

Uniting Heterogeneity, Inductiveness, and Efficiency for Graph Representation Learning

Tong Chen, Hongzhi Yin, Jie Ren, Zi Huang, Xiangliang Zhang, Hao Wang

Abstract—With the ubiquitous graph-structured data in various applications, models that can learn compact but expressive vector representations of nodes have become highly desirable. Recently, bearing the message passing paradigm, graph neural networks (GNNs) have greatly advanced the performance of node representation learning on graphs. However, a majority class of GNNs are only designed for homogeneous graphs, leading to inferior adaptivity to the more informative heterogeneous graphs with various types of nodes and edges. Also, despite the necessity of inductively producing representations for completely new nodes (e.g., in streaming scenarios), few heterogeneous GNNs can bypass the transductive learning scheme where all nodes must be known during training. Furthermore, the training efficiency of most heterogeneous GNNs has been hindered by their sophisticated designs for extracting the semantics associated with each meta path or relation. In this paper, we propose **wide** and **deep** message passing network (WIDEN) to cope with the aforementioned problems about heterogeneity, inductiveness, and efficiency that are rarely investigated together in graph representation learning. In WIDEN, we propose a novel inductive, meta path-free message passing scheme that packs up heterogeneous node features with their associated edges from both low- and high-order neighbor nodes. To further improve the training efficiency, we innovatively present an active downsampling strategy that drops unimportant neighbor nodes to facilitate faster information propagation. Experiments on three real-world heterogeneous graphs have further validated the efficacy of WIDEN on both transductive and inductive node representation learning, as well as the superior training efficiency against state-of-the-art baselines.



1 INTRODUCTION

NOWADAYS a considerable amount of real-life applications constantly involve graph-structured data, e.g., social networks, citation records, and power grids. To fully analyse the information within large-scale graphs, learning low-dimensional representations (a.k.a. embeddings) of nodes has been a well-established solution, where high-quality node embeddings are required to retain maximal information of nodes' own features and their connectivity within the graph. With the informative yet compact node embeddings, various downstream tasks like node classification and link prediction can be accurately performed.

With the widely proven effectiveness of deep neural networks, their adaptation to graph-structured data has been a great success as well. Thus, graph neural networks (GNNs) have demonstrated state-of-the-art effectiveness in a wide range of graph mining applications, such as personalized recommendation [1] and molecular fingerprint mining [2]. The rationale behind GNNs is based on the message passing scheme between different nodes in a graph. For example, the graph convolutional network (GCN) [3] extends the convolutional operation to the non-Euclidean graph-structured data, and uses the connectivity of the graph as the filter to perform neighborhood information aggregation. At the same time, the graph attention network

(GAT) [4] utilizes neighborhood sampling to simultaneously learn the topological structure of each node's neighborhood and the distribution of node features in a graph. Despite the prosperity of GNNs, a large body of GNN models are exclusively designed for learning node representations on homogeneous graphs [5] with the assumption that a graph contains only single-typed nodes and edges. When handling heterogeneous graphs (e.g., user-item graphs in e-commerce) with various types of nodes and edges, such limitation hinders a model's capability of fully capturing the diverse semantics associated with the graph's heterogeneity. Consequently, GNNs for heterogeneous graphs are introduced in recent research [5]–[8].

Essentially, to account for the diverse semantics among different types nodes and edges, a general remedy in heterogeneous GNNs is to utilize the notion of meta-paths [6], [9], [10], which are node sequences connected with heterogeneous edges, e.g., the *author-paper-conferece* meta paths in a citation graph. As a common practice, by pre-defining some specific meta path compositions, a heterogeneous graph can be decomposed into several collections of meta-paths. Then, homogeneous GNNs can thus be applied to each collection for node representation learning, such as HAN [6] and GTN [5] that respectively extend GAT and GCN to different meta paths. However, a widely acknowledged drawback of meta path-based GNNs is that, the design of meta paths are usually task-specific and requires abundant domain knowledge [5], [7], making it hard for those methods to generalize across different heterogeneous graphs. Moreover, learning node representations purely from pre-defined meta paths easily makes a model overlook the important first-order (i.e., one-hop) neighbor nodes, since meta paths are more concerned on nodes connected in two hops or further distances [11]. For instance, when learning a

- T. Chen, H. Yin, J. Ren and Z. Huang are with the School of Information Technology and Electrical Engineering, The University of Queensland. E-mail: tong.chen@uq.edu.au, h.yin1@uq.edu.au, thurenjie@foxmail.com, huang@itee.uq.edu.au
- X. Zhang is with the Machine Intelligence and Knowledge Engineering Laboratory, King Abdullah University of Science and Technology. E-mail: xiangliang.zhang@kaust.edu.sa
- H. Wang is with Alibaba AI Labs. E-mail: cashenry@126.com

Manuscript is under review. Hongzhi Yin is the corresponding author.

user’s preference for recommendation, all her/his interacted item nodes carry strong indicative signals on the user’s personal interests. Unfortunately, such information has to be sacrificed in a meta path-based learning scheme, leading to considerable amount of information loss in the produced node embeddings.

In this regard, another popular solution for heterogeneous graph representation learning is built upon the neighborhood message passing paradigm as in homogeneous GNNs, while infusing additional information about the heterogeneity of the graph. For example, given all sampled neighbors of a target node, both RGCN [12] and HGT [7] design a type-specific aggregation scheme to firstly aggregate the information within each type of neighbor nodes, then pass all summarized messages to the target node via a second-tier aggregator. Though promising performance is reported, such neighborhood-based aggregation scheme is highly vulnerable to sparse graphs where most nodes are loosely connected. Take the widely used user-item recommendation graphs (e.g., Amazon [13] and Yelp [14]) as an example, the average degree of each user node is commonly below 5, leading to a limited amount of messages available for learning user node embeddings. Though methods based on GraphSAGE [1], [15], [16] can grasp extra information from high-order neighbor nodes by increasing the neighborhood sampling depth, it brings an inevitable by-product of recursively increasing computational cost and memory footprint [17] to cope with the expanded neighborhood.

Meanwhile, a large body of previous works [3]–[5] on GNNs are dedicated to embedding nodes in a fixed graph, e.g., GCN-based approaches [3] require a full adjacency matrix of the graph. That is to say, they are based on a transductive learning paradigm, where all nodes must be present during training to facilitate downstream tasks. However, in real-world scenarios, many applications require embeddings to be quickly generated for unseen nodes. Such capability is essential for high-throughput, production machine learning systems [15], most of which involve heterogeneous graphs [8], [16], [18]. The transductive learning scheme hurts the generalizability of those methods as they have to be retained once the graph is updated. Hence, the inductive learning capability is highly desirable for heterogeneous GNNs, which benefits representation learning on evolving graphs that constantly encounter unseen nodes, e.g., new users and videos on Youtube. Recently, attempts have been made to account for both inductiveness and heterogeneity in GNNs [6], [7], but those methods heavily rely on sophisticated network structures (e.g., the meta path-specific attention network in HAN [6] and the hierarchical transformer in HGT [7]) in order to fully capture the contexts associated with different node/edge types. Therefore, practicality-wise, compared with the simplistic homogeneous GNNs, existing heterogeneous GNNs exhibit compromised training efficiency due to the complex architectures and an excessive amount of parameters to be learned.

In light of the outstanding problems from heterogeneity, inductiveness and efficiency perspectives in graph representation learning, we aim to thoroughly investigate all these three issues in a unified view. Specifically, we propose **wide and deep message passing network (WIDEN)** as an inductive and efficient solution to heterogeneous graph

representation learning. To ensure inductiveness, WIDEN essentially follows the message passing paradigm in GNNs, where information is propagated from sampled neighbor nodes to the target node to form a node’s embedding. To account for the heterogeneity of graphs, instead of using the inflexible meta paths or the cumbersome relation-specific aggregation schemes, we propose a novel heterogeneous message packaging paradigm that does not ask for separate modelling of different types of meta paths/relations. Specifically, in a sampled neighborhood, we generate heterogeneous message packs by posing interactions between a neighbor node and its associated edge. As message packs are concretized representations of heterogeneous information passed from neighbor nodes, the forward propagation in WIDEN is performed at the message pack level and can thoroughly capture the heterogeneity of the graph.

To overcome the aforementioned dilemma of utilizing only meta paths or local neighborhood for heterogeneous message passing, WIDEN specifies two types of neighborhoods, i.e., wide and deep neighborhoods. Essentially, the wide neighborhood and deep neighborhood can be respectively viewed as the width-first search node set and random walk node sequence. In a nutshell, when learning the embedding for a target node, the new message passing paradigm simultaneously makes full use of its local neighborhood information, and ensures sufficient outreach to further nodes. Besides, we introduce an innovative downsampling scheme in WIDEN to boost its training efficiency. Intuitively, for each target node, we prune its wide and deep neighbor nodes during training based on a learnable attentive weight indicating each neighbor’s relevance to the target node. Notably, we further propose a specialized pruning method for deep neighbor sequences to avoid losing important semantic information if nodes are directly removed from a complete path.

To conclude, we make the following contributions:

- We systematically investigate the heterogeneity, inductiveness, as well as efficiency when learning representations for graphs in a unified view. We propose WIDEN, which is an innovative graph neural network variant that supports inductive and efficient representation learning on heterogeneous graphs.
- The novel message packaging paradigm, coupled with our wide and deep message passing architecture, makes WIDEN free of manually crafted meta paths and complex relation-specific modelling, while ensuring a balanced blend of messages from local neighborhood and high-order connections. The proposed downsampling strategy for both wide and deep neighbor nodes further facilitates efficient training of WIDEN with minimal accuracy loss.
- We conduct extensive experiments on three real-world heterogeneous graphs, and the results show that WIDEN outperforms state-of-the-art baselines in both transductive and inductive representation learning tasks. Furthermore, WIDEN is highly efficient in terms of time consumption during training.

2 PRELIMINARIES

In this section, we present the key definitions that are frequently used in our paper.

