNN	MAGNN	HGT	HDGI	HGCN	ConCH
AMiner	Macro-F1	2%	57.06	33.89	59.91	68.25		71.47	66.49		68.44	72.38	52.23	73.10
		5%	57.19	33.91	60.24	70.66		73.87	66.84		69.01	73.24	49.39	76.05
		10%	57.25	34.00	60.41	71.60	-	74.70	66.96	-	69.37	73.62	51.27	77.11
		20%	57.26	38.33	60.50	72.00		75.44	67.07		69.72	73.87	53.94	78.33
	Micro-F1	2%	65.77	52.68	66.83	72.79		75.79	72.07		72.50	76.88	63.96	77.54
		5%	65.89	52.72	67.11	74.91		77.92	72.35		73.18	77.67	64.58	79.95
		10%	65.95	52.71	67.29	75.84	-	78.66	72.41	-	73.56	77.92	65.28	80.85
		20%	65.95	54.57	67.37	76.22		79.16	72.46		73.92	78.09	65.16	81.84




Fig. 8. Runtime comparison on AMiner

H. Parameter analysis

We end this section with a sensitivity analysis on the hyper-parameters of ConCH. In particular, we study four key hyper-parameters: the output embedding dimensionality, the input context embedding dimensionality, the number of selected relevant neighbors k and the self-supervised regularization weight λ . In our experiments, we vary one parameter each time with others fixed. Fig. 9 illustrates the results with 20% labeled objects w.r.t. the Micro-F1 scores. (Results on Macro-F1 scores exhibit similar trends, and thus are omitted for space reasons.) From the figure, we see that

- (1) As the output embedding dimensionality increases, ConCH achieves better performance. This is because when the dimensionality is small, node embeddings cannot capture enough information for classifying objects.
- (2) There is a performance drop in Freebase when the input context feature dimensionality is 128. This shows that initial context embedding vectors in a large dimensionality could contain noise that adversely affects the classification accuracy.
- (3) For the other two hyper-parameters, ConCH gives very stable performances over a wide range of parameter values.

In particular, ConCH performs well even with small k . This further shows that ConCH is an effective and efficient method for solving the classification problem on HINs.

VI. CONCLUSIONS

In this paper we studied classification in HINs and proposed ConCH, which is a graph neural network model based on meta-paths. ConCH formulates the classification problem as a multi-task learning problem that combines semi-supervised learning with self-supervised learning to enhance the model performance when the training set size is small. Further, it leverages meta-path-based contexts, which capture specific details of meta-path instances, to improve classification accuracy. It filters less relevant neighbors of an object by selecting the top- k neighbors using PathSim scores. This approach helps reduce the number of neighbors whose information is aggregated to derive an object’s embedding. We conducted extensive experiments and compared ConCH with 15 other methods. Our analysis shows that ConCH can achieve superior classification performance with high efficiency.

ACKNOWLEDGMENT

This work is supported by Hong Kong Research Grants Council grant 17253616, and partially funded by Greek national funds, under the call Research-Create-Innovate (project codes: T1EDK-04810 and T2EDK-02848), NSF III-1705169, NSF CAREER Award 1741634, NSF #1937599, DARPA HR00112090027, Okawa Foundation Grant, and Amazon Research Award.

REFERENCES

- [1] X. Li *et al.*, “On transductive classification in heterogeneous information networks,” in *CIKM*, 2016.
- [2] M. Ji *et al.*, “Graph regularized transductive classification on heterogeneous information networks,” in *ECML-PKDD*, 2010.




