

Heterogeneous Network Representation Learning: A Unified Framework With Survey and Benchmark

Carl Yang¹, Yuxin Xiao², Yu Zhang³, Yizhou Sun⁴, *Member, IEEE*, and Jiawei Han, *Fellow, IEEE*

Abstract—Since real-world objects and their interactions are often multi-modal and multi-typed, heterogeneous networks have been widely used as a more powerful, realistic, and generic superclass of traditional homogeneous networks (graphs). Meanwhile, representation learning (*a.k.a.* embedding) has recently been intensively studied and shown effective for various network mining and analytical tasks. In this work, we aim to provide a unified framework to deeply summarize and evaluate existing research on heterogeneous network embedding (HNE), which includes but goes beyond a normal survey. Since there has already been a broad body of HNE algorithms, as the first contribution of this article, we provide a generic paradigm for the systematic categorization and analysis over the merits of various existing HNE algorithms. Moreover, existing HNE algorithms, though mostly claimed generic, are often evaluated on different datasets. Understandable due to the application favor of HNE, such indirect comparisons largely hinder the proper attribution of improved task performance towards effective data preprocessing and novel technical design, especially considering the various ways possible to construct a heterogeneous network from real-world application data. Therefore, as the second contribution, we create four benchmark datasets with various properties regarding scale, structure, attribute/label availability, and *etc.* from different sources, towards handy and fair evaluations of HNE algorithms. As the third contribution, we carefully refactor and amend the implementations and create friendly interfaces for 13 popular HNE algorithms, and provide all-around comparisons among them over multiple tasks and experimental settings. By putting all existing HNE algorithms under a unified framework, we aim to provide a universal reference and guideline for the understanding and development of HNE algorithms. Meanwhile, by open-sourcing all data and code, we envision to serve the community with an ready-to-use benchmark platform to test and compare the performance of existing and future HNE algorithms (<https://github.com/yangji9181/HNE>).

Index Terms—Heterogeneous network, representation learning, survey, benchmark

1 INTRODUCTION

NETWORKS and graphs constitute a canonical and ubiquitous paradigm for the modeling of interactive objects, which has drawn significant research attention from various scientific domains [1], [2], [3], [4], [5], [6]. However, real-world objects and interactions are often multi-modal and multi-typed (e.g., authors, papers, venues and terms in a publication network [7], [8]; users, places, categories and GPS-coordinates in a location-based social network [9], [10], [11]; and genes, proteins, diseases and species in a biomedical network [12], [13]). To capture and exploit such node and link heterogeneity, heterogeneous networks have been proposed and widely used in many real-world network mining scenarios, such as meta-path based similarity search [14], [15], [16], node classification

and clustering [17], [18], [19], knowledge base completion [20], [21], [22], and recommendations [23], [24], [25].

In the meantime, current research on graphs has largely focused on representation learning (embedding), especially following the pioneer of neural network based algorithms that demonstrate revealing empirical evidence towards unprecedentedly effective yet efficient graph mining [26], [27], [28]. They aim to convert graph data (e.g., nodes [29], [30], [31], [32], [33], [34], [35], [36], links [37], [38], [39], [40], and subgraphs [41], [42], [43], [44]) into low dimensional distributed vectors in the embedding space where the graph topological information (e.g., higher-order proximity [45], [46], [47], [48] and structure [49], [50], [51], [52]) is preserved. Such *embedding vectors* are then directly executable by various downstream machine learning algorithms [53], [54], [55].

Right on the intersection of heterogeneous networks and graph embedding, heterogeneous network embedding (HNE) recently has also received significant research attention [56], [57], [58], [59], [60], [61], [62], [63], [64], [65], [66], [67], [68], [69], [70], [71], [72], [73], [74], [75], [76]. Due to the application favor of HNE, many algorithms have been separately developed in different application domains such as search and recommendations [5], [23], [77], [78]. Moreover, as knowledge bases (KBs) also fall under the general umbrella of heterogeneous networks, many KB embedding algorithms can be compared with the HNE ones [4], [20], [21], [79], [80], [81], [82], [83], [84].

- Carl Yang is with the Emory University, Atlanta, GA 30322 USA. E-mail: jiyang3@illinois.edu.
- Yuxin Xiao is with the Carnegie Mellon University, Pittsburgh, PA 15213 USA. E-mail: yuxinx2@illinois.edu.
- Yu Zhang and Jiawei Han are with the University of Illinois, Urbana Champaign, IL 61820 USA. E-mail: {yuz9, hanj}@illinois.edu.
- Yizhou Sun is with the University of California, Los Angeles, Los Angeles, CA 90095 USA. E-mail: yzsun@cs.ucla.edu.

Manuscript received 29 June 2020; revised 30 Oct. 2020; accepted 16 Dec. 2020. Date of publication 21 Dec. 2020; date of current version 12 Sept. 2022.

(Corresponding author: Carl Yang.)

Recommended for acceptance by L. Chen.

Digital Object Identifier no. 10.1109/TKDE.2020.3045924

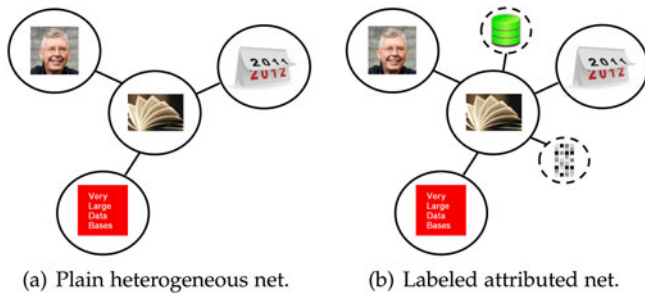


Fig. 1. Different heterogeneous networks constructed from the same real-world application data.

Unfortunately, various HNE algorithms are developed in quite disparate communities across academia and industry. They have never been systematically and comprehensively analyzed either in concepts or through experiments. In fact, due to the lack of benchmark platforms (with ready-to-use datasets and baselines), researchers often tend to construct their own datasets and re-implement a few most popular (sometimes outdated) algorithms for comparison, which renders fair performance evaluation and clear improvement attribution extremely hard, if not impossible.

Simply consider the toy examples of a publication dataset in Fig. 1.¹ Earlier HNE algorithms like metapath2vec [59] were developed on the heterogeneous network with node types of authors, papers and venues as in (a). However, one can enrich papers with a large number of terms and topics as additional nodes as in (b), which makes the random-walk based shallow embedding algorithms rather inefficient, but favors neighborhood aggregation based deep graph neural networks like R-GCN [67]. Moreover, one can further include node attributes like term embedding and labels like research fields and make them only available to the semi-supervised attributed embedding models, which may introduce even more bias [66], [75], [85], [86]. Eventually, it is often hard to clearly attribute performance gains between technical novelty and data tweaking.

In this work, we first formulate a unified yet flexible mathematical paradigm of HNE algorithms, easing the understanding of the critical merits of each model (Section 2). Particularly, based on a uniform taxonomy that clearly categorizes and summarizes the existing models (and likely future models), we propose a generic objective function of *network smoothness*, and reformulate all existing models into this uniform paradigm while highlighting their individual novel contributions (Section 3). We envision this paradigm to be helpful in guiding the development of future novel HNE algorithms, and in the meantime facilitate their conceptual contrast towards existing ones.

As the second contribution, we prepare four benchmark heterogeneous network datasets through exhaustive data collection, cleaning, analysis and curation (Section 4).² The datasets we come up with cover a wide spectrum of application domains (i.e., publication, recommendation, knowledge base, and biomedicine), which have various properties regarding scale, structure, attribute/label availability, etc. This diverse set of data, together with a series of standard network mining

tasks and evaluation metrics, constitute a handy and fair benchmark resource for future HNE algorithms.

As the third contribution, many existing HNE algorithms (including some very popular ones) either do not have a flexible implementation (e.g., hard-coded node and edge types, fixed set of meta-paths, etc.), or do not scale to larger networks (e.g., high memory requirement during training), which adds much burden to novel research (i.e., requiring much engineering effort in correct re-implementation). To this end, we focus on 13 popular HNE algorithms, where we carefully refactor and scale up the original implementations and apply additional interfaces for plug-and-run experiments on our prepared datasets (Section 5).² Based on these ready-to-use and efficient implementations, we then conduct all-around empirical evaluations of the algorithms, and report their benchmark performances. The empirical results, while providing much insight into the merits of different models that are consistent with the conceptual analysis in Section 3, also serve as the example utilization of our benchmark platform that can be followed by future studies on HNE.

Note that, although there have been several attempts to survey or benchmark heterogeneous network models [7], [8], [79], [87] and homogeneous graph embedding [26], [27], [28], [88], [89], [90], none of them has deeply looked into the intersection of the two. We advocate that our unified framework for the research and experiments on HNE is timely and necessary. First, as we will cover in this work, there has been a significant amount of research on the particular problem of HNE especially in the very recent several years, but most of them scatter across different domains, lacking proper connections and comparisons. Second, none of the existing surveys has proposed a generic mathematically complete paradigm for conceptual analysis of all HNE models. Third, existing surveys mostly do not provide systematic benchmark evaluation results, nor do they come with benchmark datasets and open-source baselines to facilitate future algorithm development.

The rest of this paper is organized as follows. Section 2 first introduces our proposed generic HNE paradigm. Subsequently, representative models in our survey are conceptually categorized and analyzed in Section 3. We then present in Section 4 our prepared benchmark datasets with detailed analysis. In Section 5, we provide a systematic empirical study over 13 popular HNE algorithms to benchmark the current state-of-the-art of HNE. Section 6 concludes the paper with visions towards future usage of our platform and research on HNE.

2 GENERIC PARADIGM

2.1 Problem Definitions

Definition 2.1. *Heterogeneous network.* A heterogeneous network $H = \{V, E, X, R, \phi, \psi\}$ is a network with multiple types of nodes and links. Particularly, within H , each node $v_i \in V$ is associated with a node type $\phi(v_i)$, and each link $e_{ij} \in E$ is associated with a link type $\psi(e_{ij})$. Each node v_i of type $o = \phi(v_i)$ is also potentially associated with attribute X_i^o , while each link e_{ij} of type $o = \psi(e_{ij})$ with attribute R_{ij}^o . It is worth noting that the type of a link e_{ij} automatically defines the types of nodes v_i and v_j on its two ends.

1. <https://dblp.uni-trier.de/>

2. <https://github.com/yangji9181/HNE>

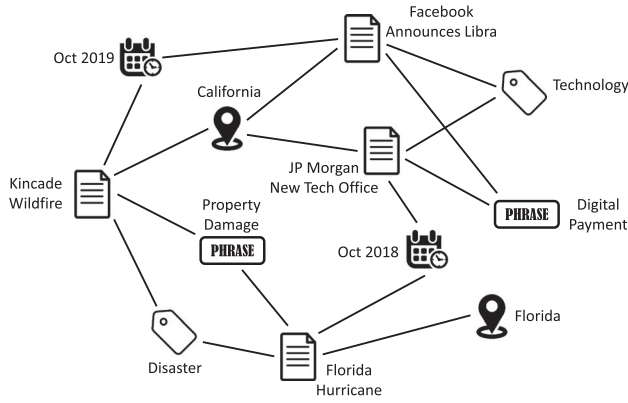


Fig. 2. Toy example of a heterogeneous network constructed from the news data.

Heterogeneous networks have been intensively studied due to its power of accommodating multi-modal multi-typed interconnected data. Besides the classic example of DBLP data used in most existing works as well as Fig. 1, consider a different yet illustrative example from NYTimes in Fig. 2.³ Nodes in this heterogeneous network include news articles, categories, phrases, locations, and datetimes. To illustrate the power of heterogeneous networks, we introduce the concept of meta-path, which has been leveraged by most existing works on heterogeneous network modeling [7], [8].

Definition 2.2. *Meta-path.* A meta-path is a path defined on the network schema denoted in the form of $o_1 \xrightarrow{l_1} o_2 \xrightarrow{l_2} \dots \xrightarrow{l_m} o_{m+1}$, where o and l are node types and link types, respectively.

Each meta-path captures the proximity among the nodes on its two ends from a particular semantic perspective. Continue with our example of heterogeneous network from news data in Fig. 2. The meta-path of article $\xrightarrow{\text{belongs to}}$ category $\xrightarrow{\text{includes}}$ article carries different semantics from article $\xrightarrow{\text{mentions}}$ location $\xrightarrow{\text{mentioned by}}$ article. Thus, the leverage of different meta-paths allows heterogeneous network models to compute the multi-modal multi-typed node proximity and relation, which has been shown beneficial to many real-world network mining applications [15], [25], [91].

Next, we introduce the problem of general network embedding (representation learning).

Definition 2.3. *Network embedding.* For a given network $G = \{V, E\}$, where V is the set of nodes (vertices) and E is the set of links (edges), a network embedding is a mapping function $\Phi : V \mapsto \mathbb{R}^{|V| \times d}$, where $d \ll |V|$. This mapping Φ defines the latent representation (a.k.a. embedding) of each node $v \in V$, which captures network topological information in E .