Definition 1: Heterogeneous Graph. A heterogeneous graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ consists of a node set \mathcal{V} and an edge set \mathcal{E} . Mathematically, an edge between nodes $v_i, v_j \in \mathcal{V}$ is defined as $e_{ij} = (v_i, v_j)$ where both nodes and edges are heterogeneous, e.g., authorship relation between author and paper nodes, or citation relation between two paper nodes in an academic graph.

Definition 2: Wide Neighbor Node Set. For a given target node $v_t \in \mathcal{V}$, its wide neighbor node set is defined as $\mathcal{W}(v_t) = \{(n, i)\}_{n=1}^{N_w}$ that stores the **indexes** i of its uniformly sampled first-order neighbor nodes, and $|\mathcal{W}(v_t)| = N_w$ where N_w is the initial sample size. In each tuple (n, i) , $i \in [1, |\mathcal{V}|]$ is the **global index** that is unique for each node $v_i \in \mathcal{V}$, while $n \in [1, N_w]$ is the **local index** used only inside the set $\mathcal{W}(v_t)$.

Definition 3: Deep Neighbor Node Set. For target node $v_t \in \mathcal{V}$, its deep neighbor set $\mathcal{D}(v_t)$ is initialized by performing a deep random walk of length N_d that starts from v_t . Hence, $\mathcal{D}(v_t) = \{(s, i)\}_{s=1}^{N_d}$ where $s \leq N_d$ and $i \leq |\mathcal{V}| = N_d$ are respectively the local and global node indexes. Note that we do not include the target node itself in both neighbor sets. In $\mathcal{D}(v_t)$, s is exactly the position of nodes in the sequence generated by the random walk, e.g., $s = 1$ and $s = N_d$ correspond to the first node (i.e., one of v_t 's first-order neighbor) and last node in the sequence, respectively.

Notations. Throughout this paper, all vectors are **row vectors** unless specified, e.g., $\mathbf{v}_t \in \mathbb{R}^{1 \times d}$. To maintain simplicity, we use superscripts \circ and \triangleright to distinguish notations for wide and deep neighbors, respectively. To make the subscripts easy to follow across computations, we resort to **local indexes** when denoting the neighbor nodes of v_t , i.e., $v_n = v_i$ for $(n, i) \in \mathcal{W}(v_t)$ and $v_s = v_i$ for $(s, i) \in \mathcal{D}(v_t)$.

Embedding Initialization. Each node is associated with a d_0 -dimensional feature vector $\mathbf{x}_t \in \mathbb{R}^{1 \times d_0}$ that is used to initialize its embedding, i.e., $\mathbf{v}_t = \mathbf{x}_t \mathbf{G}^{node}$ with linear projection weight $\mathbf{G}^{node} \in \mathbb{R}^{d_0 \times d}$. Also, for edge e_{ij} between any two nodes, we define its type-specific embedding as $\mathbf{e}_{ij} = \text{onehot}(e_{ij}) \mathbf{G}^{edge}$, where $\text{onehot}(\cdot)$ returns the one-hot encoding of e_{ij} 's type, and $\mathbf{G}^{edge} \in \mathbb{R}^{d'_0 \times d}$ is the embedding matrix for a total of d'_0 edge types.

3 WIDE AND DEEP MESSAGE PASSING NETWORK

In a wide and deep message passing step, taking the current representation \mathbf{v}_t for the target node as the input, our inductive message passing scheme aims to generate an updated representation \mathbf{v}'_t by propagating information from both the wide and deep neighbor node sets of v_t . To thoroughly account for the heterogeneity of the graph, apart from the diversified node features, we take the specific type of the edges between nodes into account. Hence, for wide and deep neighbor sets, we first define two corresponding heterogeneous message packaging functions $\text{PACK}^\circ(\cdot)$ and $\text{PACK}^\triangleright(\cdot)$ to have all node and edge information packed up in matrices \mathbf{M}_t° and $\mathbf{M}_t^\triangleright$. The generated message packs diverges from meta path-based methods that do not specifically model edge information. Then, we propose two information passing functions $\text{PASS}^\circ(\cdot)$ and $\text{PASS}^\triangleright(\cdot)$ that

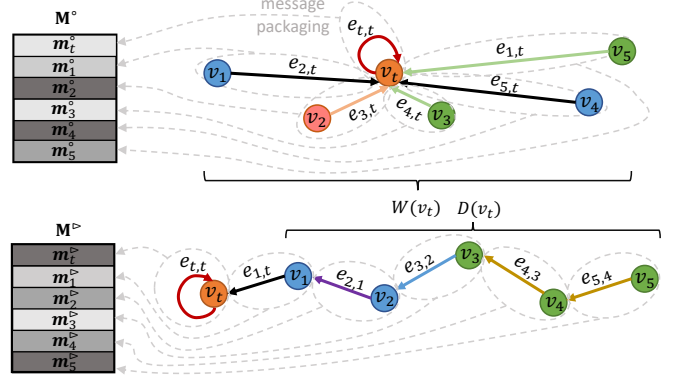


Fig. 1. A schematic view of message packaging in WIDEN. Top and bottom parts correspond to wide and deep neighbor sets, respectively. In this figure, we suppose $|\mathcal{W}(v_t)| = |\mathcal{D}(v_t)| = 5$ for illustration purposes.

respectively aggregate the information within the wide and deep message packs into two compact vectors \mathbf{h}_t° and $\mathbf{h}_t^\triangleright$. Afterwards, function $\text{FUSE}(\cdot)$ fuses both representations and generates \mathbf{v}'_t to replace the original node embedding.

Note that as both wide and deep neighbor sets are subject to downsampling in WIDEN, we use $|\mathcal{W}(v_t)|$ and $|\mathcal{D}(v_t)|$ to respectively denote the current number of neighbors in each set instead of using the fixed initial neighbor sizes N_w and N_d . The downsampling phase essentially reduces the amount of messages need to be passed into v_t , thus allowing for more efficient training. In what follows, we present the design of those key components in detail.

3.1 Heterogeneous Message Packaging

As defined in Section 2, the wide neighbor set $\mathcal{W}(t)$ contains $|\mathcal{W}(v_t)|$ sampled first-order neighbors for target node v_t . To fully preserve the heterogeneous information within $\mathcal{W}(t)$ during message passing, we define the following message packaging process:

$$\begin{aligned} \mathbf{M}^\circ &= \text{PACK}^\circ(\mathcal{W}(v_t)) \\ &= \begin{bmatrix} \mathbf{m}_t^\circ \\ \mathbf{m}_1^\circ \\ \vdots \\ \mathbf{m}_{|\mathcal{W}(v_t)|}^\circ \end{bmatrix} = \begin{bmatrix} \mathbf{v}_t \odot \mathbf{e}_{t,t} \\ \parallel_{n=1}^{|\mathcal{W}(v_t)|} \underbrace{(\mathbf{v}_n \odot \mathbf{e}_{n,t})}_{\text{a single message pack}} \end{bmatrix}, \quad (1) \end{aligned}$$

where \odot is the element-wise multiplication, and \parallel denotes the iterative operator for vertical concatenation of vectors. As illustrated in Figure 1, the resulted matrix $\mathbf{M}^\circ \in \mathbb{R}^{(|\mathcal{W}(v_t)|+1) \times d}$ carries the message packs of the target node itself \mathbf{m}_t° and all its wide neighbors $\mathbf{m}_1^\circ, \dots, \mathbf{m}_{|\mathcal{W}(v_t)|}^\circ$. Specifically, $\mathbf{e}_{n,t} \in \mathbb{R}^{1 \times d}$ denotes the trainable embedding vector of the relation (i.e., edge type) between the target node v_t and neighbor v_n . For example, when modelling academic graphs, the edge types may include scholar-paper authorship, paper-paper citation, paper-topic subordination, etc. After the element-wise interaction between the node representation and edge embedding, the resulted d -dimensional vector encodes the heterogeneous information of both parts, which we term a **message pack** \mathbf{m} . Notably, for \mathbf{m}_t° , i.e., the target node v_t , we also learn a self-loop edge embedding $\mathbf{e}_{t,t}$ between the same type of nodes (e.g., conference-conference self-loop edge).

Similarly, as Figure 1 depicts, we can construct a deep message matrix $\mathbf{M}^\triangleright \in \mathbb{R}^{(|\mathcal{D}(v_t)|+1) \times d}$ by stacking all heterogeneous message packs obtained from $\mathcal{D}(v_t)$:

$$\begin{aligned} \mathbf{M}^\triangleright &= \text{PACK}^\triangleright(\mathcal{D}(v_t)) \\ &= \begin{bmatrix} \mathbf{m}_t^\triangleright \\ \mathbf{m}_1^\triangleright \\ \vdots \\ \mathbf{m}_{|\mathcal{D}(v_t)|}^\triangleright \end{bmatrix} = \begin{bmatrix} \mathbf{v}_t \odot \mathbf{e}_{t,t} \\ \vdots \\ \mathbf{v}_s \odot \mathbf{e}_{s,s-1} \end{bmatrix}, \end{aligned} \quad (2)$$

where except for $\mathbf{m}_t^\triangleright$, each message pack is formed by fusing the s -th node representation \mathbf{v}_s with edge $\mathbf{e}_{s,s-1}$ linking its predecessor v_{s-1} in the random walk sequence. Note that $\mathbf{e}_{1,0} = \mathbf{e}_{1,t}$ connects the first node in $\mathcal{D}(v_t)$ with v_t .

3.2 Inductive Message Passing

We hereby introduce the message passing scheme in WIDEN that coordinates heterogeneity and inductivity for node representation learning.