Fig. 9. Parameter analysis

- [3] M. Wan *et al.*, “Graph regularized meta-path based transductive regression in heterogeneous information network,” in *SDM*, 2015.
- [4] T. N. Kipf and M. Welling, “Semi-supervised classification with graph convolutional networks,” in *ICLR*, 2017.
- [5] F. Xu *et al.*, “Relation-aware graph convolutional networks for agent-initiated social e-commerce recommendation,” in *CIKM*, 2019.
- [6] X. Wang *et al.*, “Heterogeneous graph attention network,” in *WWW*, 2019.
- [7] C. Zhang *et al.*, “Heterogeneous graph neural network,” in *KDD*, 2019.
- [8] Z. Hu *et al.*, “Heterogeneous graph transformer,” in *WWW*, 2020.
- [9] Z. Zhu *et al.*, “Hgen: A heterogeneous graph convolutional network-based deep learning model toward collective classification,” in *KDD*, 2020.
- [10] X. Fu *et al.*, “Magnn: Metapath aggregated graph neural network for heterogeneous graph embedding,” in *WebConf*, 2020.
- [11] A. Vaswani *et al.*, “Attention is all you need,” in *NeurIPS*, 2017.
- [12] B. Hu *et al.*, “Leveraging meta-path based context for top-n recommendation with a neural co-attention model,” in *KDD*, 2018.
- [13] Y. Dong *et al.*, “Heterogeneous network representation learning,” in *IJCAI*, 2020.
- [14] D. Zhou *et al.*, “Learning with local and global consistency,” in *NeurIPS*, 2004.
- [15] M. Belkin *et al.*, “Manifold regularization: A geometric framework for learning from labeled and unlabeled examples,” *JMLR*, 2006.
- [16] X. Zhu *et al.*, “Semi-supervised learning using gaussian fields and harmonic functions,” in *ICML*, 2003.
- [17] H. Jiang *et al.*, “Semi-supervised learning over heterogeneous information networks by ensemble of meta-graph guided random walks,” in *IJCAI*, 2017.
- [18] C. Luo *et al.*, “Hetpathmine: A novel transductive classification algorithm on heterogeneous information networks,” in *ECIR*, 2014.
- [19] D. Wang *et al.*, “Structural deep network embedding,” in *KDD*, 2016.
- [20] M. Ou *et al.*, “Asymmetric transitivity preserving graph embedding,” in *KDD*, 2016.
- [21] Z. Liu *et al.*, “Semantic proximity search on heterogeneous graph by proximity embedding,” in *AAAI*, 2017.
- [22] B. Perozzi *et al.*, “Deepwalk: online learning of social representations,” in *KDD*, 2014.
- [23] A. Grover and J. Leskovec, “node2vec: Scalable feature learning for networks,” in *KDD*, 2016.
- [24] T. Mikolov *et al.*, “Distributed representations of words and phrases and their compositionality,” in *NeurIPS*, 2013.
- [25] L. Liao *et al.*, “Attributed social network embedding,” *TKDE*, 2018.
- [26] Z. Yang *et al.*, “Revisiting semi-supervised learning with graph embeddings,” in *ICML*, vol. 48, 2016, pp. 40–48.
- [27] J. Qiu *et al.*, “Network embedding as matrix factorization: Unifying deepwalk, line, pte, and node2vec,” in *WSDM*, 2018.
- [28] J. Tang *et al.*, “Line: Large-scale information network embedding,” in *WWW*, 2015.
- [29] Y. Shi *et al.*, “Easing embedding learning by comprehensive transcription of heterogeneous information networks,” in *KDD*, 2018.
- [30] Y. Shi, H. Gui *et al.*, “Aspem: Embedding learning by aspects in heterogeneous information networks,” in *SDM*, 2018.
- [31] S. Chang *et al.*, “Heterogeneous network embedding via deep architectures,” in *KDD*, 2015.
- [32] Y. Dong *et al.*, “metapath2vec: Scalable representation learning for heterogeneous networks,” in *KDD*, 2017.
- [33] T.-y. Fu *et al.*, “Hin2vec: Explore meta-paths in heterogeneous information networks for representation learning,” in *CIKM*, 2017.