In most cases, network proximity is the major topological information to be captured. For example, DeepWalk [29] captures the random-walk based node proximity and illustrates the 2-dim node representations learned on the famous Zachary’s Karate network of small groups, where a clear correspondence between the node position in the input graph and learned embedding space can be observed. Various follow-up works have improved or extended DeepWalk, while a

complete coverage of them is beyond the scope of this work. In this work, we focus on the embedding of heterogeneous networks.

Now we define the main problem of focus in this work, heterogeneous network embedding, which lies in the intersection between Definitions 2.1 and 2.3.

Definition 2.4. *Heterogeneous network embedding.* For a given heterogeneous network H , a heterogeneous network embedding is a set of mapping functions $\{\Phi_k : V_k \mapsto \mathbb{R}^{|V_k| \times d}\}_{k=1}^K$, where K is the number of node types, $\forall v_i \in V_k, \phi(v_i) = k$, $d \ll |V|$. Each mapping Φ_k defines the latent representation (a.k.a. embedding) of all nodes of type k , which captures the network topological information regarding the heterogeneous links in E .

Compared with homogeneous networks, the definition of topological information in heterogeneous networks is even more diverse. As we will show in Section 3, the major distinctions among different HNE algorithms mostly lie in their different ways of capturing such topological information. Particularly, the leverage of meta-paths as in Definition 2.2 often plays an essential role, since many popular HNE algorithms exactly aim to model the different proximity indicated by meta-paths [59], [63], [66], [69], [70], [73], [75], [77].

2.2 Proposed Paradigm

In this work, we stress that one of the most important principles underlying HNE (as well as most other scenarios of network modeling and mining) is homophily [92]. Particularly, in the network embedding setting, homophily can be translated as ‘nodes close on a network should have similar embeddings’, which matches the requirement of Definition 2.3. In fact, we further find intrinsic connections between the well-perceived homophily principle and widely-used smoothness enforcement technique on networks [93], [94], [95], which leads to a generic mathematical paradigm covering most existing and likely many future HNE algorithms.

Based on earlier well-established concepts underlying network modeling and embedding learning [93], [94], [95], [96], [97], we introduce the following key objective function of network smoothness enforcement as follows

$$\mathcal{J} = \sum_{u,v \in V} w_{uv} d(e_u, e_v) + \mathcal{J}_R, \quad (1)$$

where $e_u = \Phi(u)$ and $e_v = \Phi(v)$ are the node embedding vectors to be learned. w_{uv} is the proximity weight, $d(\cdot, \cdot)$ is the embedding distance function, and \mathcal{J}_R denotes possible additional objectives such as regularizers, all three of which can be defined and implemented differently by the particular HNE algorithms.

3 ALGORITHM TAXONOMY

In this section, we find a universal taxonomy for existing HNE algorithms with three categories based on their common objectives, and elaborate in detail how they fit into our paradigm of Eq. (1). The main challenge of instantiating Eq. (1) on heterogeneous networks is the consideration of complex interactions regarding multi-typed links and higher-order meta-paths. In fact, our Eq. (1) also readily generalizes to homogeneous networks, though that is beyond the scope of this work.

3. <https://www.nytimes.com/>

3.1 Proximity-Preserving Methods

As mentioned above, one basic goal of network embedding is to capture network topological information. This can be achieved by preserving different types of proximity among nodes. There are two major categories of proximity-preserving methods in HNE: random walk approaches (inspired by DeepWalk [29]) and first/second-order proximity based ones (inspired by LINE [30]). Both types of proximity-preserving methods are considered as shallow network embedding, due to their essential single-layer decomposition of certain affinity matrices [98].

3.1.1 Random Walk Approaches

metapath2vec [59]. Following homogeneous network embedding [29], [31], *metapath2vec* utilizes the node paths traversed by meta-path guided random walks to model the context of a node regarding heterogeneous semantics. Formally, given a meta-path $\mathcal{M} = o_1 \xrightarrow{l_1} o_2 \xrightarrow{l_2} \dots \xrightarrow{l_{m-1}} o_m$, the transition probability at step i is defined as

$$p(v_{i+1}|v_i, \mathcal{M}) = \begin{cases} \frac{1}{|\mathcal{N}_i(v_i)|} & \phi(v_{i+1}) = o_{i+1}, \psi(v_i, v_{i+1}) = l_i \\ 0 & \text{otherwise} \end{cases}, \quad (2)$$

where $\mathcal{N}_i(v) = \{u | \psi(u, v) = l\}$ denotes the neighbors of v associated with edge type l . Assume $\mathcal{P} = \{\mathcal{P}_1, \dots, \mathcal{P}_M\}$ is the set of generated random walk sequences. The objective of *metapath2vec* is

$$\mathcal{J} = \sum_{v \in V} \sum_{u \in \mathcal{C}(v)} \log \frac{\exp(e_u^T e_v)}{\sum_{u' \in V} \exp(e_{u'}^T e_v)}, \quad (3)$$

where $\mathcal{C}(v)$ is the context (i.e., skip-grams) of v in \mathcal{P} . For example, if $\mathcal{P}_1 = v_1 v_2 v_3 v_4 v_5 \dots$ and the context window size is 2, then $\{v_1, v_2, v_4, v_5\} \subseteq \mathcal{C}(v_3)$. Let w_{uv} be the number of times that $u \in \mathcal{C}(v)$, and we can rewrite Eq. (3) as

$$\mathcal{J} = \sum_{u, v \in V} w_{uv} \log \frac{\exp(e_u^T e_v)}{\sum_{u' \in V} \exp(e_{u'}^T e_v)}.$$

Calculating the denominator in this objective requires summing over all nodes, which is computationally expensive. In actual computation, it is approximated using negative sampling [30], [99].

HIN2Vec [63]. *HIN2Vec* considers the probability that there is a meta-path \mathcal{M} between nodes u and v . Specifically,

$$p(\mathcal{M}|u, v) = \sigma\left(\mathbf{1}^T \left(W_X^T u \odot W_Y^T v \odot f_{01}(W_R^T m) \right)\right),$$

where $\mathbf{1}$ is an all-ones vector; \odot is the Hadamard product; f_{01} is a normalization function. Here $e_u = W_X^T u$, $e_v = W_Y^T v$ and $e_{\mathcal{M}} = f_{01}(W_R^T m)$ can be viewed as the embeddings of u , v and \mathcal{M} , respectively. Let $A_{\mathcal{M}} = \text{diag}(e_{\mathcal{M}1}, \dots, e_{\mathcal{M}d})$. We have

$$p(\mathcal{M}|u, v) = \sigma\left(\mathbf{1}^T (e_u \odot e_v \odot e_{\mathcal{M}})\right) = \sigma(e_u^T A_{\mathcal{M}} e_v).$$

σ is the sigmoid function, so we have

$$1 - p(\mathcal{M}|u, v) = 1 - \sigma(e_u^T A_{\mathcal{M}} e_v) = \sigma(-e_u^T A_{\mathcal{M}} e_v).$$

HIN2Vec generates positive tuples (u, v, \mathcal{M}) (i.e., u connects with v via the meta-path \mathcal{M}) using homogeneous random walks [29] regardless of node/link types. For each positive tuple (u, v, \mathcal{M}) , it generates several negative tuples by replacing u with a random node u' . Its objective is

$$\begin{aligned} \mathcal{J}_0 &= \sum_{(u, v, \mathcal{M})} \log p(\mathcal{M}|u, v) + \sum_{(u', v, \mathcal{M})} \log (1 - p(\mathcal{M}|u', v)) \\ &= \sum_{(u, v, \mathcal{M})} \left(\log \sigma(e_u^T A_{\mathcal{M}} e_v) + \sum_{u'} \log \sigma(-e_{u'}^T A_{\mathcal{M}} e_v) \right). \end{aligned}$$

This is essentially the negative sampling approximation of the following objective

$$\mathcal{J} = \sum_{\mathcal{M}} \sum_{u, v \in V} w_{uv}^{(\mathcal{M})} \log \frac{\exp(e_u^T A_{\mathcal{M}} e_v)}{\sum_{u' \in V} \exp(e_{u'}^T A_{\mathcal{M}} e_v)},$$

where $w_{uv}^{(\mathcal{M})}$ is the number of path instances between u and v following the meta-path \mathcal{M} .

Other random walk approaches are summarized in Table 1. To be specific, MRWNN [57] incorporates content priors into DeepWalk for image retrieval; SHNE [69] incorporates additional node information like categorical attributes, images, *etc.* by leveraging domain-specific deep encoders; HHNE [73] extends *metapath2vec* to the hyperbolic space; GHE [19] proposes a semi-supervised meta-path weighting technique; MNE [100] conducts random walks separately for each view in a multi-view network; JUST [65] proposes random walks with Jump and Stay strategies that do not rely on pre-defined meta-paths; HeteSpaceyWalk [101] introduces a scalable embedding framework based on heterogeneous personalized spacey random walks; TapEm [102] proposes a task-guided node pair embedding approach for author identification.

3.1.2 First/Second-Order Proximity Based Approaches

PTE [103]. *PTE* proposes to decompose a heterogeneous network into multiple bipartite networks, each of which describes one edge type. Its objective is the sum of log-likelihoods over all bipartite networks:

$$\begin{aligned} \mathcal{J} &= \sum_{l \in \mathcal{T}_E} \sum_{u, v \in V} w_{uv}^{(l)} \log \frac{\exp(e_u^T e_v)}{\sum_{u' \in V_{\phi(u)}} \exp(e_{u'}^T e_v)} \\ &= \sum_{u, v \in V} w_{uv} \log \frac{\exp(e_u^T e_v)}{\sum_{u' \in V_{\phi(u)}} \exp(e_{u'}^T e_v)}. \end{aligned}$$

Here \mathcal{T}_E is the set of edge types; $w_{uv}^{(l)}$ is the type- l edge weight of (u, v) (if there is no edge between u and v with type l , then $w_{uv}^{(l)} = 0$); $w_{uv} = \sum_l w_{uv}^{(l)}$ is the total edge weight between u and v .

AspEm [60]. *AspEm* assumes that each heterogeneous network has multiple aspects, and each aspect is defined as a subgraph of the network schema [14]. An incompatibility measure is proposed to select appropriate aspects for embedding learning. Given an aspect a , its objective is

$$\mathcal{J} = \sum_a \frac{1}{Z_l} \sum_{u, v \in V} w_{uv}^{(l)} \log \frac{\exp(e_{u,a}^T e_{v,a})}{\sum_{u' \in V_{\phi(u)}} \exp(e_{u',a}^T e_{v,a})},$$

TABLE 1
A Summary of Proximity-Preserving Based HNE Algorithms

Algorithm	$w_{uv} / w_{uv}^{(l)} / w_{uv}^{(M)}$	$d(e_u, e_v)$	\mathcal{J}_{R0}	Applications
MRWNN [57]	number of times that $u \in \mathcal{C}(v)$ in homogeneous random walks	$\ e_u - e_v\ ^2$	$-\sum_v \ e_v - f_{\text{ENC}}(X_v)\ ^2$	image retrieval
metapath2vec [59]	number of times that $u \in \mathcal{C}(v)$ in heterogeneous random walks following \mathcal{M}	$\ e_u - e_v\ ^2$	N/A	node classification, node clustering, link prediction, recommendation
SHNE [69]		$e_u = f_{\text{ENC}}(\mathbf{x}_u)$		
HHNE [73]		$d_{\mathbb{D}}(e_u, e_v)$		
GHE [19]	number of meta-path instances between u and v following \mathcal{M}	$\ e_u - e_v\ ^2$	N/A	author identification
HIN2Vec [63]	$\ \sqrt{\mathbf{A}_{\mathcal{M}}}(e_u - e_v)\ ^2$			
MNE [100]	number of times that $u \in \mathcal{C}(v)$ in homogeneous random walks in (V, E_l)	$\ e_u - e_{v,l}\ ^2$	N/A	node classification, node clustering, link prediction, recommendation
JUST [65]	number of times that $u \in \mathcal{C}(v)$ in Jump/Stay random walks [65]	$\ e_u - e_v\ ^2$		
HeteSpaceyWalk [101]	number of times that $u \in \mathcal{C}(v)$ in heterogeneous spacey random walks [101]			
TapEm [102]	number of times that $u \in \mathcal{C}(v)$ in heterogeneous random walks following \mathcal{M}	$\ f_{\text{MAP}}(e_u, e_v) - h_{uv}^{\mathcal{M}}\ ^2$ $h_{uv}^{\mathcal{M}}$: representation of the node sequence from u to v following \mathcal{M}	supervised loss $\sum y \log \hat{y} + (1 - y) \log(1 - \hat{y})$	author identification
HNE [56]	$\mathbf{1}_{(u,v) \in E}$	$\ e_u - e_v\ ^2$, $e_u = \mathbf{A}_{\phi(u)} \mathbf{x}_u$	$-\sum_{o \in \mathcal{T}_V} \ \mathbf{A}_o\ _F^2$	text classification, image retrieval
PTE [103]	edge weight of (u, v)	$\ e_u - e_v\ ^2$	N/A	text classification
DRL [68]			$-\sum_v \ e_v\ ^2$	node classification
CMF [58]				relatedness measurement of Wikipedia entities
HEBE [62]	edge weight of hyperedge (u, C) [62]	$\ e_u - e_C\ ^2$, $e_C = \sum_{v \in C} e_v / C $	N/A	event embedding
Phine [11]	number of times that u and v co-occur in a meta-graph instance	$\ e_u - e_v\ ^2$	supervised loss $-\sum \hat{y} - y ^2$	user rating prediction
MVE [104]	edge weight of (u, v) with type l	$\ e_{u,l} - e_v\ ^2$	$-\sum_v \sum_l \lambda_{v,l} \ e_{v,l} - e_v\ ^2$	node classification, link prediction
AspEm [60]		$\ e_{u,a} - e_{v,a}\ ^2$, a : aspect		aspect mining
HEER [61]		$\ \sqrt{\mathbf{A}_l}(e_u - e_v)\ ^2$		relation prediction in knowledge graphs
GERM [105]	edge weight of (u, v) with type l (if l is activated by the genetic algorithm [105])	$\ e_u - e_v\ ^2$	N/A	citation recommendation coauthor recommendation feeds recommendation
mg2vec [106]	1st: number of meta-graph instances containing node u 2nd: number of meta-graph instances containing both nodes u and v	1st: $\ e_u - e_m\ ^2$ 2nd: $\ f_{\text{MAP}}(e_u, e_v) - e_m\ ^2$ m : meta-graph		relationship prediction, relationship search

(Additional notations: f_{ENC} : a function to encode text/image/attribute information; $d_{\mathbb{D}}$: distance between two points in the Poincaré ball; f_{MAP} : a function to map two d -dimensional embeddings to one d -dimensional vector.)

where \mathcal{T}_E^a is the set of edge types in aspect a ; $Z_l = \sum_{u,v} w_{uv}^{(l)}$ is a normalization factor; $e_{u,a}$ is the aspect-specific embedding of u .