Wide Attentive Message Passing. Intuitively, given $\mathcal{W}(v_t)$, we learn the representation of target node v_t by modelling the combinatorial message passed from its neighbors as well as itself. Among the literatures of homogeneous graphs, the most common message passing scheme is defined in [15], where an aggregator (e.g., average pooling) is applied for merging the information from v_t itself and its neighborhood. However, as pointed out in [4], instead of treating all neighbor nodes equally, different neighbor nodes should be weighted proportionally based on their importance to target node v_t . In the context of heterogeneous graphs, the capability of distinguishing the varied contributions from all heterogeneous message packs becomes especially crucial owing to the diversity of node and edge semantics. Hence, in WIDEN, we devise the following self-attention unit [19] to selectively transfer the information within \mathbf{M}° to v_t :

$$\begin{aligned} \mathbf{h}_t^\circ &= \text{PASS}^\circ(\mathbf{M}^\circ, \mathbf{m}_t^\circ) \\ &= \text{softmax}\left(\frac{\mathbf{m}_t^\circ \mathbf{W}_Q^\circ \cdot (\mathbf{M}^\circ \mathbf{W}_K^\circ)^\top}{\sqrt{d}}\right) \cdot \mathbf{M}^\circ \mathbf{W}_V^\circ, \end{aligned} \quad (3)$$

where $\mathbf{W}_Q^\circ, \mathbf{W}_K^\circ, \mathbf{W}_V^\circ \in \mathbb{R}^{d \times d}$ are respectively the query, key, and value weight matrices, while \sqrt{d} is the scaling factor to smooth the row-wise softmax output and avoid extremely large values of the inner product. Note that we only take the message pack of target node, i.e., \mathbf{m}_t° as the query of Eq.(3), hence the resulted $|\mathcal{W}(v_t)| + 1$ attentive weights $\{a_{t,t}^\circ, a_{t,1}^\circ, a_{t,2}^\circ, \dots, a_{t,|\mathcal{W}(v_t)|}^\circ\} = \text{softmax}\left(\frac{\mathbf{v}_t \mathbf{W}_Q^\circ \cdot (\mathbf{M}^\circ \mathbf{W}_K^\circ)^\top}{\sqrt{d}}\right)$ carries the probability distribution indicating the importance of each heterogeneous message pack in \mathbf{M}_t° to v_t . Then, through a weighted sum of all message packs, the self-attention unit summarized message for v_t , denoted by \mathbf{h}_t° . Unlike existing relation-specific or meta path-specific graph embedding paradigms [5]–[7], our self-attentive scheme avoids the need for repeating the computationally expensive aggregations for each pre-defined relations/meta paths and merging all information afterwards. As the message packs in \mathbf{M}° essentially encodes heterogeneous semantics from both nodes and edges, \mathbf{h}_t° is already an expressive representation of the heterogeneous connectivity around v_t .

Successive Self-Attention for Deep Message Passing. Learning node embeddings via sampled node sequences (i.e., paths) involves two widely used strategies. One originates from SkipGram [20], where the core is to maximize the log-likelihood of observing a sampled homogeneous node path [21], [22]. The other adopts recurrent neural networks (RNNs) which allows a node’s message to be propagated to distantly connected nodes, and is proven effective in capturing the information cascade within node sequences [23], [24]. Speaking of heterogeneity, the majority of state-of-the-art message passing networks are based on the notion of meta paths. However, meta path-based approaches are heavily constrained by their inflexibility and high dependent of domain expertise. To overcome this dilemma, we propose to subsume $\mathbf{M}^\triangleright$, which essentially a sequence of the deep message packs, under a sequential modelling paradigm with our proposed successive self-attentive message passing operation:

$$\mathbf{H}^\triangleright = \text{softmax}\left(\frac{\mathbf{M}^\triangleright \mathbf{W}_Q^\triangleright \cdot (\mathbf{M}^\triangleright \mathbf{W}_K^\triangleright)^\top}{\sqrt{d}} + \Theta\right) \cdot \mathbf{M}^\triangleright \mathbf{W}_V^\triangleright, \quad (4)$$

which is followed by:

$$\begin{aligned} \mathbf{h}_t^\triangleright &= \text{PASS}^\triangleright(\mathbf{H}^\triangleright, \mathbf{m}_t^\triangleright) \\ &= \text{softmax}\left(\frac{\mathbf{m}_t^\triangleright \mathbf{W}_Q^{\triangleright'} \cdot (\mathbf{H}_t^\triangleright \mathbf{W}_K^{\triangleright'})^\top}{\sqrt{d}}\right) \cdot \mathbf{M}^\triangleright \mathbf{W}_V^{\triangleright'}. \end{aligned} \quad (5)$$

Essentially, the first self-attention in Eq.(4) updates each deep message pack in \mathbf{M}° by capturing the sequential dependencies when information is propagated from the end of the sequence (i.e., $v_{|\mathcal{D}(v_t)|}$) to v_t . This provides the sequential characteristics of RNNs for node sequence modelling [23], [24], while offering a substantially lower computational complexity. Then, \mathbf{H}° can be viewed as a refined version of deep message packs with awareness of sequentiality. With the joint effect of $\text{PASS}^\triangleright(\cdot, \cdot)$ in Eq.(5), the resulted representation $\mathbf{h}_t^\triangleright \in \mathbb{R}^{1 \times d}$ will in return gather richer semantic information from the deep neighbors into the target node. However, conventional self-attention allows each element in the sequence to receive information from both preceding and succeeding ones, which is an inappropriate assumption for message passing in graph modelling. To ensure the information propagation is one-directional in a sequence, we incorporate an attention mask Θ into Eq.(4). Each element $\theta_{row,col} \in \Theta$ is defined as:

$$\theta_{row,col} = \begin{cases} 0, & \text{if } row \leq col \\ -\infty, & \text{otherwise} \end{cases}. \quad (6)$$

Rationale of The Attention Mask. We denote the matrix product of the query and key matrices in Eq.(4) as \mathbf{A} , i.e., $\mathbf{A} = \frac{\mathbf{M}^\triangleright \mathbf{W}_Q^\triangleright \cdot (\mathbf{M}^\triangleright \mathbf{W}_K^\triangleright)^\top}{\sqrt{d}} \in \mathbb{R}^{(|\mathcal{D}(v_t)|+1) \times (|\mathcal{D}(v_t)|+1)}$. By adding the attention mask $\mathbf{M}^\triangleright$, the interaction scores $a_{row,col} \in \mathbf{A}$ turn to $-\infty$ if $row > col$ and the softmax-normalized attentive weights consequently approximate 0. The interaction scores on other positions remain unchanged, ensuring that each message pack only yields influence to previous ones in the deep neighbor sequence.

Fusing Wide and Deep Representations. After obtaining \mathbf{h}_t° and $\mathbf{h}_t^\triangleright$, we combine the message passing results from both wide and deep neighbors to update node v_t ’s

embedding. It is worth noting that, to better compensate for the insufficiency of low-order neighbors, in practice we will sample $\Phi \geq 1$ deep neighbor sets for v_t . With the deep message passing scheme, one representation can be derived for each deep neighbor set, which is denoted by $\mathbf{h}_\phi^\triangleright$ with $1 \leq \phi \leq \Phi$. For all Φ learned message representations, we merge them into a unified representation with a simplistic average pooling operation that does not introduce excessive computational costs. Lastly, to generate the updated embedding for v_t , we send both wide and deep messages into a feed-forward network with a normalization operation:

$$\mathbf{h}_t = \text{ReLU}(\mathbf{W}[\mathbf{h}_t^\circ; \frac{1}{\Phi} \sum_{\phi=1}^{\Phi} \mathbf{h}_{t(\phi)}^\triangleright] + \mathbf{b}),$$

$$\mathbf{v}'_t = \frac{\mathbf{h}_t}{\|\mathbf{h}_t\|}, \quad (7)$$

where $[\cdot; \cdot]$ denotes horizontal concatenation of two vectors, $\mathbf{W} \in \mathbb{R}^{2d \times d}$ and $\mathbf{b} \in \mathbb{R}^{1 \times d}$ are respectively the weight and bias, and \mathbf{v}'_t is the output embedding of target node v_t in this deep and wide message passing step.

3.3 Efficient Training via Downsampling

To facilitate fast model learning, on top of the computationally lightweight self-attention network to facilitate efficient feed-forward, we propose a downsampling strategy to progressively shrink the number of message packs need to be propagated in every training iteration of WIDEN. Intuitively, by filtering out nodes (and connected edges) from $\mathcal{W}(v_t)$ and $\mathcal{D}(v_t)$ that make the least contribution to the learning of \mathbf{v}'_t . In short, by reducing the amount of message to be passed to node v_t , we can speed up the training process. Furthermore, with a selective downsampling process, the discarded message packs are of less importance in the current iteration. As such, though less neighborhood information is involved for learning \mathbf{v}'_t , downsampling high-quality node embeddings can be guaranteed in WIDEN.

Shrinking Wide Message Packs. In WIDEN, we advocate to fully utilize the attentive weights that are assigned to each neighbor node in $\mathcal{W}(v_t)$ for downsampling. Specifically, a single step for shrinking wide messages \mathbf{M}° is presented in Algorithm 1. Recall that $\{a_{t,t}^\circ, a_{t,1}^\circ, a_{t,2}^\circ, \dots, a_{t,|\mathcal{W}(v_t)|}^\circ\} = \text{softmax}(\frac{v_t \mathbf{W}_Q \cdot (\mathbf{M}^\circ \mathbf{W}_K)^\top}{\sqrt{d}})$ are the $|\mathcal{W}(v_t)| + 1$ attentive weights indicating the importance of message packs $\mathbf{m}_t^\circ, \mathbf{m}_1^\circ, \mathbf{m}_2^\circ, \dots, \mathbf{m}_{|\mathcal{W}(v_t)|}^\circ$ towards the

Algorithm 1 A Wide Message Shrinking Step

- 1: **Input:** Message matrix \mathbf{M}° for node v_t , attentive weights $\{a_{t,t}^\circ, a_{t,1}^\circ, a_{t,2}^\circ, \dots, a_{t,|\mathcal{W}(v_t)|}^\circ\}$ computed by Eq.(3);
 - 2: **Output:** Updated message matrix \mathbf{M}'° ;
 - 3: $\mathbf{a}^\circ \leftarrow \{a_{t,1}^\circ, a_{t,2}^\circ, \dots, a_{t,|\mathcal{W}(v_t)|}^\circ\}$; (exclude v_t itself)
 - 4: $n' \leftarrow \text{argmin}(\mathbf{a}^\circ)$;
 - 5: $\mathcal{W}'(v_t) = \emptyset$;
 - 6: **for each** $(n, i) \in \mathcal{W}(v_t)$ **do**
 - 7: $\mathcal{W}'(v_t) \cup \begin{cases} (n, i), & \text{if } n < n' \\ (n-1, i), & \text{if } n > n' \end{cases}$
 - 8: **end for** (update $\mathcal{W}(v_t)$ with new local indexes)
 - 9: $\mathbf{M}'_t \leftarrow \text{PACK}^\circ(\mathcal{W}'(v_t))$;
-




Fig. 2. A demonstration of why directly dropping message packs from deep neighbors causes semantic loss, and how our contextualized relay edges alleviate this issue when downsampling. The “Target Author” node is v_t in this figure.

learning of context \mathbf{h}_t° . Hence, after a message passing step, taking \mathbf{M}° and the attentive weights as the input, Algorithm 1 actively removes v_t 's one wide neighbor $v_n \in \mathcal{W}(v_t)$ that has the smallest attentive score $a_{t,n}^\circ$.