- [34] T. Chen and Y. Sun, “Task-guided and path-augmented heterogeneous network embedding for author identification,” in *WSDM*, 2017.
- [35] J. Tang *et al.*, “Pte: Predictive text embedding through large-scale heterogeneous text networks,” in *KDD*, 2015.
- [36] K. He *et al.*, “Momentum contrast for unsupervised visual representation learning,” in *CVPR*, 2020.
- [37] T. Chen *et al.*, “A simple framework for contrastive learning of visual representations,” *arXiv preprint arXiv:2002.05709*, 2020.
- [38] Z. Wu *et al.*, “Unsupervised feature learning via non-parametric instance discrimination,” in *CVPR*, 2018.
- [39] J. Devlin *et al.*, “Bert: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
- [40] Z. Yang *et al.*, “Xlnet: Generalized autoregressive pretraining for language understanding,” in *NeurIPS*, 2019.
- [41] A. Radford *et al.*, “Language models are unsupervised multitask learners,” *OpenAI Blog*, 2019.
- [42] Y. You *et al.*, “When does self-supervision help graph convolutional networks?” *arXiv preprint arXiv:2006.09136*, 2020.
- [43] Z. Hu *et al.*, “Gpt-gnn: Generative pre-training of graph neural networks,” in *KDD*, 2020.
- [44] J. Qiu *et al.*, “Gcc: Graph contrastive coding for graph neural network pre-training,” in *KDD*, 2020.
- [45] P. Velickovic *et al.*, “Deep graph infomax,” in *ICLR*, 2019.
- [46] F.-Y. Sun *et al.*, “Infograph: Unsupervised and semi-supervised graph-level representation learning via mutual information maximization,” *arXiv preprint arXiv:1908.01000*, 2019.
- [47] W. Hu *et al.*, “Strategies for pre-training graph neural networks,” *arXiv preprint arXiv:1905.12265*, 2019.
- [48] K. Hassani and A. H. Khasahmadi, “Contrastive multi-view representation learning on graphs,” *arXiv preprint arXiv:2006.05582*, 2020.
- [49] Y. Ren *et al.*, “Heterogeneous deep graph infomax,” *arXiv preprint arXiv:1911.08538*, 2019.
- [50] J. Bruna *et al.*, “Spectral networks and locally connected networks on graphs,” in *ICLR*, 2014.
- [51] M. Defferrard *et al.*, “Convolutional neural networks on graphs with fast localized spectral filtering,” in *NeurIPS*, 2016.
- [52] W. L. Hamilton *et al.*, “Inductive representation learning on large graphs,” in *NeurIPS*, 2017.
- [53] P. Velickovic *et al.*, “Graph attention networks,” in *ICLR*, 2018.
- [54] Y. Zhang *et al.*, “Deep collective classification in heterogeneous information networks,” in *WebConf*, 2018.
- [55] Y. Cen *et al.*, “Representation learning for attributed multiplex heterogeneous network,” in *KDD*, 2019.
- [56] S. Yun *et al.*, “Graph transformer networks,” in *NeurIPS*, 2019.
- [57] S. Yang *et al.*, “Domain adaptive classification on heterogeneous information networks,” in *IJCAI*, 2020.
- [58] Y. Sun *et al.*, “Pathsim: Meta path-based top-k similarity search in heterogeneous information networks,” *PVLDB*, 2011.
- [59] X. Li *et al.*, “Semi-supervised clustering in attributed heterogeneous information networks,” in *WWW*, 2017.
- [60] Y. Sun *et al.*, “Ranking-based clustering of heterogeneous information networks with star network schema,” in *KDD*, 2009.
- [61] C. Yang *et al.*, “Relation learning on social networks with multi-modal graph edge variational autoencoders,” in *WSDM*, 2020.
- [62] X. Glorot and Y. Bengio, “Understanding the difficulty of training deep feedforward neural networks,” in *AISTATS*, 2010.
- [63] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *ICLR*, 2015.