HEER [61]. HEER extends PTE by considering *typed closeness*. Specifically, each edge type l has an embedding μ_l , and its objective is

$$\mathcal{J} = \sum_{l \in \mathcal{T}_E} \sum_{u,v \in V} w_{uv}^{(l)} \log \frac{\exp(\mu_l^T g_{uv})}{\sum_{(u',v') \in P_l(u,v)} \exp(\mu_l^T g_{u'v'})},$$

where g_{uv} is the edge embedding of (u, v) ; $P_l(u, v) = \{(u', v) | \psi(u', v) = l\} \cup \{(u, v') | \psi(u, v') = l\}$. In [61], g_{uv} has different definitions for directed and undirected edges based on the Hadamard product. To simplify our discussion, we assume $g_{uv} = e_u \odot e_v$. Let $\mathbf{A}_l = \text{diag}(\mu_{l1}, \dots, \mu_{ld})$. Then we have $\mu_l^T g_{uv} = e_u^T \mathbf{A}_l e_v$ and

$$\mathcal{J} = \sum_{l \in \mathcal{T}_E} \sum_{u,v \in V} w_{uv}^{(l)} \log \frac{\exp(e_u^T \mathbf{A}_l e_v)}{\sum_{(u',v') \in P_l(u,v)} \exp(e_{u'}^T \mathbf{A}_l e_{v'})}.$$

Other first/second-order proximity based approaches are summarized in Table 1. To be specific, Chang *et al.* [56] introduce a node type-aware content encoder; CMF [58] performs joint matrix factorization over the decomposed bipartite networks; HEBE [62] preserves proximity regarding each meta-graph; Phine [11] combines additional regularization towards semi-supervised training; MVE [104] proposes an attention based framework to consider multiple views of the same nodes; DRL [68] proposes to learn one type of edges at each step and uses deep reinforcement learning approaches to select the edge type for the next step; GERM [105] adopts a genetic evolutionary approach to preserve the critical edge-types while removing the noisy or incompatible ones given specific tasks; mg2vec [106] jointly

embeds nodes and meta-graphs into the same space by exploiting both first-order and second-order proximity.

3.1.3 Unified Objectives

Based on the discussions above, the objective of most proximity-preserving methods can be unified as

$$\max \mathcal{J} = \sum_{u,v} w_{uv} \log \frac{\exp(s(u,v))}{\sum_{u'} \exp(s(u',v))} + \mathcal{J}_{R0}. \quad (4)$$

Here, w_{uv} is the weight of node pair (u, v) in the objective;⁴ \mathcal{J}_{R0} is an algorithm-specific regularization term (summarized in Table 1); $s(u, v)$ is a proximity function between u and v .

Negative Sampling. Directly optimizing the objective in Eq. (4) is computationally expensive because it needs to traverse all nodes when computing the softmax function. Therefore, in actual computation, most studies adopt the negative sampling strategy [30], [99], which modifies the objective as follows:

$$\sum_{u,v} w_{uv} \left(\log \sigma(s(u,v)) + b \mathbb{E}_{u' \sim P_N} [\log \sigma(s(u',v))] \right) + \mathcal{J}_{R0}.$$

Here, b is the number of negative samples; P_N is known as the noise distribution that generates negative samples.

Negative sampling serves as a generic paradigm to unify network embedding approaches. For example, starting from the negative sampling objective, Qiu *et al.* [98] unify DeepWalk [29], LINE [30], PTE [103] and node2vec [31] into a matrix factorization framework. In this paper, as mentioned in Section 2.2, we introduce another objective (i.e., Eq. (1)) that is equivalent to Eq. (4) and consider network embedding from the perspective of network smoothness enforcement.

Network Smoothness Enforcement. Note that in most cases, we can write $s(u, v)$ in Eq. (4) as $f(e_u)^T f(e_v)$. For example, $f(e_u) = e_u$ in metapath2vec, PTE, *etc.*; $f(e_u) = \sqrt{A_{\mathcal{M}}} e_u$ in HIN2Vec; $f(e_u) = \sqrt{A_l} e_u$ in HEER. In these cases,

$$\begin{aligned} \mathcal{J} &= \sum_{u,v} w_{uv} s(u,v) - \sum_{u,v} w_{uv} \log \sum_{u'} \exp(s(u',v)) + \mathcal{J}_{R0} \\ &= \sum_{u,v} w_{uv} f(e_u)^T f(e_v) \\ &\quad - \sum_{u,v} w_{uv} \log \sum_{u'} \exp(f(e_{u'})^T f(e_v)) + \mathcal{J}_{R0} \\ &= \sum_{u,v} \frac{w_{uv}}{2} \left(\|f(e_u)\|^2 + \|f(e_v)\|^2 - \|f(e_u) - f(e_v)\|^2 \right) \\ &\quad - \sum_{u,v} w_{uv} \log \sum_{u'} \exp(f(e_{u'})^T f(e_v)) + \mathcal{J}_{R0}. \end{aligned}$$

The last step holds because $\|x - y\|^2 = (x - y)^T (x - y) = \|x\|^2 + \|y\|^2 - 2x^T y$. Therefore, our goal is equivalent to

4. w_{uv} can be specific to a meta-path \mathcal{M} or an edge type l , in which cases we can denote it as $w_{uv}^{(\mathcal{M})}$ or $w_{uv}^{(l)}$ accordingly. In Eq. (4) and following derivations, for ease of notation, we omit the superscript.

$$\begin{aligned} \min - \mathcal{J} &= \sum_{u,v} \frac{w_{uv}}{2} \underbrace{\|f(e_u) - f(e_v)\|^2}_{d(e_u, e_v)} - \mathcal{J}_{R0} \\ &\quad - \underbrace{\sum_{u,v} \frac{w_{uv}}{2} \left(\|f(e_u)\|^2 + \|f(e_v)\|^2 \right)}_{\mathcal{J}_{R1}} \\ &\quad + \underbrace{\sum_{u,v} w_{uv} \log \sum_{u'} \exp(f(e_{u'})^T f(e_v))}_{\mathcal{J}_{R2}}. \end{aligned} \quad (5)$$

Here \mathcal{J}_{R1} and \mathcal{J}_{R2} are two regularization terms. Without \mathcal{J}_{R1} , $d(e_u, e_v)$ can be minimized by letting $\|f(e_u)\| \rightarrow 0$ ($\forall u \in V$); without \mathcal{J}_{R2} , $d(e_u, e_v)$ can be minimized by letting $e_u \equiv c$ ($\forall u \in V$). \mathcal{J}_R in Eq. (1) is then the sum of $-\mathcal{J}_{R0}$, $-\mathcal{J}_{R1}$ and \mathcal{J}_{R2} . Among them, $-\mathcal{J}_{R0}$ is algorithm-specific, while $-\mathcal{J}_{R1}$ and \mathcal{J}_{R2} are commonly shared across most HNE algorithms in the proximity-preserving group. We summarize the choices of w_{uv} , $d(e_u, e_v)$ and \mathcal{J}_{R0} in Table 1.

Although Eq. (4) can cover most existing and likely many future approaches, we would like to remark that there are also studies adopting other forms of proximity-preserving objectives. For example, SHINE [107] uses reconstruction loss of autoencoders; HINSE [64] adopts spectral embedding based on adjacency matrices with different meta-graphs; MetaDynaMix [108] introduces a low-rank matrix factorization framework for dynamic HIN embedding; HeGAN [72] proposes an adversarial learning approach with a relation type-aware discriminator.

3.2 Message-Passing Methods

Each node in a network can have attribute information represented as a feature vector x_u . Message-passing methods aim to learn node embeddings e_u based on x_u by aggregating the information from u 's neighbors. In recent studies, Graph Neural Networks (GNNs) [33] are widely adopted to facilitate this aggregation/message-passing process. Compared to the proximity based HNE methods, message-passing methods, especially the GNN based ones are often considered as deep network embedding, due to their multiple layers of learnable projection functions.

R-GCN [67]. R-GCN has K convolutional layers. The initial node representation $h_u^{(0)}$ is just the node feature x_u . In the k th convolutional layer, each representation vector is updated by accumulating the vectors of neighboring nodes through a normalized sum.

$$h_u^{(k+1)} = \sigma \left(\sum_{l \in \mathcal{T}_E} \sum_{v \in \mathcal{N}_l(u)} \frac{1}{|\mathcal{N}_l(u)|} W_l^{(k)} h_v^{(k)} + W_0^{(k)} h_u^{(k)} \right).$$

Different from the regular GCN model [33], R-GCN considers edge heterogeneity by learning multiple convolution matrices W 's, each of which corresponding to one edge type. During message passing, neighbors under the same edge type will be aggregated and normalized first. The node embedding is the output of the K th layer (i.e., $e_v = h_v^{(K)}$).

In unsupervised settings, message-passing approaches use link prediction as their downstream task to train GNNs. To be specific, the likelihood of observing edges in the heterogeneous network is maximized. R-GCN optimizes a cross-entropy loss through negative sampling. Essentially, it is the

approximation of the following objective:

$$\mathcal{J} = \sum_{l \in \mathcal{T}_E} \sum_{u, v \in V} w_{uv}^{(l)} \log \frac{\exp(\mathbf{e}_u^T \mathbf{A}_l \mathbf{e}_v)}{\sum_{u' \in V} \exp(\mathbf{e}_u^T \mathbf{A}_l \mathbf{e}_{v'})},$$

where $w_{uv}^{(l)} = \mathbf{1}_{(u,v) \in E_l}$.

Recently, CompGCN [109] extends R-GCN by leveraging a variety of entity-relation composition operations so as to jointly embed nodes and relations.

HAN [75]. Instead of considering one-hop neighbors, HAN utilizes meta-paths to model higher-order proximity. Given a meta-path \mathcal{M} , the representation of node u is aggregated from its meta-path based neighbors $\mathcal{N}_{\mathcal{M}}(u) = \{u\} \cup \{v | v \text{ connects with } u \text{ via the meta-path } \mathcal{M}\}$. HAN proposes an attention mechanism to learn the weights of different neighbors

$$\alpha_{uv}^{\mathcal{M}} = \frac{\exp(\sigma(\mathbf{a}_{\mathcal{M}}^T [\mathbf{x}'_u || \mathbf{x}'_v]))}{\sum_{v' \in \mathcal{N}_{\mathcal{M}}(u)} \exp(\sigma(\mathbf{a}_{\mathcal{M}}^T [\mathbf{x}'_u || \mathbf{x}'_{v'}]))},$$

$$\mathbf{h}_u^{\mathcal{M}} = \sigma \left(\sum_{v \in \mathcal{N}_{\mathcal{M}}(u)} \alpha_{uv}^{\mathcal{M}} \mathbf{x}'_v \right),$$

where $\mathbf{a}_{\mathcal{M}}$ is the node-level attention vector of \mathcal{M} ; $\mathbf{x}'_u = \mathbf{M}_{\phi(u)} \mathbf{x}_u$ is the projected feature vector of node u ; $||$ is the concatenation operator. Given the meta-path specific embedding $\mathbf{h}_u^{\mathcal{M}}$, HAN uses a semantic-level attention to weigh different meta-paths

$$\beta_{\mathcal{M}} = \frac{\exp(\frac{1}{|\mathcal{V}|} \sum_{v \in V} \mathbf{q}^T \tanh(\mathbf{W} \mathbf{h}_v^{\mathcal{M}} + \mathbf{b}))}{\sum_{\mathcal{M}'} \exp(\frac{1}{|\mathcal{V}|} \sum_{v \in V} \mathbf{q}^T \tanh(\mathbf{W} \mathbf{h}_v^{\mathcal{M}'} + \mathbf{b}))},$$

$$\mathbf{e}_u = \sum_{\mathcal{M}} \beta_{\mathcal{M}} \mathbf{h}_u^{\mathcal{M}},$$

where \mathbf{q} is the semantic-level attention vector.