Pruning Deep Message Packs with Contextualized Relay Edges. For notation simplicity, we avoid individually referring to each ϕ -th deep neighbor set of v_t but keep using the single message matrix $\mathbf{M}^\triangleright$ in this section. The basic idea behind deep message pack downsampling is similar to the attentive weight-based method for wide message packs. However, an important difference of message passing with $\mathcal{D}(v_t)$ against $\mathcal{W}(v_t)$ is that, the node neighbors form a long-range sequence that carries a complete, high-order context to be transmitted into v_t . Hence, straightforwardly deleting node and its associated edge (i.e., a message pack) will lead to the connectivity on the node sequence ill-posed as the modified node sequence no longer reflects how the message is passed from the last node $v_{|\mathcal{D}(v_t)|}$ to v_t via heterogeneous relationships. As shown in Figure 2, when learning representation for the target author node, directly dropping the message pack with “GCN Paper” node and “cited” edge and rejoining the remaining content in the node sequence will create a path “T. Kipf authored ResNet Paper”, which carries wrong semantics and will impede the correctness of messages passed to v_t .

In WIDEN, we propose to generate contextualized relay edges in order to fully preserve the semantic information within $\mathcal{D}(v_t)$ during the downsampling process. To be specific, let $\mathbf{m}_{s'}^\triangleright$ be the message pack we will delete from $\mathbf{M}^\triangleright$ ($s' \in \{1, 2, \dots, |\mathcal{D}(v_t)|\}$). If $\mathbf{m}_{s'}^\triangleright$ is not the last element in the deep node sequence (i.e., $s' < |\mathcal{D}(v_t)|$), then we need to compute an additional relay edge to preserve the context of $\mathbf{m}_{s'}^\triangleright$ before deleting it. As depicted in Figure 2, to ensure that semantic information is retained for the entire path of $\mathcal{D}(v_t)$, we update its subsequent message pack $\mathbf{m}_{s'+1}$ via the following:

$$\mathbf{m}_{s'+1} \leftarrow \mathbf{v}_{s'+1} \odot \underbrace{\text{maxpool}(\mathbf{e}_{s'+1, s'}, \mathbf{m}_{s'}^\triangleright)}_{\text{relay edge}}, \quad (8)$$

where $\text{maxpool}(\cdot, \cdot)$ is the element-wise maxpooling operation between two vectors. With Eq.(8), we present Algorithm

Algorithm 2 A Deep Message Pruning Step

```

1: Input: Message matrix  $\mathbf{M}^\triangleright$  for node  $v_t$ , attentive weights
    $\{a_{t,t}^\triangleright, a_{t,1}^\triangleright, a_{t,2}^\triangleright, \dots, a_{t,|\mathcal{D}(v_t)|}^\triangleright\}$  computed by Eq.(5);
2: Output: Updated message matrix  $\mathbf{M}^{\triangleright'}$ ;
3:  $\mathbf{a}^\triangleright \leftarrow \{a_{t,1}^\triangleright, a_{t,2}^\triangleright, \dots, a_{t,|\mathcal{D}(v_t)|}^\triangleright\}$ ; (exclude  $v_t$  itself)
4:  $s' \leftarrow \text{argmin}(\mathbf{a}^\triangleright)$ ;
5: if  $s' < |\mathcal{D}(v_t)|$  do
6:    $\mathbf{m}_{s'+1}^\triangleright \leftarrow \text{Eq.}(8)$  where  $\mathbf{m}_{s'+1}^\triangleright \in \mathbf{M}^\triangleright$ ;
7: end if (renew the message pack with relay edge)
8:  $\mathcal{D}'(v_t) = \emptyset$ ;
9: for each  $(s, i) \in \mathcal{D}(v_t)$  do
10:   $\mathcal{D}'(v_t) \cup \begin{cases} (s, i), & \text{if } s < s' \\ (s-1, i), & \text{if } s > s' \end{cases}$ 
11: end for (update  $\mathcal{D}(v_t)$  with new local indexes)
12:  $\mathbf{M}_t^{\triangleright'} \leftarrow \text{PACK}^\triangleright(\mathcal{D}'(v_t))$ ;
13:  $\mathbf{m}_{s'}^{\triangleright'} \leftarrow \mathbf{m}_{s'+1}^\triangleright$  where  $\mathbf{m}_{s'}^{\triangleright'} \in \mathbf{M}^\triangleright$ ;
   (update  $s'$ -th message pack with relay edge information)

```

2 for pruning one deep message pack in a single iteration. In short, the relay edge binds the useful information from both the deprecated message pack \mathbf{m}_s and the edge embedding \mathbf{e}_{s+1} , thus generating a contextualized relay edge between node v_{s+1} and v_{s-1} prior to the deletion of v_s . Then, the new message pack \mathbf{m}_{s+1} can be computed by replacing the original \mathbf{e}_{s+1} with the relay edge. Intuitively, the relay edge greatly helps minimize the loss of semantics while still enabling deep message pack downsampling to facilitate efficient training of WIDEN.

3.4 Training WIDEN

We introduce the training process of WIDEN in this section.

Training Process. Algorithm 3 describes the efficient training procedure of our model. In Algorithm 3, lines 5-16 essentially iterate the mini-batch training by firstly engaging wide and deep message passing and then performing active downsampling for both neighbors. Noticeably, the downsampling process is subject to two additional constraints, i.e., the downsampling lower bounds k° and k^\triangleright and the downsampling thresholds r° and r^\triangleright . The lower bounds $k^\circ, k^\triangleright \geq 1$ controls the minimum neighbor set size that v_t preserves during the downsampling process. Once there are only k°/k^\triangleright remain in the wide/deep neighbor set, WIDEN stops the downsampling so that neighbor nodes will no longer be dropped in future iterations. Meanwhile, the downsampling thresholds r° and r^\triangleright determines circumstances where downsampling is triggered. We further elaborate on the downsampling triggering mechanism below.

Triggering Downsampling with KL Divergence. The downsampling process should be used with caution as suboptimal performance might be caused by the aggressively deleted data. This is because that if we do not allow sufficient training iterations to let the model “warm up” and grasp sufficient knowledge about a target node’s neighbor sets, directly reducing the amount of node neighbors will likely lead to high instability in the training phase. A straightforward solution is to put the downsampling process on hold until WIDEN’s training loss stabilizes, however this is a costly option due to the potentially large time

Algorithm 3 Training WIDEN

```

1: Input: Heterogeneous graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , raw features
   of all nodes  $\{\mathbf{x}_t\}_{t=1}^{|\mathcal{V}|}$ , neighborhood functions  $\mathcal{W}(\cdot)$  and
    $\mathcal{D}(\cdot)$ , initial neighbor sizes  $N_w$  and  $N_d$ , batch size  $B$ ,
   learning rate  $\tau$ , maximum training epoch  $Z$ , down-
   sampling thresholds  $r^\circ$  and  $r^\triangleright$ ; downsampling lower
   bounds  $k^\circ$  and  $k^\triangleright$ ;
2: Output: Vector representations  $\mathbf{v}_t$  for all  $v_t \in \mathcal{V}$ ;
3: Sample  $\mathcal{W}(v_t)$  and  $\mathcal{D}(v_t)$  for all  $v_t \in \mathcal{V}$ ;
4:  $\mathbf{M}^\circ \leftarrow \text{Eq.}(1)$ ,  $\mathbf{M}^\triangleright \leftarrow \text{Eq.}(2)$ ,  $z \leftarrow 0$ ; (initialization)
5: repeat
6:   Draw  $B$  target nodes from  $\mathcal{V}$ ;
7:   for each target node  $v_t$  in the training batch do
8:      $\mathbf{v}_t \leftarrow \text{Eq.}(3)$ -(7);
9:     if  $z > 1$  and  $KL_z^\circ < r^\circ$  and  $|\mathcal{W}(v_t)| > k^\circ$  do
10:       $\mathbf{M}^\circ \leftarrow \text{Algorithm 1}$ ;
11:    end if
12:    if  $z > 1$  and  $KL_z^\triangleright < r^\triangleright$  and  $|\mathcal{D}(v_t)| > k^\triangleright$  do
13:       $\mathbf{M}^\triangleright \leftarrow \text{Algorithm 2}$ ;
14:    end if (we assume  $\Phi = 1$  to be succinct)
15:     $L \leftarrow \text{Eq.}(10)$ ;
16:    Take a gradient step to optimize  $L$  with rate  $\tau$ ;
17:    if the end of  $\mathcal{V}$  is reached do
18:       $z \leftarrow z + 1$ ;
19:    end if
20: until  $L$  converges or  $z = Z$ ;