In the original HAN paper, the authors mainly consider the task of semi-supervised node classification. For unsupervised learning (i.e., without any node labels), according to [112], HAN can use the link prediction loss introduced in GraphSAGE [34], which is the negative sampling approximation of the following objective:

$$\mathcal{J} = \sum_{u, v \in V} w_{uv} \log \frac{\exp(\mathbf{e}_u^T \mathbf{e}_v)}{\sum_{u' \in V} \exp(\mathbf{e}_u^T \mathbf{e}_{v'})}. \quad (6)$$

Here, w_{uv} is the edge weight of (u, v) .

MAGNN [112]. MAGNN extends HAN by considering both the meta-path based neighborhood $\mathcal{N}_{\mathcal{M}}(u) = \{v | v \text{ connects with } u \text{ via meta-path } \mathcal{M}\}$ and the nodes along the meta-path instances. Given a meta-path \mathcal{M} , MAGNN first employs an encoder to transform all the node features along an instance of \mathcal{M} into a single vector.

$$\mathbf{h}_{uv}^{\mathcal{M}} = f_{ENC}(\{\mathbf{x}'_t | t \in \mathcal{P}_{u \rightarrow v}^{\mathcal{M}}\}),$$

where $\mathbf{x}'_u = \mathbf{M}_{\phi(u)} \mathbf{x}_u$ is the projected feature vector of node u ; $\mathcal{P}_{u \rightarrow v}^{\mathcal{M}}$ denotes a meta-path instance of \mathcal{M} connecting u and v ; f_{ENC} is a relational rotation encoder inspired by [114]. After encoding each meta-path instance, MAGNN proposes an intra-meta-path aggregation to learn the weight of different neighbors

$$\alpha_{uv}^{\mathcal{M}} = \frac{\exp(\text{LeakyReLU}(\mathbf{a}_{\mathcal{M}}^T [\mathbf{x}'_v || \mathbf{h}_{uv}^{\mathcal{M}}]))}{\sum_{v' \in \mathcal{N}_{\mathcal{M}}(u)} \exp(\text{LeakyReLU}(\mathbf{a}_{\mathcal{M}}^T [\mathbf{x}'_{v'} || \mathbf{h}_{uv'}^{\mathcal{M}}]))},$$

$$\mathbf{h}_u^{\mathcal{M}} = \sigma \left(\sum_{v \in \mathcal{N}_{\mathcal{M}}(u)} \alpha_{uv}^{\mathcal{M}} \mathbf{h}_{uv}^{\mathcal{M}} \right),$$

where $\mathbf{a}_{\mathcal{M}}$ is the node-level attention vector of \mathcal{M} . This attention mechanism can also be extended to multiple heads. After aggregating the information within each meta-path, MAGNN further combines the semantics revealed by all meta-paths using an inter-meta-path aggregation

$$\beta_{\mathcal{M}} = \frac{\exp(\frac{1}{|\mathcal{V}_{\phi(u)}|} \sum_{v \in \mathcal{V}_{\phi(u)}} \mathbf{q}_{\phi(u)}^T \tanh(\mathbf{W}_{\phi(u)} \mathbf{h}_v^{\mathcal{M}} + \mathbf{b}_{\phi(u)}))}{\sum_{\mathcal{M}'} \exp(\frac{1}{|\mathcal{V}_{\phi(u)}|} \sum_{v \in \mathcal{V}_{\phi(u)}} \mathbf{q}_{\phi(u)}^T \tanh(\mathbf{W}_{\phi(u)} \mathbf{h}_v^{\mathcal{M}'} + \mathbf{b}_{\phi(u)}))},$$

$$\mathbf{e}_u = \sigma \left(\mathbf{W} \left(\sum_{\mathcal{M}} \beta_{\mathcal{M}} \mathbf{h}_u^{\mathcal{M}} \right) \right),$$

where $\mathbf{q}_{\phi(u)}$ is the semantic-level attention vector for node type $\phi(u)$. In comparison with HAN, MAGNN employs an additional projection to get the final representation \mathbf{e}_u .

For link prediction, MAGNN adopts the loss introduced in GraphSAGE [34], which is equivalent to the objective in Eq. (6).

HGT [85]. Inspired by the success of Transformer [115], [116] in text representation learning, Hu *et al.* propose to use each edge's type to parameterize the Transformer-like self-attention architecture. To be specific, for each edge (u, v) , their Heterogeneous Graph Transformer (HGT) model maps v into a *Query* vector, and u into a *Key* vector, and calculate their dot product as attention

$$\mathbf{Q}_v^i = \mathbf{Q}_{\phi(v)}^i(\mathbf{h}_v^{(k)}), \quad \mathbf{K}_u^i = \mathbf{K}_{\phi(u)}^i(\mathbf{h}_u^{(k)})$$

$$\text{Head}_i^{\text{ATT}}(u, v) = \left(\frac{\mathbf{K}_u^i \mathbf{W}_{\psi(u,v)}^{\text{ATT}} \mathbf{Q}_v^i}{\sqrt{d}} \right) \mu(\phi(u), \psi(u, v), \phi(v)),$$

$$\text{Attention}(u, v) = \text{Softmax}_{x_u \in N(v)} \left(\left\| \text{Head}_i^{\text{ATT}}(u, v) \right\| \right).$$

Here, $\mathbf{h}_u^{(k)}$ is the output of the k th HGT layer ($\mathbf{h}_u^{(0)} = \mathbf{x}_u$); $\mathbf{Q}_{\phi(v)}^i$ and $\mathbf{K}_{\phi(u)}^i$ are node type-aware linear mappings; $\text{Head}_i^{\text{ATT}}$ is the i th attention head; μ is a prior tensor representing the weight of each edge type in the attention. Parallel to the calculation of attention, the message passing process can be computed in a similar way by incorporating node and edge types

$$\text{Head}_i^{\text{MSG}}(u, v) = \mathbf{M}_{\phi(u)}^i(\mathbf{h}_u^{(k)}) \mathbf{W}_{\psi(u,v)}^{\text{MSG}},$$

$$\text{Message}(u, v) = \left\| \text{Head}_i^{\text{MSG}}(u, v) \right\|,$$

where $\mathbf{M}_{\phi(u)}^i$ is also a node type-aware linear mapping. To aggregate the messages from v 's neighborhood, the attention vector serves as the weight to get the updated vector

$$\hat{\mathbf{h}}_v^{(k+1)} = \sum_{u \in N(v)} \text{Attention}(u, v) \odot \text{Message}(u, v),$$

and following the residual connection [117], the output of the k th layer is

$$\mathbf{h}_v^{(k+1)} = \mathbf{A}_{\phi(v)}(\sigma(\hat{\mathbf{h}}_v^{(k+1)})) + \mathbf{h}_v^{(k)}.$$

TABLE 2
A Summary of Message-Passing Based HNE Algorithms

Algorithm	$w_{uv} / w_{uv}^{(l)} / w_{uv}^{\mathcal{M}}$	$d(e_u, e_v)$	Aggregation Function	Applications
R-GCN [67]	$\mathbf{1}_{(u,v) \in E_l}$	$\ \sqrt{A_l}(e_u - e_v)\ ^2$	$\mathbf{h}_u^{(k+1)} = \sigma\left(\sum_{l \in \mathcal{T}_E} \sum_{v \in \mathcal{N}_l(u)} \frac{1}{ \mathcal{N}_l(u) } \mathbf{W}_r^{(k)} \mathbf{h}_v^{(k)} + \mathbf{W}_0^{(k)} \mathbf{h}_u^{(k)}\right)$ $\mathbf{h}_u^{(0)} = \mathbf{x}_u, \mathbf{e}_u = \mathbf{h}_u^{(K)}$	entity classification, KB completion
HEP [110]	$\mathbf{1}_{(u,v) \in E}$	$\ \sqrt{A}(e_u - e_v)\ ^2$	$\tilde{e}_u = \sigma\left(\mathbf{W}_{\phi(u)} \left(\sum_{o \in \mathcal{T}_V} \sum_{v \in \mathcal{N}_o(u)} \alpha_{uv} e_v\right) + \mathbf{b}_{\phi(u)}\right)$	user alignment
HAN [75]	edge weight of (u, v)	$\ e_u - e_v\ ^2$	$e_u = \sum_{\mathcal{M}} \beta_{\mathcal{M}} \sigma\left(\sum_{v \in \mathcal{N}_{\mathcal{M}}(u)} \alpha_{uv}^{\mathcal{M}} \mathbf{M}_{\phi(u)} \mathbf{x}_v\right)$	node classification, node clustering, link prediction, recommendation
HetGNN [71]	number of times that $u \in \mathcal{C}(v)$ in homogeneous random walks		$\mathbf{h}_u = f_{\text{ENC}}(\mathbf{x}_u), \mathbf{h}_{u,o} = f_{\text{AGG}}(\{\mathbf{h}_v v \in \mathcal{N}_{\text{RWR}}(u), \phi(v) = o\})$ $e_u = \alpha_u \mathbf{h}_u + \sum_{o \in \mathcal{T}_V} \alpha_{u,o} \mathbf{h}_{u,o}$	
GATNE [70]	number of times that $u \in \mathcal{C}(v)$ in random walks following \mathcal{M}	$\ e_{u,l} - e_v\ ^2$	$\mathbf{h}_{u,l}^{(k+1)} = f_{\text{AGG}}(\{\mathbf{h}_{v,l}^{(k)} v \in \mathcal{N}_l(u)\}), \mathbf{h}_{u,l}^{(0)} = \mathbf{x}_u$ $e_{u,l} = \alpha_l \mathbf{W}_l \left(\sum_{l' \in \mathcal{T}_E} \mathbf{h}_{u,l'}^{(K)}\right) \mathbf{a}_{u,l} + \mathbf{b}_u$	node classification, node clustering, link prediction, recommendation
MV-ACM [111]	number of times that $u \in \mathcal{C}(v)$ in random walks following \mathcal{M}		$\mathbf{h}_{u,l}^{(k+1)} = \sigma\left(\mathbf{W}_l^{(k)} \frac{1}{ \mathcal{N}_l(u) } \sum_{v \in \mathcal{N}_l(u)} \mathbf{h}_{v,l}^{(k)}\right), \mathbf{h}_{u,l}^{(0)} = \mathbf{x}_u$ $e_{u,l} = \mathbf{M}_l \left(\mathbf{h}_{u,l}^{(K)} + \sum_{l' \in \mathcal{T}_E} \alpha_{u,l,l'} \mathbf{h}_{u,l'}^{(K)}\right) + \mathbf{b}_u$	
MAGNN [112]	edge weight of (u, v)	$\ e_u - e_v\ ^2$	$\mathbf{h}_u^{\mathcal{M}} = \sigma\left(\sum_{v \in \mathcal{N}_{\mathcal{M}}(u)} \alpha_{uv}^{\mathcal{M}} f_{\text{ENC}}(\{\mathbf{M}_{\phi(t)} \mathbf{x}_t t \in \mathcal{P}_{u \rightarrow v}^{\mathcal{M}}\})\right)$ $e_u = \sigma\left(\mathbf{W} \left(\sum_{\mathcal{M}} \beta_{\mathcal{M}} \mathbf{h}_u^{\mathcal{M}}\right)\right)$	
HGT [85]	$\mathbf{1}_{(u,v) \in E}$	refer to NTN [20] in Table 3	$\hat{\mathbf{h}}_v^{(k+1)} = \sum_{u \in \mathcal{N}(v)} \text{Attention}(u, v) \odot \text{Message}(u, v)$ $\mathbf{h}_v^{(k+1)} = A_{\phi(v)}(\sigma(\hat{\mathbf{h}}_v^{(k+1)})) + \mathbf{h}_v^{(k)}, \mathbf{h}_u^{(0)} = \mathbf{x}_u, \mathbf{e}_u = \mathbf{h}_u^{(K)}$	node classification, author identification

(Additional notations: $\mathcal{N}_{\text{RWR}}(v)$: neighbors of v defined through random walk with restart [113]; $\mathcal{N}_l(u)$: neighbors of u with edge type l ; $\mathcal{N}_o(u)$: neighbors of u with node type o ; $\mathcal{N}_{\mathcal{M}}(u)$: nodes connects with u via meta-path \mathcal{M} ; $\mathcal{P}_{u \rightarrow v}^{\mathcal{M}}$: a meta-path instance of \mathcal{M} connecting u and v ; f_{AGG} : a function to aggregate information from neighbors. We show the transductive version of GATNE.)

Here, $A_{\phi(u)}$ is a linear function mapping v 's vector back to its node type-specific distribution.

For the unsupervised link prediction task, HGT borrows the objective function from Neural Tensor Network [20], which will be further discussed in Section 3.3.

Some other message-passing approaches are summarized in Table 2. For example, HEP [110] aggregates u 's representation from $\mathcal{N}_o(u)$ (i.e., the *node type-aware* neighborhood) to reconstruct u 's own embedding; HetGNN [71] also adopts a node type-aware neighborhood aggregation, but the neighborhood is defined through random walk with restart [113]; GATNE [70] aggregates u 's representation from $\mathcal{N}_l(u)$ (i.e., the *edge type-aware* neighborhood) and is applicable to both transductive and inductive network embedding settings; MV-ACM [111] aggregates u 's representation from $\mathcal{N}_{\mathcal{M}}(u)$ (i.e., the *meta-path-aware* neighborhood) and utilizes an adversarial learning framework to learn the reciprocity between different edge types.

The objective of message-passing approaches mentioned above can also be written as Eq. (4) (except HGT, whose objective is the same as NTN [20] and will be discussed in Section 3.3), where $s(u, v)$ is a function of e_u and e_v . The only difference is that e_u here is aggregated from x_v using GNNs. Following the derivation of proximity-preserving approaches, if we still write \mathcal{J}_R in Eq. (1) as the sum of $-\mathcal{J}_{R_0}, -\mathcal{J}_{R_1}$ and \mathcal{J}_{R_2} , we can get the exactly same \mathcal{J}_{R_1} and \mathcal{J}_{R_2} as in Eq. (5). We summarize the choices of w_{uv} , $d(e_u, e_v)$ and the aggregation function in Table 2.

Within this group of algorithms, HEP has an additional reconstruction loss $\mathcal{J}_{R_0} = \sum_v \|\tilde{e}_v - e_v\|^2$, and MV-ACM [111] has an adversarial loss $\mathcal{J}_{R_0} = \sum_{v \in V, l \in \mathcal{T}_E} \mathbb{E}_{Z \sim \mathcal{G}(\cdot | l, v)} [\log(1 - \mathcal{D}(Z, l, v))]$, where \mathcal{G} and \mathcal{D} are the generator and discriminator of complementary edge types, respectively. All the other algorithms in this group have $\mathcal{J}_{R_0} = 0$.

Similar to the case of proximity-preserving approaches, there are also message-passing methods adopting other forms of objectives. For example, GraphInception [66]

studies collective classification (where the labels within a group of instances are correlated and should be inferred collectively) using meta-path-aware aggregation; HDGI [86] maximizes local-global mutual information to improve unsupervised training based on HAN; NEP [74] performs semi-supervised node classification using an edge type-aware propagation function to mimic the process of label propagation; NLAH [124] also considers semi-supervised node classification by extending HAN with several pre-computed non-local features; HGCN [125] explores the collective classification task and improves GraphInception by considering semantics of different types of edges/nodes and relations among different node types; HetETA [126] studies a specific task of estimating the time of arrival in intelligent transportation using a heterogeneous graph network with fast localized spectral filtering.

3.3 Relation-Learning Methods

As discussed above, knowledge graphs can be regarded as a special case of heterogeneous networks, which are schema-rich [127]. To model network heterogeneity, existing knowledge graph embedding approaches explicitly model the relation types of edges via parametric algebraic operators, which are often shallow network embedding [79], [128]. Compared to the shallow proximity-preserving HNE models, they often focus on the designs of triplet based scoring functions instead of meta-paths or meta-graphs due to the large numbers of entity and relation types.

Each edge in a heterogeneous network can be viewed as a triplet (u, l, v) composed of two nodes $u, v \in V$ and an edge type $l \in \mathcal{T}_E$ (i.e., entities and relations, in the terminology of KG). The goal of relation-learning methods is to learn a scoring function $s_l(u, v)$ which evaluates an arbitrary triplet and outputs a scalar to measure the acceptability of this triplet. This idea is widely adopted in KB embedding. Since there are surveys of KB embedding algorithms already [79],

we only cover the most popular approaches here and highlight their connections to HNE.

TransE [4]. TransE assumes that the relation induced by l -labeled edges corresponds to a translation of the embedding (i.e., $e_u + e_l \approx e_v$) when (u, l, v) holds. Therefore, the scoring function of TransE is defined as

$$s_l(u, v) = -\|e_u + e_l - e_v\|_p, \quad (7)$$

where $p = 1$ or 2 . The objective is to minimize a margin based ranking loss.

$$\mathcal{J} = \sum_{(u,l,v) \in T} \sum_{(u',l',v') \in T'_{(u,l,v)}} \max(0, \gamma - s_l(u, v) + s_l(u', v')), \quad (8)$$

where T is the set of positive triplets (i.e., edges); $T'_{(u,l,v)}$ is the set of corrupted triplets, which are constructed by replacing either u or v with an arbitrary node. Formally,

$$T'_{(u,l,v)} = \{(u', l, v) | u' \in V\} \cup \{(u, l, v') | v' \in V\}.$$

TransE is the most representative model using “a translational distance” to define the scoring function. It has many extensions. For example, TransH [118] projects each entity vector to a relation-specific hyperplane when calculating the distance; TransR [80] further extends relation-specific hyperplanes to relation-specific spaces; RHINE [76] distinguishes affiliation relations from interaction relations and adopts different objectives for the two types of relations. For more extensions of TransE, please refer to [83].

Recently, Sun *et al.* [114] propose the RotatE model, which defines each relation as a rotation (instead of a translation) from the source entity to the target entity in the complex vector space. Their model is able to describe various relation patterns including symmetry/antisymmetry, inversion, and composition.

DistMult [81]. In contrast to translational distance models [4], [80], [118], DistMult exploits a similarity based scoring function. Each relation is represented by a diagonal matrix $A_l = \text{diag}(e_{l1}, \dots, e_{ld})$, and $s_l(u, v)$ is defined using a bilinear function

$$s_l(u, v) = e_u^T A_l e_v.$$

Note that $s_l(u, v) = s_l(v, u)$ for any u and v . Therefore, DistMult is mainly designed for symmetric relations.

Using the equation $\|x - y\|^2 = (x - y)^T(x - y) = \|x\|^2 + \|y\|^2 - 2x^T y$, we have

$$\begin{aligned} s_l(u, v) &= (\sqrt{A_l} e_u)^T (\sqrt{A_l} e_v) \\ &= \frac{1}{2} (\|\sqrt{A_l} e_u\|^2 + \|\sqrt{A_l} e_v\|^2 - \|\sqrt{A_l} (e_u - e_v)\|^2). \end{aligned}$$

Complex [121]. Instead of considering a real-valued embedding space, ComplEx introduces complex-valued representations for e_u , e_v and e_l . Similar to DistMult, it utilizes a similarity based scoring function

$$s_l(u, v) = \text{Re}(e_u^T A_l \bar{e}_v),$$

where $\text{Re}(\cdot)$ is the real part of a complex number; \bar{e} is the complex conjugate of e ; $A_l = \text{diag}(e_{l1}, \dots, e_{ld})$. Here, it is

possible that $s_l(u, v) \neq s_l(v, u)$, which allows ComplEx to capture asymmetric relations.

In the complex space, using the equation $\|x - y\|^2 = (x - y)^T(\bar{x} - \bar{y}) = \|x\|^2 + \|y\|^2 - x^T \bar{y} - y^T \bar{x} = \|x\|^2 + \|y\|^2 - 2\text{Re}(x^T \bar{y})$, we have

$$\begin{aligned} s_l(u, v) &= \text{Re}((A_l e_u)^T \bar{e}_v) \\ &= \frac{1}{2} (\|A_l e_u\|^2 + \|e_v\|^2 - \|A_l e_u - e_v\|^2). \end{aligned}$$

Besides ComplEx, there are many other extensions of DistMult. RESCAL [119] uses a bilinear scoring function similar to the one in DistMult, but A_l is no longer restricted to be diagonal; HolE [120] composes the node representations using the circular correlation operation, which combines the expressive power of RESCAL with the efficiency of DistMult; Simple [122] considers the inverse of relations and calculates the average score of (u, l, v) and (v, l^{-1}, u) ; TuckER [123] proposes to use a three-way Tucker tensor decomposition approach to learning node and relation embeddings.

ConvE [82]. ConvE goes beyond simple distance or similarity functions and proposes deep neural models to score a triplet. The score is defined by a convolution over 2D shaped embeddings. Formally,

$$s_l(u, v) = \sigma(\text{vec}(\sigma([E_u; E_r] * \omega))W)e_v,$$

where E_u and E_r denote the 2D reshaping matrices of node embedding and relation embedding, respectively; vec is the vectorization operator that maps a m by n matrix to a mn -dimensional vector; “*” is the convolution operator.

There are several other models leveraging deep neural scoring functions. For example, NTN [20] proposes to combine the two node embedding vectors by a relation-specific tensor $\mathcal{M}_l \in \mathbb{R}^{d \times d \times d_l}$, where d_l is the dimension of e_l ; NKG [83] develops a deep memory network to encode information from neighbors and employs a gating mechanism to integrate structure representation e_u^s and neighbor representation e_u^n ; SACN [84] encodes node representations using a graph neural network and then scores the triplet using a convolutional neural network with the translational property.

Most relation-learning approaches adopt a margin based ranking loss with some regularization terms that generalizes Eq. (8)

$$\sum_{(u,l,v)} w_{uv}^{(l)} \sum_{(u',l',v')} \max(0, \gamma - s_l(u, v) + s_l(u', v')) + \mathcal{J}_{R0}. \quad (9)$$

In [129], Qiu *et al.* point out that the margin based loss shares a very similar form with the following negative sampling loss:

$$- \sum_{(u,l,v)} \left(\log(\sigma(s_l(u, v))) - bE_{(u',l',v')} [\log(\sigma(s_l(u', v')))] \right).$$

Following [129], if we use the negative sampling loss to rewrite Eq. (9), we are approximately maximizing

$$\mathcal{J} = \sum_{(u,l,v)} w_{uv}^{(l)} \log \frac{\exp(s_l(u, v))}{\sum_{(u',l',v')} \exp(s_l(u', v'))} + \mathcal{J}_{R0}.$$

For translational models [4], [76], [80], [118] whose $s_l(u, v)$ is described by a distance function, maximizing \mathcal{J} is

TABLE 3
A Summary of Relation-Learning Based HNE Algorithms

Algorithm	$w_{uv}^{(l)}$	$d(e_u, e_v)$	\mathcal{J}_{R0}	Applications
TransE [4]	$\mathbf{1}_{(u,v) \in E_l}$	$\ e_u + e_l - e_v\ $	$\sum_v (\ e_v\ - 1)$	KB completion, relation extraction from text
TransH [118]		$\ e_{u,l} + e_l - e_{v,l}\ ^2,$ $e_{v,l} = e_v - w_l^T e_v w_l$	$\sum_l (\ w_l\ - 1) + \sum_v (\ e_v\ - 1) +$ $+\sum_l \left[\frac{(w_l^T e_l)^2 - \epsilon^2}{\ e_l\ ^2} \right]_+$	
TransR [80]		$\ e_{u,l} + e_l - e_{v,l}\ ^2,$ $e_{v,l} = A_l e_v$	$\sum_v (\ e_v\ - 1) + \sum_l (\ A_l\ - 1) +$ $+\sum_v (\ A_l e_v\ - 1)_+$	
RHINE [76]	edge weight of (u, v) with type l	$\ e_u - e_v\ ^2$ if l models affiliation, $\ e_u + e_l - e_v\ $ if l models interaction	N/A	link prediction, node classification
RotatE [114]	$\mathbf{1}_{(u,v) \in E_l}$	$\ e_u \odot e_l - e_v\ ^2, e_u, e_v, e_l \in \mathbb{C}^n$	$\sum_l (\ e_l\ - 1)$	KB completion, relation pattern inference
RESCAL [119]		$\ \sqrt{A_l}(e_u - e_v)\ ^2$	$\sum_v (\ e_v\ - 1) + \sum_l (\ A_l\ _F - 1)_+$	entity resolution, link prediction
DistMult [81]		$\ \sqrt{A_l}(e_u - e_v)\ ^2, A_l = \text{diag}(e_l)$	$\sum_v (\ e_v\ - 1) + \sum_l (\ e_l\ - 1)_+$	KB completion, triplet classification
HoLE [120]		$\ e_r - \mathcal{F}^{-1}(\mathcal{F}(e_u) \odot \mathcal{F}(e_v))\ ^2$	$\sum_v (\ e_v\ - 1) + \sum_l (\ e_l\ - 1)_+$	
CompLex [121]		$\ A_l e_u - e_v\ ^2, A_l = \text{diag}(e_l), e_u, e_v, e_l \in \mathbb{C}^n$	$\sum_v \ e_v\ ^2 + \sum_l \ e_l\ ^2$	
Simple [122]		$\frac{1}{2} (\ \sqrt{A_l}(e_u - e_v)\ ^2 + \ \sqrt{A_{l-1}}(e_v - e_u)\ ^2),$ $A_l = \text{diag}(e_l), A_{l-1} = \text{diag}(e_{l-1})$		
TuckER [123]		$C - W \times_1 e_u \times_2 e_l \times_3 e_v$	N/A	
NTN [20]		$C - e_l^T \tanh(e_u^T M_l e_v + M_{l,1} e_u + M_{l,2} e_v + b_l)$	$\ \Theta\ _2^2$	
ConvE [82]		$C - \sigma(\text{vec}(\sigma([E_u; E_l] * w))W) e_v$	N/A	
NKGE [83]		same as TransE or ConvE, where $e_u = \sigma(g_u) \odot e_u^s + (1 - \sigma(g_u)) \odot e_u^n$	$\ \Theta\ _2^2$	
SACN [84]	$C - f(\text{vec}(M(e_u, e_l))W) e_v$	N/A		