```

consumption before the loss converges. Hence, we propose to dynamically quantify WIDEN’s confidence in learning from v_t ’s neighbors in the current training epoch. Taking the wide neighbor set as an example, for the same target node v_t , if the same $\mathcal{W}(v_t)$ is used for two consecutive epochs $z-1$ and z , then we can compare WIDEN’s information gain in the current (i.e., z -th) over the previous (i.e., $(z-1)$ -th) epoch from neighbor set $\mathcal{W}(v_t)$. Intuitively, as each epoch learns a probability distribution $\{a_{t,t}^\circ, a_{t,1}^\circ, a_{t,2}^\circ, \dots, a_{t,|\mathcal{W}(v_t)|}^\circ\}$ over all $|\mathcal{W}(v_t)| + 1$ message packs, the information gain can be concretized by comparing the distributions learned in two epochs. As such, we formulate this process as computing the Kullback-Leibler (KL) divergence between two sets of attentive weights:

$$KL_z^\circ = \begin{cases} a_{t,t(z-1)}^\circ \ln \frac{a_{t,t(z)}^\circ}{a_{t,t(z-1)}^\circ} + \sum_{n=1}^{|\mathcal{W}(v_t)|} a_{t,n(z-1)}^\circ \ln \frac{a_{t,n(z)}^\circ}{a_{t,n(z-1)}^\circ}, & \text{if } \mathcal{W}_z(v_t) = \mathcal{W}_{z-1}(v_t) \\ +\infty, & \text{otherwise} \end{cases} \quad (9)$$

where we append subscript z or $z-1$ to differentiate notations for the current or previous training epoch. Notably, by feeding Eq.(9) with $\{a_{t,t}^\triangleright, a_{t,1}^\triangleright, a_{t,2}^\triangleright, \dots, a_{t,|\mathcal{D}(v_t)|}^\triangleright\}$, we can calculate the KL divergence regarding each deep neighbor set KL_z^\triangleright at epoch z . In general, a sufficiently small KL_z means that WIDEN acquires low information gain by performing attentive message passing on the same neighbor set again, and is hence highly confident in capturing information within the current neighbor set. Then, given $KL_z^\circ < r_z^\circ$ or $KL_z^\triangleright < r_z^\triangleright$, we can safely downsample the neighbors and let WIDEN re-capture the attentive weights with updated neighbors.

Loss Function. As a versatile and generic heterogeneous

graph embedding model, WIDEN can be optimized for different downstream tasks. In this paper, we focus on semi-supervised node classification, which is one of the most popular applications for heterogeneous graph embedding models [5]–[7]. Given the embedding \mathbf{v}_t^i of an arbitrary node (e.g., an author in a citation network), we aim to predict its categorical label (e.g., the research field of an author). To optimize WIDEN towards the task goal, we enforce the model to output a probability distribution over all classes, and adopt cross-entropy to quantify the training error for Algorithm 3:

$$L = - \sum_{t \in \mathcal{Y}_{train}} \mathbf{y}_t^\top \log(\text{softmax}(\mathbf{v}_t^i \mathbf{C})), \quad (10)$$

where $\mathbf{y}_t = \{0, 1\}^{1 \times c}$ is node v_t 's one-hot label among c different classes, \mathcal{Y}_{train} is the full training set, while we use projection weight $\mathbf{C} \in \mathbb{R}^{d \times c}$ followed by $\text{softmax}(\cdot)$ to estimate v_t 's probability distribution over all classes from \mathbf{v}_t^i . In semi-supervised circumstances where we only have labels for a fraction of nodes, WIDEN iterates over all nodes in \mathcal{V} while masking out unlabelled nodes in the graph.

4 EXPERIMENTS

In this section, we conduct experiments to evaluate WIDEN in terms of both effectiveness and efficiency. In particular, we aim to answer the following research questions (RQs) via experiments:

- RQ1:** How effectively can WIDEN perform heterogeneous graph representation learning?
- RQ2:** Can WIDEN still learn expressive node representations in an inductive manner?
- RQ3:** Is WIDEN efficient and scalable for training?
- RQ4:** How WIDEN benefits from each component of the proposed model structure?
- RQ5:** What is the impact of key hyperparameters to the performance of WIDEN?

4.1 Datasets

We perform node classification task with three heterogeneous graph datasets, namely DBLP, ACM, and Yelp. The key statistics of our experimental datasets are listed in Table 1. DBLP and ACM are both academic graphs, and we use the source data from [5], [6]. Yelp¹ is a social business review graph that allows users to comment and rate different businesses. We briefly describe their characteristics below:

- **DBLP:** In DBLP, the node types include *paper*, *author*, *conference*, and *terms*; and edge types include *paper-author*, *paper-conference*, and *paper-term*. Each author node in DBLP is labeled by one of four research areas (i.e., *database*, *data mining*, *machine learning*, and *information retrieval*). Following [5], [6], we use bag-of-words representations of each node's keywords as its raw features.
- **ACM:** ACM contains *paper*, *author*, and *subject* nodes, as well as *paper-author* and *paper-subject* edges. Labels are assigned to paper nodes based on the type of conferences (i.e., *database*, *wireless communication*, and

TABLE 1
Statistics of datasets in use.

Property	Dataset			
	ACM	DBLP	Yelp	
#Nodes	8,994	18,405	2,179,470	
#Node Types	3	4	4	
#Edges	25,922	67,946	37,776,380	
#Edge Types	2	3	4	
#Features	1,902	334	184	
#Class Labels	3	4	3	
Transductive Learning	#Training Nodes	600	800	60,000
	#Validation Nodes	300	400	30,000
	#Test Nodes	2,125	2,857	119,392
Inductive Learning	#Training Nodes	2,420	3,245	167,513
	#Validation Nodes	-	-	-
	#Test Nodes	605	812	41,879

data mining) they are accepted to. Similar to DBLP, node features are formulated as bag-of-words representations.

- **Yelp:** Yelp is a million-scale heterogeneous graph with *user*, *business*, *category* and *attribute* nodes, where the edge types include *user-business*, *user-user*, *business-category*, and *business-attribute*. In Yelp, we label each business node's service quality as *low*, *medium*, and *high*, which respectively corresponds to an overall rating below 3, between 3 and 4, and above 4. Such prediction is especially useful for evaluating new businesses where customer feedback is sparse. For user and business nodes, we extract Google's pre-trained word embeddings² of all words from their associated reviews, which are then averaged as the input features. Word embeddings of keywords are also used as the raw features of category and attribute nodes.

4.2 Baseline Methods

We compare WIDEN with the following state-of-the-art baseline methods:

- **Node2Vec:** This method [22] extends the Word2Vec to the graph embedding task, which maximizes the co-occurrence probability among nodes from sequences generated from random walks.
- **GCN:** The graph convolutional network [3] introduces spectral convolutions on graphs to support semi-supervised representation learning.
- **FastGCN:** As a parallelizable model, FastGCN [17] adopts a sampling strategy for mini-batch training while retaining similar performance as GCN.
- **GraphSAGE:** Using the sampling-and-aggregating scheme, GraphSAGE [15] is a widely used message passing network for graph representation learning.
- **GAT:** The graph attention network [4] employs attention mechanism to selectively aggregate information from neighbor nodes based on their importance.
- **GTN:** The graph transformer network [5] first constructs meta-paths by multiplying relation-specific adjacency matrices, then performs graph convolutions on those transformed adjacency matrices.

1. <https://www.yelp.com/dataset>

2. <https://code.google.com/archive/p/word2vec/>

- **HAN**: By attentively selecting relevant information from the sampled meta-paths, the heterogeneous attention network [6] is currently the state-of-the-art heterogeneous graph embedding approach.
- **HGT**: With a relation-aware transformer model and a subgraph sampling strategy, the heterogeneous graph transformer [7] is the latest graph embedding model for balancing performance and efficiency.

4.3 Evaluation Protocols

Transductive Graph Representation Learning. To avoid imbalance among node classes, we use **micro-averaged F1 score** to evaluate all methods’ classification performance. We further vary the amount of training labels (i.e., 25%, 50%, 75% and 100% of the full training set) to test the robustness of each method with different supervision strengths. The split of training, validation, and test sets are given in Table 1.

Inductive Graph Representation Learning. We further evaluate the performance of all models that support inductive learning, i.e., a part of the full graph (containing labeled test nodes) is completely unseen during training. This setting is highly relevant to high-throughput production systems [15] which constantly encounter unseen data and evolving graph structures. Hence, similar to [4], [15], we randomly pick 20% labeled nodes from DBLP, ACM and Yelp, and exclude them from the graph during training. All model hyperparameters are inherited from the optimal configurations tuned with transductive setting, hence we no longer sample validation nodes in this experiment.

4.4 Experimental Settings

To be consistent, we report the performance of WIDEN with a unified hyperparameter set $\{d = 128, N_w = 20, N_d = 20, N_t = 10\}$. We also set the learning rate $\tau = 0.0001$, downsampling thresholds $r^\circ = r^\triangleright = 0.001$ and lower bounds $k^\circ = k^\triangleright = 5$. For DBLP and ACM, we adopt $L2$ regularization with strength $\gamma = 0.01$ to prevent overfitting while we skip the regularization term on Yelp as the dataset is sufficiently large. The effect of different hyperparameter settings will be further discussed in Section 4.9. For all baselines, we optimize their hyperparameters via grid search with the validation set. Except for Node2Vec that is trained in a solely unsupervised manner, we adopt cross-entropy loss on the labeled training nodes to facilitate semi-supervised learning for all baselines. All effectiveness results are averaged over 5 executions. It is worth mentioning that some graph embedding baselines require the presence of the entire graph during training (i.e., Node2Vec, GCN, GAT, GTN and HAN), however the Yelp graph is too large to fit into the GPU memory. In this case, we adopt a well-established graph partitioning toolkit Metis [25] to divide the full graph into multiple subgraphs by minimizing edge cuts, such that those models can iterate over all subgraphs in the training phase.

4.5 Effectiveness on Transductive Heterogeneous Graph Embedding (RQ1)

Table 2 reports the transductive node classification results of all methods on three datasets. It is worth noting that, GTN

incurs huge time consumption when training on the largest dataset Yelp (one training epoch consumes more than 10 hours) due to its CPU-only implementation, hence its performance is not reported on Yelp. Based on the transductive learning effectiveness, we draw the following observations.

First, in almost all cases, WIDEN consistently and significantly outperforms all the baseline methods. On the relatively small graphs ACM and DBLP, when 100% of the labelled nodes in the training sets are used, WIDEN achieves a relative improvement of 1.98% and 20.12% over HGT, which is the latest heterogeneous GNN method. The margin becomes more apparent when tested on the large-scale graph Yelp with the full training set, where WIDEN advances the performance of the best homogeneous GNN (i.e., FastGCN) and heterogeneous GNN (i.e., HGT) by 8.15% and 20.86%, respectively. This verifies that, by incorporating heterogeneous message passing schemes with both wide and deep neighborhoods, WIDEN is able to learn more expressive node representations for downstream tasks.

Second, homogeneous graph embedding methods does not always underperform on heterogeneous. For example, GAT and FastGCN respectively yield comparative performance on ACM and Yelp. A highly possible reason is the varied information richness regarding nodes and edges in a heterogeneous graph. Commonly the nodes are associated with sufficient raw features (e.g., user demographics) for learning high-quality representations while the edges connecting them do not. Consequently, compared with homogeneous methods, heterogeneous GNNs that additionally consider node and edge types sometimes show only marginal performance improvement. In contrast, by constructing message packs with the node and edge representations, WIDEN is capable of fully utilizing the heterogeneity of different edges for message passing.

Third, as we reduce the amount of labelled nodes for training, WIDEN demonstrates the slightest performance drop among all the methods. This is a highly desirable property in applications based on semi-supervised learning, as real-world data usually faces difficulties in acquiring sufficient training labels.

4.6 Effectiveness on Inductive Heterogeneous Graph Embedding (RQ2)

We follow the evaluation protocols in Section 4.3 to test all methods’ performance on inductive learning. Note that Node2Vec is excluded from this test as its design requires all node IDs to be known beforehand. For inductive GNNs (i.e., GCN, FastGCN and GTN), we mask out the corresponding nodes in the input feature matrix and adjacency matrix during training, and recover them for evaluation to approximate the effect of inductive learning. As inductivity has been an increasingly important aspect for graph-based applications, we evaluate WIDEN’s performance in inductive settings both quantitatively and qualitatively.

Quantitative Results. We list the quantitative classification results on inductively learned node embeddings in Table 3. First of all, the inductive learning results place WIDEN in the leading position compared with all baselines. Hence, with the network parameters fully trained, WIDEN can effectively generalize to cold-start nodes that never

TABLE 2

Transductive node classification results. Different percentages indicate the proportion of labelled nodes used for training. Numbers in bold face are the best results in each column. Note that on Yelp, GTN incurs large training time cost (one epoch needs more than 10 hours) hence its performance is not reported. The results show that WIDEN significantly outperforms the best baseline in each column with p -value < 0.05 (underscored) or p -value < 0.01 (double-underscored) in paired t -test.

Method	ACM				DBLP				Yelp			
	25%	50%	75%	100%	25%	50%	75%	100%	25%	50%	75%	100%
Node2Vec [22]	0.7797	0.7665	0.7906	0.7910	0.9077	0.9020	0.9140	0.9181	0.4098	0.4074	0.4060	0.4069
GCN [3]	0.8058	0.8115	0.8133	0.8219	0.8381	0.8675	0.8683	0.8685	0.4523	0.4595	0.4765	0.4953
FastGCN [17]	0.7807	0.8875	0.9129	0.9188	0.7049	0.7871	0.8067	0.8039	0.5030	0.5265	0.6285	0.6638
GraphSAGE [15]	0.7567	0.7712	0.7816	0.8193	0.9062	0.9212	0.9205	0.9296	0.4629	0.5224	0.5475	0.5766
GAT [4]	0.8811	0.9076	0.9162	0.9128	0.8148	0.8443	0.8519	0.8667	0.4729	0.4748	0.5181	0.5208
GTN [5]	0.8844	0.8835	0.8912	0.9021	0.9032	0.9044	0.9195	0.9310	-	-	-	-
HAN [6]	0.8859	0.8950	0.9016	0.9052	0.9030	0.9219	0.9255	0.9247	0.4778	0.4910	0.5016	0.4832
HGT [7]	0.8757	0.8875	0.9054	0.9089	0.7315	0.7693	0.7767	0.7778	0.4629	0.5224	0.5742	0.5940
WIDEN	0.8870	0.9083	0.9196	0.9269	0.9076	0.9110	0.9293	0.9330	0.6892	0.7037	0.7125	0.7179

TABLE 3

Inductive node classification results. Numbers in bold face are the best results in each column. The results show that WIDEN significantly outperforms the best baseline on all datasets p -value < 0.01 (double-underscored) in paired t -test.

Method	ACM	DBLP	Yelp
GCN [3]	0.5735	0.4921	0.3523
FastGCN [17]	0.5826	0.5237	0.3616
GraphSAGE [15]	0.8016	0.9185	0.4214
GAT [4]	0.9044	0.8543	0.5829
GTN [5]	0.7829	0.8384	-
HAN [6]	0.9005	0.9210	0.5315
HGT [7]	0.9091	0.8264	0.6424
WIDEN	0.9175	0.9251	0.7613

appear in the training phase. Noticeably, as the inductive setting involves 80% labelled nodes for training, this leads to a substantial increase of the training data size on three graphs, considering the transductive setting only uses a small portion of the labelled nodes to train the models. As a side effect, on Yelp graph with a 100,000 bump in the size of training nodes, a slight performance gain is observed from several deep methods, e.g., GAT, HAN, HGT and WIDEN.

Qualitative Results. To qualitatively understand WIDEN’s capability of performing inductive node representation learning on heterogeneous graphs, we organize a case study by visualizing the inductively learned node embeddings on all three datasets. Specifically, we translate the learned embeddings of inductive nodes into 2-dimensional vectors with t-SNE [26], which are then plotted in Figure 3. We use different colours to distinguish nodes from different classes. As can be seen from Figure 3, nodes from the same class tend to form their own cluster, and there are clear boundaries between two difference classes. Thus, even if all these nodes are new to WIDEN, it is still able to generate high-quality node representations owing to its comprehensive message passing architecture with the important edge information, and both wide and deep neighbors.

4.7 Training Efficiency and Scalability (RQ3)

Unlike most existing heterogeneous GNNs [6], [7], one advantage of WIDEN is that it bypasses the sophisticated design for modelling relation-specific message passing processes. The message packaging scheme allows different semantics to be carried and passed into target nodes with a unified message passing architecture. In this section, we test the training efficiency and scalability of WIDEN given the




Fig. 3. t-SNE visualization of inductively learned node embeddings on ACM, DBLP, and Yelp. Each node class in the dataset corresponds to one colour in the plot. Note that for Yelp, we randomly sampled 1,000 inductive nodes for better clarity.

importance of heterogeneous GNN’s practicality in real-life applications. We discuss both aspects below.

Training Efficiency of WIDEN. To make the training efficiency measurable, we examine all methods’ training time consumption per epoch on ACM and DBLP³, and further evaluate their performance on (transductive) node classification after 10 training epochs. As such, if it takes less time for a model to be trained over one epoch and the model yields better performance with the same amount of training epochs, the model is regarded highly efficient to train. As discussed in Section 4.4, on Yelp graph, most baseline methods can only be trained on one sampled sub-graph at a time, hence this part of the test is only conducted on the small-scale ACM and DBLP graphs. We demonstrate the results on training efficiency of all models in Figure 4. On both ACM and DBLP, WIDEN consumes as few as 0.8964s and 0.9213s to be trained over one epoch, which is faster than the relatively lightweight GraphSAGE and FastGCN (both need more than 1s). The reasons behind are two-fold. First, WIDEN deploys the highly efficient self-attention modules to facilitate message passing between nodes. Second, the downsampling strategy directly reduces the amount of computations needed in each iteration, thus greatly cutting down the time consumption during training. Meanwhile, by outperforming all baselines in node classification on both graphs after only 10 training epochs, WIDEN has demonstrated highly competitive training efficiency.

Scalability of WIDEN. We test the training efficiency and scalability of WIDEN by varying the proportions of the nodes in the graph. This is achieved by randomly sampling nodes from the full graph with a ratio of $\{0.2, 0.4, 0.6, 0.8, 1.0\}$ and then report the corresponding

3. Hardware configuration: Intel Xeon Gold 6128 for CPU-enabled methods and Nvidia GeForce RTX 2080 Ti for GPU-enabled methods.

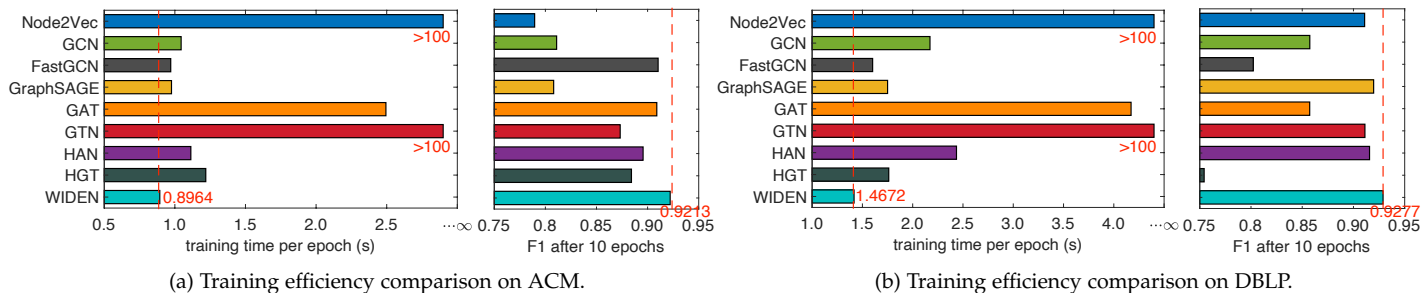


Fig. 4. Training efficiency results. F1 scores after 10 training epochs are reported for reference, and the training time per epoch is averaged over all the training epochs taken by each method.




Fig. 5. Training time of WIDEN on Yelp w.r.t varied data proportions.

time cost for model training. The Yelp graph is used for scalability test since it is the largest graph in our experiments. The growth of training time along with the data size is shown in Figure 5. When the ratio of training data gradually extends from 0.2 to 1.0, the training time for SeqFM increases from 0.61×10^3 seconds to 3.38×10^3 seconds. It shows that the dependency of training time on the data scale is approximately linear. Hence, we conclude that WIDEN is scalable to even larger heterogeneous graphs.