(Additional notations: f : a non-linear activation function; $[x]_+$: $\max\{x, 0\}$; $\|\cdot\|_F$: the Frobenius norm of a matrix; $\mathcal{F}, \mathcal{F}^{-1}$: the fast Fourier transform and its inverse; \bar{x} : the complex conjugate of x ; l^{-1} : the inverse of relation l ; \times_i : the tensor product along the i th mode; Θ : the set of learned parameters; E_u, E_l : 2D reshaping matrices of e_u and e_l [82]; vec : the vectorization operator; $*$: the convolution operator; $M(e_u, e_l)$: a matrix aligning the output vectors from the convolution with all kernels [84].)

equivalent to

$$\min - \mathcal{J} = \sum_{(u,l,v)} w_{uv}^{(l)} \underbrace{\|e_u + e_l - e_v\|^{1 \text{ or } 2}}_{d(e_u, e_v)} - \mathcal{J}_{R0} + \underbrace{\sum_{(u,l,v)} w_{uv}^{(l)} \log \sum_{(u',l',v')} \exp(s_l(u', v'))}_{\mathcal{J}_{R1}}. \quad (10)$$

In this case, we can write \mathcal{J}_R as $-\mathcal{J}_{R0} + \mathcal{J}_{R1}$. For RotatE [114], the objective is the same except that $d(e_u, e_v) = \|e_u \odot e_l - e_v\|^2$.

For similarity based approaches [81], [119], [121], [122], we follow the derivation of Eq. (5), and the objective will be

$$\min - \mathcal{J} = \sum_{(u,l,v)} \frac{w_{uv}^{(l)}}{2} \underbrace{\|f(e_u) - f(e_v)\|^2}_{d(e_u, e_v)} - \mathcal{J}_{R0} - \underbrace{\sum_{(u,l,v)} \frac{w_{uv}^{(l)}}{2} (\|f(e_u)\|^2 + \|f(e_v)\|^2)}_{\mathcal{J}_{R1}} + \underbrace{\sum_{(u,l,v)} w_{uv}^{(l)} \log \sum_{(u',l',v')} \exp(s_l(u', v'))}_{\mathcal{J}_{R2}}.$$

Here, $f(e_u) = \sqrt{A_l} e_u$ in RESCAL [119] and DistMult [81]; $f(e_u) = \sqrt{A_l} e_u$ or $\sqrt{A_{l-1}} e_u$ in Simple [122]; $f(e_u) = A_l e_u$ and $f(e_v) = e_v$ in CompLex [121]. The regularization term is $\mathcal{J}_R = -\mathcal{J}_{R0} - \mathcal{J}_{R1} + \mathcal{J}_{R2}$.

For neural triplet scorers [20], [21], [82], [84], the forms of $s_l(u, v)$ are more complicated than translational distances or

bilinear products. In these cases, since distance (or dissimilarity) and proximity can be viewed as reverse metrics, we define $d(e_u, e_v)$ as $C - s_l(u, v)$, where C is an constant upper bound of $s_l(\cdot, \cdot)$. Then the derivation of the loss function is similar to that of Eq. (10), i.e.,

$$\min - \mathcal{J} = \sum_{(u,l,v)} w_{uv}^{(l)} \underbrace{(C - s_l(u, v))}_{d(e_u, e_v)} - \mathcal{J}_{R0} + \underbrace{\sum_{(u,l,v)} w_{uv}^{(l)} \log \sum_{(u',l',v')} \exp(s_l(u', v'))}_{\mathcal{J}_{R1}}.$$

In this case, $\mathcal{J}_R = -\mathcal{J}_{R0} + \mathcal{J}_{R1}$.

We summarize the choices of $w_{uv}^{(l)}$, $d(e_u, e_v)$ and \mathcal{J}_{R0} in Table 3. Note that for relation learning methods, $d(\cdot, \cdot)$ is usually *not* a distance metric. For example, $d(e_u, e_v) \neq d(e_v, e_u)$ in most translational distance models and deep neural models. This is intuitive because (u, l, v) and (v, l, u) often express different meanings.

4 BENCHMARK

4.1 Dataset Preparation

Towards off-the-shelf evaluation of HNE algorithms with standard settings, in this work, we collect, process, analyze, and publish four new real-world heterogeneous network datasets from different domains, which we aim to set up as a handy and fair benchmark for existing and future HNE algorithms.²

DBLP. We construct a network of authors, papers, venues, and phrases from DBLP. Phrases are extracted by the

TABLE 4
A Summary of the Statistics on Four Real-World Heterogeneous Network Datasets

Dataset	#node type	#node	#link type	#link	#attributes	#attributed nodes	#label type	#labeled node
DBLP	4	1,989,077	6	275,940,913	300	ALL	13	618
Yelp	4	82,465	4	30,542,675	N/A	N/A	16	7,417
Freebase	8	12,164,758	36	62,982,566	N/A	N/A	8	47,190
PubMed	4	63,109	10	244,986	200	ALL	8	454




Fig. 3. Portion of different node types in four real-world heterogeneous network datasets.

popular AutoPhrase [130] algorithm from paper texts and further filtered by human experts. We compute word2vec [99] on all paper texts and aggregate the word embeddings to get 300-dim paper and phrase features. Author and venue features are the aggregations of their corresponding paper features. We further manually label a relatively small portion of authors into 12 research groups from four research areas by crawling the web. Each labeled author has only one label.

Yelp. We construct a network of businesses, users, locations, and reviews from Yelp.⁵ Nodes do not have features, but a large portion of businesses are labeled into sixteen categories. Each labeled business has one or multiple labels.

Freebase. We construct a network of books, films, music, sports, people, locations, organizations, and businesses from Freebase.⁶ Nodes are not associated with any features, but a large portion of books are labeled into eight genres of literature. Each labeled book has only one label.

PubMed. We construct a network of genes, diseases, chemicals, and species from PubMed.⁷ All nodes are extracted by AutoPhrase [130], typed by bioNER [131], and further filtered by human experts. The links are constructed through open relation pattern mining [132] and manual selection. We compute word2vec [99] on all PubMed papers and aggregate the word embeddings to get 200-dim features for all types of nodes. We further label a relatively small portion of diseases into eight categories. Each labeled disease has only one label.

The four datasets we prepare are from four different domains, which have been individually studied in some existing works on HNE. Among them, DBLP has been most commonly used, because all information about authors, papers, etc. is public and there is no privacy issue, and the results are often more interpretable to researchers in computer science related domains. Other real-world networks like Yelp and IMDB have been commonly studied for recommender systems. These networks are naturally heterogeneous, including at least users and items (e.g., businesses and movies), as well as some additional item descriptors

(e.g., categories of businesses and genres of movies). Freebase is one of the most popular open-source knowledge graph, which is relatively smaller but cleaner compared with the others (e.g., YAGO [133] and Wikidata [134]), where most entity and relation types are well defined. One major difference between conventional heterogeneous networks and knowledge graphs is the number of types of nodes and links. We further restrict the types of entities and relations inside Freebase, so as to get a heterogeneous network that is closer to a knowledge graph, while in the meantime does not have too many types of nodes and links. Therefore, most conventional HNE algorithms can be applied on this dataset and properly compared against the KB embedding ones. PubMed is a novel biomedical network we directly construct through text mining and manual processing on biomedical literature. This is the first time we make it available to the public, and we hope it to serve both the evaluation of HNE algorithms and novel downstream tasks in biomedical science such as biomedical information retrieval and disease evolution study.

We notice that several other heterogeneous network datasets such as OAG [85], [87], [135] and IMDB [75], [112], [125] have been constructed recently in parallel to this work. Due to their similar nature and organization as some of our datasets (e.g., DBLP, Yelp), our pipeline can be easily adopted on these datasets, so we do not copy them here.

4.2 Structure Analysis

A summary of the statistics on the four datasets is provided in Table 4 and Fig. 3. As can be observed, the datasets have different sizes (numbers of nodes and links) and heterogeneity (numbers and distributions of node/link types). Moreover, due to the nature of the data sources, DBLP and PubMed networks are attributed, whereas Yelp and Freebase networks are abundantly labeled. A combination of these four datasets thus allows researchers to flexibly start by testing an HNE algorithm in the most appropriate settings, and eventually complete an all-around evaluation over all settings.

We also provide detailed analysis regarding several most widely concerned properties of heterogeneous networks, i.e., degree distribution (Fig. 4), clustering coefficient (Fig. 5), and

5. <https://www.yelp.com/dataset/challenge>

6. <http://www.freebase.com/>

7. <https://www.ncbi.nlm.nih.gov/pubmed/>

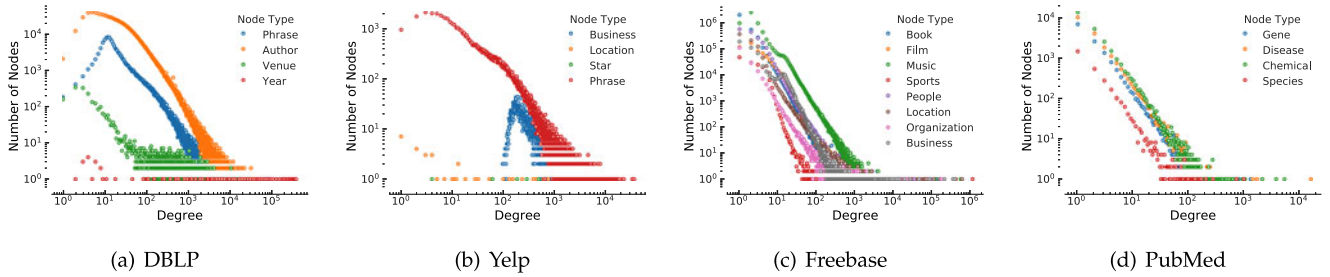


Fig. 4. Degree distribution of different node types in four real-world heterogeneous network datasets.

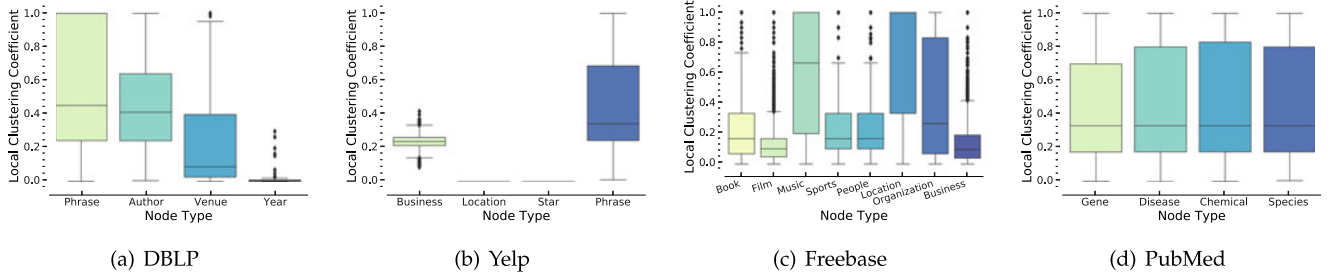


Fig. 5. Local clustering coefficient of different node types in four real-world heterogeneous networks.




Fig. 6. Number of five most frequent 2-hop meta-paths in four real-world heterogeneous network datasets.

number of frequent meta-paths (Fig. 6). In particular, degree distribution is known to significantly influence the performance of HNE algorithms due to the widely used node sampling process, whereas clustering coefficient impacts HNE algorithms that utilize latent community structures. Moreover, since many HNE algorithms rely on meta-paths, the skewer distribution of meta-paths can bias towards algorithms using fewer meta-paths.

As we can see, the properties we concern are rather different across the four datasets we prepare. For example, there are tighter links and more labels in Yelp, while there are more types of nodes and links in Freebase; compared with nodes in Freebase and PubMed which clearly follow the long-tail degree distribution, certain types of nodes in DBLP and Yelp are always well connected (e.g., phrases in DBLP and businesses in Yelp), forming more *star-shaped* subgraphs; the type-wise clustering coefficients and meta-path distributions are the most skewed in DBLP and most balanced in PubMed. The set of four datasets together provide a comprehensive benchmark towards the robustness and generalizability of various HNE algorithms (as we will also see in Section 5.2).