4.8 Importance of Key Components (RQ4)

To better understand the performance gain from the major components proposed in WIDEN, we conduct ablation tests on different degraded versions of WIDEN. Each variant removes one key component from the model, and the corresponding results on node classification are reported. Table 4 summarizes node classification outcomes with different model architectures. In what follows, we introduce the variants and analyze their effect respectively.

No Downsampling. The adaptive downsampling strategy aims to selectively locate important node/edge information and improve the training efficiency with optimal performance. Removing the downsampling component will make the wide and deep node neighbor sets remain unchanged during the entire training process. As we can infer from Table 4, the resulted model performance is very similar on DBLP and shows a slight increase on ACM and Yelp. This is within our expectation as retaining all neighbor nodes during training will offer more information for message passing. As the same time, it indicates that our proposed downsampling strategy is fully able to ensure high-quality node representation learning.

Removing Wide or Deep Neighbors. The second and third variant of WIDEN removes the message passing from wide and deep neighbors, respectively. Both variants suffer from inferior performance, especially when deep neighbors are removed from WIDEN since a severe (over 5%) performance drop is observed on both DBLP and Yelp graphs. Also, when wide neighbors are totally removed, WIDEN

TABLE 4
Ablation test with different model architectures, where “↓” marks a severe (over 5%) performance drop.

Architecture	ACM	DBLP	Yelp
Default	0.9269	0.9330	0.7179
No Downsampling	0.9352	0.9323	0.7334
Removing Wide Neighbors	0.9046	0.9023	0.7024
Removing Deep Neighbors	0.8976	0.8126↓	0.6720↓
Removing Successive Self-Attention	0.9035	0.8832↓	0.6913
Removing Relay Edges	0.8885	0.8915	0.6947
Random Downsampling for $\mathcal{W}(t)$	0.9192	0.9110	0.7111
Random Downsampling for $\mathcal{D}(t)$	0.8743↓	0.8537↓	0.6867

can still learn expressive node representations solely from its deep neighbor sets. This verifies our motivation of enriching the messages passed to the target node with sampled deep node sequences.

Removing Successive Self-Attention. After removing the successive self-attention in Eq.(4), the deep message passing part is simplified into an attentive aggregation of all deep neighbor nodes w.r.t. the target node. This leads to worse model performance on all three datasets as removing this components limits WIDEN’s ability to comprehensively capture the sequential dependencies among the message packs within the sampled deep walk sequences.

Removing Relay Edges. Without the relay edges, the deprecated message packs in the downsampling stage are directly discarded. As the relay edges help retain the semantics within the original deep walk sequences, deleting this component apparently makes the performance deteriorates.

Random Downsampling for $\mathcal{W}(t)$ or $\mathcal{D}(t)$. Instead of using the computed attention scores for identifying irrelevant nodes, we randomly drop one node from either $\mathcal{W}(t)$ or $\mathcal{D}(t)$ during the downsampling process. Note that the KL-divergence triggering mechanism is also removed in this variant. As can be seen from Table 4, randomly downsampling deep neighbors leads to more significant decrease on WIDEN’s performance compared with randomly dropping wide neighbors. One reason is that, as $\mathcal{D}(t)$ contains higher-order connectivity information regarding the target node, using random downsampling instead of the attentive strategy is likely to drop out informative deep neighbors at an early stage, causing a substantial loss of useful information.

4.9 Hyperparameter Sensitivity (RQ5)

We answer the forth research question by investigating the performance fluctuations of WIDEN with varied hyperparameters. Particularly, as mentioned in Section 4.4, we study our model’s sensitivity to the latent dimension d , the number of wide neighbors N_w , the number of deep neighbors N_d , as well as the number of deep walk sequences for each

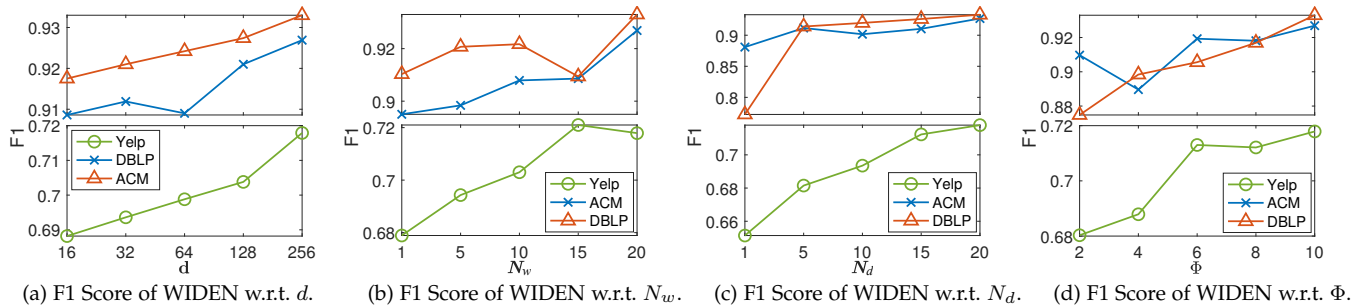


Fig. 6. Hyperparameter sensitivity analysis.

target node N_t . For each test, based on the standard setting $\{d = 256, N_w = 20, N_d = 20, N_t = 10\}$, we vary the value of one hyperparameter while keeping the others unchanged, and record the new F1 score achieved on transductive node classification. Figure 6 lays out the results with different parameter settings.

Impact of d . The value of the embedding dimension d is examined in $\{16, 32, 64, 128, 256\}$. As an important hyperparameter in deep neural networks, the latent dimension is apparently associated with the model’s expressiveness. In general, WIDEN benefits from a relatively larger d when learning node representations, hence setting $d = 256$ yields the highest F1 scores for node classification tasks. It is worth mentioning that on Yelp graph, even with $d = 16$, WIDEN still outperforms all the baselines, which further proves the effectiveness of our proposed model on large-scale heterogeneous graphs.

Impact of N_w . We study the impact of the number of initial neighbor nodes in each target node’s wide neighbor set with $N_w \in \{1, 5, 10, 15, 20\}$. For both ACM and DBLP graphs, a total of 20 initial neighbors in the wide neighbor set leads to the best performance, while inferior performance is observed when N_w is reduced. Interestingly, on Yelp graph, when $N_w = 15$, WIDEN achieves slightly better performance than having $N_w = 20$. This is possibly because setting N_w to 15 also helps remove neighbor nodes that might contain noisy information.

Impact of N_d . As can be concluded from Figure 6, WIDEN’s performance is positively associated with the number of deep neighbor nodes N_d , which are examined in $\{1, 5, 10, 15, 20\}$. Compared with smaller graphs ACM and DBLP, a larger N_d brings more significant performance increase on Yelp graph. This validates that on large and sparse heterogeneous graphs, passing information from remotely connected nodes is actually beneficial for node representation learning.

Impact of Φ . The impact of different numbers of deep walk sequences for a target node is investigated via $\Phi \in \{2, 4, 6, 8, 10\}$. Overall, the best prediction performance of SeqFM is reached when $\Phi = 10$ and more sampled deep walk sequences generally yield better results. We do notice that the performance gain starts to become less significant when Φ reaches a certain scale (6 in our case), indicating chances for a further boost in WIDEN’s efficiency with small trade-offs in performance.

5 RELATED WORK

In this section, we provide a summary of works related to our paper. We will start with classic (i.e., deep learning-free) graph representation learning methods, and then extend to state-of-the-art graph neural networks.

5.1 Deep Learning-free Graph Representation Learning Methods

Graph representation learning, i.e., graph embedding is the task of embedding a graph into a low-dimensional space while preserving the network structure and property to support downstream tasks [6]. Most existing graph embedding tasks focus on learning representations for nodes, where methods based on random walk [21], [22], [27] are initially proposed to learn node embeddings on homogeneous graphs. Originated from the skip-gram learning algorithm [20], random walk-based methods samples node sequences from a graph, then learns node embeddings by optimizing the co-occurrence probability of nodes within the same sequence. Meanwhile, LINE [28] introduces two hypothesis where two nodes are highly similar in the embedding space if they have direction connections (i.e., first-order proximity) and/or similar neighborhood compositions (i.e., second-order proximity), based on which two effective loss functions are designed for graph embedding.

To cope with the heterogeneity of nodes and edges, [29] proposes PTE model that extends the notion of second-order proximity [28] to heterogeneous graphs by separating a graph into multiple bipartite/homogeneous graphs. In a similar spirit, Chen et al. proposes to learn relation-specific projection matrices [30] in the divided bipartite/homogeneous graphs in order to learn more expressive node embeddings. Among heterogeneous graph embedding methods, a large body of recent works [31]–[34] make use of the heterogeneous contexts within meta paths [9] to facilitate node representation learning for downstream tasks. Essentially, a meta path is a sequence of heterogeneous nodes linked via various types of relations [35]. For example, [31] directly considers meta paths as contextual information when learning embeddings for nodes, while He et al. [33] expand the meta path-based random walk by proposing a heterogeneous personalized spacey random walk algorithm based on Markov chains.

Most deep learning-free graph embedding methods focus on improving the sampling strategies to acquire richer contextual information (e.g., generating high-quality random walk sequences [22], [33]) for learning node embeddings, or enhancing the optimization paradigm to maximize

the structural and property information preserved in the learned embedding (e.g., defining contrastive losses [28], [36]). Though proven effective, these methods are subject to limited expressiveness for representation learning. In what follows, we will discuss the deep graph neural networks which have been highly advantageous in a considerable amount of graph embedding tasks.

5.2 Graph Neural Networks

Recent years have witnessed the prominence of graph neural networks (GNNs) in modelling structural [37] and relational data [12]. Different from aforementioned shallow embedding methods, GNNs are often empowered by more complex encoders, usually a deep neural network, enabling the natural modeling of both structures and vertex attributes [17], [35], [38], [39]. In general, the inner mechanism of GNNs can be viewed as a message passing process, like propagating node or edge information via the graph structure in MPNN [40] or treating the neighbors of the target node as a receptive field for iterative feature aggregation in GCN [3], GraphSAGE [15] and GAT [4].

As an early variant of GNNs for modelling heterogeneous graphs, the relational GCN [12] learns a distinct linear projection weight for each edge type. Similarly, to deal with heterogeneous graph structures and node attributes, [10] uses type-specific RNNs to encode features for each type of neighbor vertices, followed by a subsequent RNN to aggregate different types of neighbor representations. More recently, the heterogeneous graph transformer (HGT) [7] uses each edge’s meta relation to parameterize the transformer network [19], so that both the common and specific patterns of different relationships are effectively captured.

Concurrently, an increasing amount of efforts have been devoted to enhancing heterogeneous GNNs with meta paths and attention mechanisms [5]–[7], [41]–[43]. For example, [42], [44] utilize the rich contexts within heterogeneous meta paths to enhance the representation learning capability of deep neural networks for recommendation tasks. Afterwards, the heterogeneous graph attention network (HAN) [5] defines different attentive aggregators and weight matrices for each meta path when embedding the target node. To bypass the inflexibility of manually defined meta paths, Yun et al. propose the graph transformer network (GTN) [5], which learns a soft selection of edge types and composite relations for generating useful multi-hop meta paths to augment the input graph. After that, plain GCNs are directly adopted in GTN for learning contextualized node embeddings.

5.3 Conclusion

This paper presents our effort in uniting heterogeneity, inductiveness, and efficiency for graph representation learning. To address those three practical problems, we propose WIDEN, which establishes a heterogeneous message packaging and passing paradigm for both wide and deep neighbor nodes that fully supports inductive representation learning. Furthermore, we design an active downsampling approach to facilitate efficient training while preserving the expressiveness of the learned node embeddings. Experimental results on real-world datasets have demonstrated

the efficacy of WIDEN in resolving the three important challenges when learning node representations on graphs.

REFERENCES

- [1] X. Wang, X. He, M. Wang, F. Feng, and T.-S. Chua, “Neural graph collaborative filtering,” in *SIGIR*, 2019.
- [2] K. T. Butler, D. W. Davies, H. Cartwright, O. Isayev, and A. Walsh, “Machine learning for molecular and materials science,” *Nature*, vol. 559, no. 7715, pp. 547–555, 2018.
- [3] T. N. Kipf and M. Welling, “Semi-supervised classification with graph convolutional networks,” *ICLR*, 2017.
- [4] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, “Graph attention networks,” in *ICLR*, 2018.
- [5] S. Yun, M. Jeong, R. Kim, J. Kang, and H. J. Kim, “Graph transformer networks,” in *NIPS*, 2019, pp. 11 983–11 993.
- [6] X. Wang, H. Ji, C. Shi, B. Wang, Y. Ye, P. Cui, and P. S. Yu, “Heterogeneous graph attention network,” in *WWW*, 2019, pp. 2022–2032.
- [7] Z. Hu, Y. Dong, K. Wang, and Y. Sun, “Heterogeneous graph transformer,” in *The Web Conference*, 2020, pp. 2704–2710.
- [8] X. Wang, X. He, Y. Cao, M. Liu, and T.-S. Chua, “Kgat: Knowledge graph attention network for recommendation,” in *SIGKDD*, 2019, pp. 950–958.
- [9] Y. Sun, J. Han, X. Yan, P. S. Yu, and T. Wu, “Pathsim: Meta path-based top-k similarity search in heterogeneous information networks,” *Vldb Endowment*, vol. 4, no. 11, pp. 992–1003, 2011.
- [10] C. Zhang, D. Song, C. Huang, A. Swami, and N. V. Chawla, “Heterogeneous graph neural network,” in *SIGKDD*, 2019, pp. 793–803.
- [11] S.-w. Park, B. J. Bae, J. Yeo, and S.-w. Hwang, “Meta-path free semi-supervised learning for heterogeneous networks,” *arXiv preprint arXiv:2010.08924*, 2020.
- [12] M. Schlichtkrull, T. N. Kipf, P. Bloem, R. Van Den Berg, I. Titov, and M. Welling, “Modeling relational data with graph convolutional networks,” in *ESWC*, 2018, pp. 593–607.
- [13] W.-C. Kang and J. McAuley, “Self-attentive sequential recommendation,” *ICDM*, pp. 197–206, 2018.
- [14] S. Nilizadeh, H. Aghakhani, E. Gustafson, C. Kruegel, and G. Vigna, “Think outside the dataset: Finding fraudulent reviews using cross-dataset analysis,” in *WWW*, 2019, pp. 3108–3115.
- [15] W. Hamilton, Z. Ying, and J. Leskovec, “Inductive representation learning on large graphs,” in *NIPS*, 2017, pp. 1024–1034.
- [16] R. Ying, R. He, K. Chen, P. Eksombatchai, W. L. Hamilton, and J. Leskovec, “Graph convolutional neural networks for web-scale recommender systems,” in *SIGKDD*, 2018.
- [17] J. Chen, T. Ma, and C. Xiao, “Fastgcn: fast learning with graph convolutional networks via importance sampling,” *ICLR*, 2018.
- [18] M. Gao, L. Chen, X. He, and A. Zhou, “Bine: Bipartite network embedding,” in *SIGIR*, 2018, pp. 715–724.
- [19] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” in *NIPS*, 2017.
- [20] T. Mikolov, K. Chen, G. Corrado, and J. Dean, “Efficient estimation of word representations in vector space,” *ICLR*, 2013.
- [21] B. Perozzi, R. Al-Rfou, and S. Skiena, “Deepwalk: Online learning of social representations,” in *SIGKDD*, 2014, pp. 701–710.
- [22] A. Grover and J. Leskovec, “node2vec: Scalable feature learning for networks,” in *SIGKDD*, 2016, pp. 855–864.
- [23] C. Li, J. Ma, X. Guo, and Q. Mei, “Deepcas: An end-to-end predictor of information cascades,” in *WWW*, 2017, pp. 577–586.
- [24] C. Yang, J. Tang, M. Sun, G. Cui, and Z. Liu, “Multi-scale information diffusion prediction with reinforced recurrent networks,” in *IJCAI*, 2019, pp. 4033–4039.
- [25] G. Karypis and V. Kumar, “A fast and high quality multilevel scheme for partitioning irregular graphs,” *SIAM Journal on Scientific Computing*, vol. 20, no. 1, pp. 359–392, 1998.
- [26] L. v. d. Maaten and G. Hinton, “Visualizing data using t-sne,” *JMLR*, vol. 9, no. Nov, pp. 2579–2605, 2008.
- [27] J. Qiu, Y. Dong, H. Ma, J. Li, K. Wang, and J. Tang, “Network embedding as matrix factorization: Unifying deepwalk, line, pte, and node2vec,” in *WSDM*, 2018, pp. 459–467.
- [28] J. Tang, M. Qu, M. Wang, M. Zhang, J. Yan, and Q. Mei, “Line: Large-scale information network embedding,” in *WWW*, 2015, pp. 1067–1077.

- [29] J. Tang, M. Qu, and Q. Mei, "Pte: Predictive text embedding through large-scale heterogeneous text networks," in *SIGKDD*, 2015, pp. 1165–1174.
- [30] H. Chen, H. Yin, W. Wang, H. Wang, Q. V. H. Nguyen, and X. Li, "Pme: projected metric embedding on heterogeneous networks for link prediction," in *SIGKDD*, 2018.
- [31] T.-y. Fu, W.-C. Lee, and Z. Lei, "Hin2vec: Explore meta-paths in heterogeneous information networks for representation learning," in *CIKM*, 2017, pp. 1797–1806.
- [32] Y. Cen, X. Zou, J. Zhang, H. Yang, J. Zhou, and J. Tang, "Representation learning for attributed multiplex heterogeneous network," in *SIGKDD*, 2019, pp. 1358–1368.
- [33] Y. He, Y. Song, J. Li, C. Ji, J. Peng, and H. Peng, "Hetspaceywalk: a heterogeneous spacey random walk for heterogeneous information network embedding," in *CIKM*, 2019, pp. 639–648.
- [34] C. Park, D. Kim, Q. Zhu, J. Han, and H. Yu, "Task-guided pair embedding in heterogeneous network," in *CIKM*, 2019, pp. 489–498.
- [35] Y. Dong, Z. Hu, K. Wang, Y. Sun, and J. Tang, "Heterogeneous network representation learning," in *IJCAI*, 2020, pp. 4861–4867.
- [36] A. Garcia Duran and M. Niepert, "Learning graph representations with embedding propagation," *NIPS*, pp. 5119–5130, 2017.
- [37] J. Bruna, W. Zaremba, A. Szlam, and Y. Lecun, "Spectral networks and locally connected networks on graphs," in *ICLR*, 2014.
- [38] P. Velickovic, W. Fedus, W. L. Hamilton, P. Liò, Y. Bengio, and R. D. Hjelm, "Deep graph infomax," in *ICLR*, 2019.
- [39] H. Chen, H. Yin, T. Chen, Q. V. H. Nguyen, W.-C. Peng, and X. Li, "Exploiting centrality information with graph convolutions for network representation learning," in *ICDE*, 2019, pp. 590–601.
- [40] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, "Neural message passing for quantum chemistry," *ICML*, 2017.
- [41] H. Chen, H. Yin, T. Chen, W. Wang, X. Li, and X. Hu, "Social boosted recommendation with folded bipartite network embedding," *TKDE*, 2020.
- [42] C. Shi, B. Hu, W. X. Zhao, and S. Y. Philip, "Heterogeneous information network embedding for recommendation," *TKDE*, vol. 31, no. 2, pp. 357–370, 2018.
- [43] S. Zhang, H. Yin, T. Chen, Q. V. N. Hung, Z. Huang, and L. Cui, "Gcn-based user representation learning for unifying robust recommendation and fraudster detection," in *SIGIR*, 2020, pp. 689–698.
- [44] B. Hu, C. Shi, W. X. Zhao, and P. S. Yu, "Leveraging meta-path based context for top-n recommendation with a neural co-attention model," in *SIGKDD*, 2018.