4.3 Settings, Tasks, and Metrics

We mainly compare all 13 algorithms under the setting of unsupervised unattributed HNE over all datasets, where the essential goal is to preserve different types of edges in

the heterogeneous networks. Moreover, for message-passing algorithms that are particularly designed for attributed and semi-supervised HNE, we also conduct additional experiments for them in the corresponding settings. Particularly, due to the nature of the datasets, we evaluate attributed HNE on DBLP and PubMed datasets where node attributes are available, and semi-supervised HNE on Yelp and Freebase where node labels are abundant. We always test the computed network embeddings on the two standard network mining tasks of node classification and link prediction. Note that, while most existing studies on novel HNE algorithms have been focusing on these two standard tasks [59], [61], [63], [75], [85], [103], [112], we notice that there are also various other tasks due to the wide usage of heterogeneous networks in real-world applications [11], [126], [136], [137], [138], [139], [140]. While the performance of HNE there can be rather task-dependant and data-dependant, to first simplifying the task into either a standard node classification or link prediction problem can often serve to provide more insights into the task and dataset, which can help the further development of novel HNE algorithms.

For the standard unattributed unsupervised HNE setting, we first randomly hide 20 percent links and train all HNE algorithms with the remaining 80 percent links. For node classification, we then train a separate linear Support Vector Machine (LinearSVC) [141] based on the learned embeddings on 80 percent of the labeled nodes and predict on the remaining 20 percent. We repeat the process for five

times and compute the average scores regarding macro-F1 (across all labels) and micro-F1 (across all nodes). For link prediction, we use the Hadamard function to construct feature vectors for node pairs, train a two-class LinearSVC on the 80 percent training links and evaluate towards the 20 percent held out links. We also repeat the process for five times and compute the two metrics of AUC (area under the ROC curve) and MRR (mean reciprocal rank). AUC is a standard measure for classification, where we regard link prediction as a binary classification problem, and MRR is a standard measure for ranking, where we regard link prediction as a link retrieval problem. Since exhaustive computation over all node pairs is too heavy, we always use the two-hop neighbors as the candidates for all nodes. For attributed HNE, node features are used during the training of HNE algorithms, whereas for semi-supervised HNE, certain amounts of node labels are used (80 percent by default).

5 EXPERIMENTAL EVALUATIONS

5.1 Algorithms and Modifications

We amend the implementations of 13 popular HNE algorithms for seamless and efficient experimental evaluations on our prepared datasets. The algorithms we choose and the modifications we make are as follows.

- *metapath2vec* [59]: Since the original implementation contains a large amount of hard-coded data-specific settings such as node types and meta-paths, and the optimization is unstable and limited as it only examines one type of meta-path based context, we completely reimplement the algorithm. In particular, we first run random walks to learn the weights of different meta-paths based on the number of sampled instances, and then train the model using the unified loss function, which is a weighted sum over the loss functions of individual meta-paths. Both the random walk and meta-path-based embedding optimization are implemented with multi-threads in parallel.
- *PTE* [103]: Instead of accepting labeled texts as input and working on text networks with the specific three types of nodes (word, document, and label) and three types of links (word-word, document-word, and label-word), we revise the original implementation and allow the model to consume heterogeneous networks directly with arbitrary types of nodes and links by adding more type-specific objective functions.
- *HIN2Vec* [63]: We remove unnecessary data preprocessing codes and modify the original implementation so that the program first generates random walks, then trains the model, and finally outputs node embeddings only.
- *AspEm* [60]: We clean up the hard-coded data-specific settings in the original implementation and write a script to connect the different components of automatically selecting the aspects with the least incompatibilities, as well as learning, matching, and concatenating the embeddings based on different aspects.
- *HEER* [61]: We remove the hard-coded data-specific settings and largely simplify the data preprocessing step in the original implementation by skipping the

knockout step and disentangling the graph building step.

- *R-GCN* [67]: The existing implementation from DGL [142] is only scalable to heterogeneous networks with thousands of nodes, due to the requirement of putting the whole graphs into memory during graph convolutions. To scale up R-GCN, we perform fixed-sized node and link sampling for batch-wise training following the framework of GraphSAGE [34].
- *HAN* [75]: Since the original implementation of HAN contains a large amount of hard-coded data-specific settings such as node types and meta-paths, and is unfeasible for large-scale datasets due to the same reason as R-GCN, we completely reimplement the HAN algorithm based on our implementation of R-GCN. In particular, we first automatically construct meta-path based adjacency lists for the chosen node type, and then sample the neighborhood for the seed nodes during batch-wise training.
- *MAGNN* [112]: The original DGL-based [142] unsupervised implementation of MAGNN solely looks at predicting the links between users and artists in a music website dataset with the use of a single neural layer. Hence, we remove all the hard-coded data-specific settings and refactor the entire pipeline so that the model can support multi-layer mini-batch training over arbitrary link types.
- *HGT* [85]: The original PyG-based [143] implementation of HGT targets on the specific task of author disambiguation among the associated papers in a dynamic academic graph [135]. Therefore, we refactor it by removing hard-coded data-specific settings, assigning the same timestamp to all the nodes, and conducting training over all types of links.
- *TransE* [4]: We modify the OpenKE [144] implementation so that the model outputs node embeddings only.
- *DistMult* [81]: We remove the hard-coded data-specific settings and largely simplify the data preprocessing step in the original implementation.
- *ComplEx* [121]: Same as for TransE.
- *ConvE* [82]: Same as for DistMult.

We set the embedding size of all algorithms to 50 by default, and tune other hyperparameters following the original papers through standard five-fold cross validation on all datasets. We have put the implementation of all compared algorithms in a python package and released them together with the datasets to constitute an open-source ready-to-use HNE benchmark.²

5.2 Performance Benchmarks

We provide systematic experimental comparisons of the 13 popular state-of-the-art HNE algorithms across our four datasets, on the scenarios of unsupervised unattributed HNE, attributed HNE, and semi-supervised HNE.

Table 5 shows the performance of compared algorithms on unsupervised unattributed HNE, evaluated towards node classification and link prediction. We have the following observations.

From the perspective of compared algorithms:

TABLE 5
Performance Comparison (%) Under the Standard Setting of Unattributed Unsupervised HNE

Model	Node classification (Macro-F1/Micro-F1)				Link prediction (AUC/MRR)			
	DBLP	Yelp	Freebase	PubMed	DBLP	Yelp	Freebase	PubMed
metapath2vec	43.85/55.07	5.16/23.32	20.55/46.43	12.90/15.51	65.26/90.68	80.52/99.72	56.14/78.24	69.38/84.79
PTE	43.34/54.53	5.10/23.24	10.25/39.87	09.74/12.27	57.72/77.51	50.32/68.84	57.89/78.23	70.36/89.54
HIN2Vec	12.17/25.88	5.12/23.25	17.40/41.92	10.93/15.31	53.29/75.47	51.64/66.71	58.11/81.65	69.68/84.48
AspEm	33.07/43.85	5.40/23.82	23.26/45.42	11.19/14.44	67.20/91.46	76.10/95.18	55.80/77.70	68.31/87.43
HEER	09.72/27.72	5.03/22.92	12.96/37.51	11.73/15.29	53.00/72.76	73.72/95.92	55.78/78.31	69.06/88.42
R-GCN	09.38/13.39	5.10/23.24	06.89/38.02	10.75/12.73	50.50/73.35	72.17/97.46	50.18/74.01	63.33/81.19
HAN	07.91/16.98	5.10/23.24	06.90/38.01	09.54/12.18	50.24/73.10	N/A	51.50/74.13	65.85/85.33
MAGNN	06.74/10.35	5.10/23.24	06.89/38.02	10.30/12.60	50.10/73.26	50.03/69.81	50.12/74.18	61.11/90.01
HGT	15.17/32.05	5.07/23.12	23.06/46.51	11.24/18.72	59.98/83.13	79.00/99.66	55.68/79.46	73.00/88.05
TransE	22.76/37.18	5.05/23.03	31.83/52.04	11.40/15.16	63.53/86.29	69.13/83.66	52.84/75.80	67.95/84.69
DistMult	11.42/25.07	5.04/23.00	23.82/45.50	11.27/15.79	52.87/74.84	80.28/99.73	54.91/78.04	70.61/90.64
ComplEx	20.48/37.34	5.05/23.03	35.26/52.03	09.84/18.51	65.92/90.01	80.11/99.73	60.43/84.22	75.96/92.47
ConvE	12.42/26.42	5.09/23.02	24.57/47.61	13.00/14.49	54.03/75.31	78.55/99.70	54.29/76.11	71.81/89.82

- Proximity-preserving algorithms often perform well on both tasks under the unsupervised unattributed HNE setting, which explains why proximity-preserving is the most widely used HNE or even general network embedding framework when node attributes and labels are unavailable. Among the proximity-preserving methods, HIN2Vec and HEER show reasonable results on link prediction but perform not so well on node classification (especially on DBLP and Freebase). In fact, these two methods focus on modeling link representations in their objectives (A_M in HIN2Vec and A_l in HEER), thus are more suitable for link prediction.
- Under the unsupervised unattributed HNE setting, message-passing methods perform poorly except for HGT, especially on node classification. As we discuss before, message-passing methods are known to excel due to their integration of node attributes, link structures, and training labels. When neither of node attributes and labels are available, we use random vectors as node features and adopt a link prediction loss, which largely limits the performance of R-GCN, HAN, and MAGNN. We will focus our evaluation on the message-passing algorithms in the attributed and semi-supervised HNE settings later. On the contrary, HGT exhibits competitive results on both node classification and link prediction. This is attributed to the usage of node-type and link-type dependent parameters which maintains dedicated representations for different types of nodes and links. In addition, the heterogeneous mini-batch graph sampling algorithm designed by HGT further reduces the loss of structural information due to sampling and boosts the performance to a greater extent. Finally, the link prediction result of HAN on the Yelp dataset is not available. This is because HAN can only embed one type of nodes at a time (we embed Business in Yelp) and thus predict the link between two nodes with the same type (i.e., Business-Business). However, all links in Yelp connect distinct types of nodes (e.g., Business-Location, Business-User), and HAN cannot predict such links (thus marked as N/A).
- Relation-learning methods such as TransE and ComplEx perform better on Freebase and PubMed on both tasks, especially on link prediction. In fact, in Table 4 and Fig. 3 we can observe that both datasets (especially Freebase) have more link types. Relation-learning approaches, which are mainly designed to embed knowledge graphs (e.g., Freebase), can better capture the semantics of numerous types of direct links. From the perspective of datasets:
 - All approaches have relatively low F1 scores on Yelp and PubMed (especially Yelp) on node classification. This is because both datasets have larger numbers of classes (i.e., 16 in Yelp and 8 in PubMed) as shown in Table 4. Moreover, unlike the cases of the other datasets, a node in Yelp can have multiple labels, which makes the classification task more challenging.
 - In Fig. 4, we can observe that the degree distribution of Freebase is more skewed. Therefore, when we conduct link sampling or random walks on Freebase during representation learning, nodes with lower degrees will be sampled less frequently and their representations may not be learned accurately. This observation may explain why the link prediction metrics on Freebase are in general lower than those on Yelp and PubMed.
 - As we can see in Fig. 3, 4, 5, and 6, most studied network properties are more balanced on Freebase and PubMed (especially PubMed) across different types of nodes and links. This in general makes both the node classification and link prediction tasks harder for all algorithms, and also makes the gaps among different algorithms smaller.

5.3 Ablation Studies

To provide an in-depth performance comparison among various HNE algorithms, we further conduct controlled experiments by varying the embedding sizes and randomly removing links from the training set.

In Fig. 7, we show the micro-F1 scores for node classification and AUC scores for link prediction computed on the PubMed dataset. We omit the other results here, which can be easily computed in our provided benchmark package.




Fig. 7. Performance comparison under controlled experiments with varying emb. sizes and link removals (PubMed).




Fig. 8. Performance comparison under controlled experiments with varying label amount and attr. noise (PubMed).

As we can observe, some algorithms are more robust to varying settings while some others are more sensitive. In general, varying embedding size and link removal can significantly impact the performance of most algorithms on both tasks, and sometimes can even lead to different ordering of certain algorithms. This again emphasizes the importance of setting up standard benchmark including datasets and evaluation protocols for systematic HNE algorithm evaluation. In particular, on PubMed, larger embedding sizes like over 50 can harm the performance of most algorithms especially on node classification, probably due to overfitting with the limited labeled data. Interestingly, the random removal of links does have a negative impact on link prediction, but it does not necessarily harm node classification. This means that node classes and link structures may not always be tightly correlated, and even parts of the links already provide the necessary information useful enough for node classification.

Towards the evaluation of the message-passing HNE algorithms designed to integrate node attributes and labels into representation learning like R-GCN, HAN, MAGNN, and HGT, we also conduct controlled experiments by adding random Gaussian noises to the node attributes and masking different amounts of training labels.

In Fig. 8, we show the results on the PubMed dataset. As we can observe, the scores in most subfigures are significantly higher than the scores in Table 5, indicating the effectiveness of R-GCN, HAN, MAGNN, and HGT in integrating node attributes and labels for HNE. In particular, the incorporation of node attributes boosts the node classification results of R-GCN, HAN, and MAGNN significantly (almost tripling the F1 scores and significantly higher than all algorithms that do not use node attributes), but it offers very little help to HGT. This suggests that R-GCN, HAN, and MAGNN can effectively leverage the semantic information associated with attributes, whereas HGT relies more on the network structures and type information of nodes and links. Moreover, MAGNN achieves the highest node classification results with large margins as it successfully incorporates the attributes of intermediate nodes along

the meta-paths (which are ignored by HAN). In addition, when random noises with larger variances are added to node attributes, the performance of node classification significantly drops, while the performance of link prediction is less affected. As more training labels become available, without a surprise, the node classification results of all four algorithms increase, but surprisingly the link prediction results are almost not affected. These observations again reveal the different natures of the two tasks, where node classes are more related to node contents, whereas links should be typically inferred from structural information.

6 FUTURE

In this work, we present a comprehensive survey on various existing HNE algorithms, and provide benchmark datasets and baseline implementations to ease future research in this direction. While HNE has already demonstrated strong performance across a variety of downstream tasks, it is still in its infancy with many open challenges. To conclude this work and inspire future research, we now briefly discuss the limitation of current HNE and several specific directions potentially worth pursuing.

Beyond Homophily. As we formulate in Eq. (1), current HNE algorithms focus on the leverage of network homophily. Due to recent research on homogeneous networks that study the combination of *positional* and *structural* embedding [49], [145], it would be interesting to explore how to generalize such design principles and paradigms to HNE. Particularly, in heterogeneous networks, relative positions and structural roles of nodes can both be measured under different meta-paths or meta-graphs, which are naturally more informative and diverse. However, such considerations also introduce harder computational challenges.

Beyond Accuracy. Most, if not all, existing research on HNE has primarily focused on the accuracy towards different downstream tasks. It would be interesting to further study the *scalability and efficiency* (for large-scale networks) [34], [146], *temporal adaptability* (for dynamic evolving networks) [147], [148], *robustness* (towards adversarial attacks)

[32], [149], *interpretability* [150], *uncertainty* [151], *fairness* [152], [153] of HNE, and so on.

Beyond Node Embedding. Graph- and subgraph-level embeddings have been intensively studied on homogeneous networks to enable graph-level classification [43], [154] and unsupervised message-passing model training [36], [155], but they are hardly studied on heterogeneous networks [156]. Although existing works like HIN2Vec [63] study the embedding of meta-paths to improve the embedding of nodes, direct applications of graph- and subgraph-level embeddings in the context of heterogeneous networks largely remain nascent.

Revisiting KB Embedding. The difference between KB embedding and other types of HNE is mainly due to the numbers of node and link types. Direct application of KB embedding to heterogeneous networks fails to consider meta-paths with rich semantics, whereas directly applying HNE to KB is unrealistic due to the exponential number of meta-paths. However, it would still be interesting to study the intersection between these two groups of methods (as well as two types of data). For example, how can we combine the ideas of meta-paths on heterogeneous networks and embedding transformation on KB for HNE with more semantic-aware transformations? How can we devise truncated random walk based methods for KB embedding to include higher-order relations?

Modeling Heterogeneous Contexts. Heterogeneous networks mainly model different types of nodes and links. However, networks nowadays are often associated with rich contents, which provide contexts of the nodes, links, and subnetworks [157], [158], [159]. There have been studies exploiting text-rich heterogeneous information networks for taxonomy construction [160], [161] and text classification [162], [163]. Meanwhile, how to model heterogeneous interactions under multi-faceted contexts through the integration of multi-modal content and structure could be a challenging but rewarding research area.

Understanding the Limitation. While HNE (as well as many neural representation learning models) has demonstrated strong performance in various domains, it is worthwhile to understand its potential limits. For example, when do modern HNE algorithms work better compared with traditional network mining approaches (e.g., path counting, subgraph matching, non-neural or linear propagation)? How can we join the advantages of both worlds? Moreover, while there has been intensive research on the mathematical mechanisms behind neural networks for homogeneous network data (e.g., smoothing, low-pass filtering, invariant and equivariant transformations), by unifying existing models on HNE, this work also aims to stimulate further theoretical studies on the power and limitation of HNE.

ACKNOWLEDGMENTS

The authors would like to thank Dr. Hanghang Tong for his generous help in the polish of this work. This work was supported in part by the U.S. Army Research Laboratory under Cooperative Agreement No. W911NF-09-2-0053 and W911NF-13-1-0193, DARPA No. W911NF-17-C-0099 and FA8750-19-2-1004, National Science Foundation IIS 16-18481, IIS 17-04532, and IIS-17-41317, DTRA HDTRA11810026, and Grant 1U54 GM114838 awarded by NIGMS through funds provided by

the trans-NIH Big Data to Knowledge (BD2K) initiative. Carl Yang, Yuxin Xiao, and Yu Zhang are contribute equally to this work.

REFERENCES

- [1] Y. Seo, M. Defferrard, P. Vandergheynst, and X. Bresson, "Structured sequence modeling with graph convolutional recurrent networks," in *Proc. Int. Conf. Neural Inf. Process.*, 2018, pp. 362–373.
- [2] J. He, M. Li, H.-J. Zhang, H. Tong, and C. Zhang, "Manifold-ranking based image retrieval," in *Proc. 12th Annu. ACM Int. Conf. Multimedia*, 2004, pp. 9–16.
- [3] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, "Neural message passing for quantum chemistry," in *Proc. 34th Int. Conf. Mach. Learn.*, 2017, pp. 1263–1272.
- [4] A. Bordes, N. Usunier, A. Garcia-Duran, J. Weston, and O. Yakhnenko, "Translating embeddings for modeling multi-relational data," in *Proc. 26th Int. Conf. Neural Inf. Process. Syst.*, 2013, pp. 2787–2795.
- [5] C. Yang, L. Bai, C. Zhang, Q. Yuan, and J. Han, "Bridging collaborative filtering and semi-supervised learning: A neural approach for POI recommendation," in *Proc. 23rd ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2017, pp. 1245–1254.
- [6] Y. Xiao, A. Krishnan, and H. Sundaram, "Discovering strategic behaviors for collaborative content-production in social networks," in *Proc. Web Conf.*, 2020, pp. 2078–2088.
- [7] Y. Sun and J. Han, "Mining heterogeneous information networks: Principles and methodologies," *Synthesis Lectures Data Mining Knowl. Discov.*, vol. 3, no. 2, pp. 1–159, 2012.
- [8] C. Shi, Y. Li, J. Zhang, Y. Sun, and S. Y. Philip, "A survey of heterogeneous information network analysis," *IEEE Trans. Knowl. Data Eng.*, vol. 29, no. 1, pp. 17–37, Jan. 2017.
- [9] J.-D. Zhang and C.-Y. Chow, "GeoSoCa: Exploiting geographical, social and categorical correlations for point-of-interest recommendations," in *Proc. 38th Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 2015, pp. 443–452.
- [10] C. Yang, D. H. Hoang, T. Mikolov, and J. Han, "Place deduplication with embeddings," in *Proc. World Wide Web Conf.*, 2019, pp. 3420–3426.
- [11] C. Yang, C. Zhang, X. Chen, J. Ye, and J. Han, "Did you enjoy the ride: Understanding passenger experience via heterogeneous network embedding," in *Proc. IEEE 34th Int. Conf. Data Eng.*, 2018, pp. 1392–1403.
- [12] Y. Li and J. C. Patra, "Genome-wide inferring gene-phenotype relationship by walking on the heterogeneous network," *Bioinformatics*, vol. 26, no. 9, pp. 1219–1224, 2010.
- [13] A. P. Davis *et al.*, "The comparative toxicogenomics database: Update 2017," *Nucleic Acids Res.*, vol. 45, no. D1, pp. D972–D978, 2016.
- [14] Y. Sun, J. Han, X. Yan, P. S. Yu, and T. Wu, "PathSim: Meta path-based top-k similarity search in heterogeneous information networks," in *Proc. 37th Int. Conf. Very Large Data Bases*, 2011, pp. 992–1003.
- [15] C. Shi, X. Kong, Y. Huang, S. Y. Philip, and B. Wu, "HeteSim: A general framework for relevance measure in heterogeneous networks," *IEEE Trans. Knowl. Data Eng.*, vol. 26, no. 10, pp. 2479–2492, Oct. 2014.
- [16] C. Yang, M. Liu, F. He, X. Zhang, J. Peng, and J. Han, "Similarity modeling on heterogeneous networks via automatic path discovery," in *Proc. Joint Eur. Conf. Mach. Learn. Knowl. Discov. Databases*, 2018, pp. 37–54.
- [17] L. Dos Santos, B. Piwowarski, and P. Gallinari, "Multilabel classification on heterogeneous graphs with Gaussian embeddings," in *Proc. Joint Eur. Conf. Mach. Learn. Knowl. Discov. Databases*, 2016, pp. 606–622.
- [18] D. Eswaran, S. Günnemann, C. Faloutsos, D. Makhija, and M. Kumar, "ZooBP: Belief propagation for heterogeneous networks," in *Proc. Int. Conf. Very Large Data Bases*, 2017, pp. 625–636.
- [19] T. Chen and Y. Sun, "Task-guided and path-augmented heterogeneous network embedding for author identification," in *Proc. 10th ACM Int. Conf. Web Search Data Mining*, 2017, pp. 295–304.
- [20] R. Socher, D. Chen, C. D. Manning, and A. Ng, "Reasoning with neural tensor networks for knowledge base completion," in *Proc. 26th Int. Conf. Neural Inf. Process. Syst.*, 2013, pp. 926–934.
- [21] B. Oh, S. Seo, and K.-H. Lee, "Knowledge graph completion by context-aware convolutional learning with multi-hop neighborhoods," in *Proc. 27th ACM Int. Conf. Inf. Knowl. Manage.*, 2018, pp. 257–266.

- [22] W. Zhang *et al.*, "Iteratively learning embeddings and rules for knowledge graph reasoning," in *Proc. World Wide Web Conf.*, 2019, pp. 2366–2377.
- [23] X. Geng, H. Zhang, J. Bian, and T.-S. Chua, "Learning image and user features for recommendation in social networks," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015, pp. 4274–4282.
- [24] H. Zhao, Q. Yao, J. Li, Y. Song, and D. L. Lee, "Meta-graph based recommendation fusion over heterogeneous information networks," in *Proc. 23rd ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2017, pp. 635–644.
- [25] S. Hou, Y. Ye, Y. Song, and M. Abdulhayoglu, "HinDroid: An intelligent android malware detection system based on structured heterogeneous information network," in *Proc. 23rd ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2017, pp. 1507–1515.
- [26] P. Goyal and E. Ferrara, "Graph embedding techniques, applications, and performance: A survey," *Knowl.-Based Syst.*, vol. 151, pp. 78–94, 2018.
- [27] H. Cai, V. W. Zheng, and K. C.-C. Chang, "A comprehensive survey of graph embedding: Problems, techniques, and applications," *IEEE Trans. Knowl. Data Eng.*, vol. 30, no. 9, pp. 1616–1637, Sep. 2018.
- [28] P. Cui, X. Wang, J. Pei, and W. Zhu, "A survey on network embedding," *IEEE Trans. Knowl. Data Eng.*, vol. 31, no. 5, pp. 833–852, May 2019.
- [29] B. Perozzi, R. Al-Rfou, and S. Skiena, "DeepWalk: Online learning of social representations," in *Proc. 20th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2014, pp. 701–710.
- [30] J. Tang, M. Qu, M. Wang, M. Zhang, J. Yan, and Q. Mei, "LINE: Large-scale information network embedding," in *Proc. 24th Int. Conf. World Wide Web*, 2015, pp. 1067–1077.
- [31] A. Grover and J. Leskovec, "node2vec: Scalable feature learning for networks," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2016, pp. 855–864.
- [32] H. Wang *et al.*, "GraphGAN: Graph representation learning with generative adversarial nets," in *Proc. 32nd AAAI Conf. Artif. Intell.*, 2018.
- [33] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *Proc. Int. Conf. Learn. Representations*, 2017.
- [34] W. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 1025–1035.
- [35] J. Chen, T. Ma, and C. Xiao, "FastGCN: Fast learning with graph convolutional networks via importance sampling," in *Proc. Int. Conf. Learn. Representations*, 2018.
- [36] P. Veličković, W. Fedus, W. L. Hamilton, P. Liò, Y. Bengio, and R. D. Hjelm, "Deep graph infomax," in *Proc. Int. Conf. Learn. Representations*, 2019.
- [37] N. Zhao, H. Zhang, M. Wang, R. Hong, and T.-S. Chua, "Learning content–social influential features for influence analysis," *Int. J. Multimedia Inf. Retrieval*, vol. 5, pp. 137–149, 2016.
- [38] S. Abu-El-Haija, B. Perozzi, and R. Al-Rfou, "Learning edge representations via low-rank asymmetric projections," in *Proc. ACM Conf. Inf. Knowl. Manage.*, 2017, pp. 1787–1796.
- [39] B. Perozzi, V. Kulkarni, H. Chen, and S. Skiena, "Don't walk, skip!: Online learning of multi-scale network embeddings," in *Proc. IEEE/ACM Int. Conf. Advances Soc. Netw. Anal. Mining*, 2017, pp. 258–265.
- [40] C. Yang *et al.*, "Relation learning on social networks with multimodal graph edge variational autoencoders," in *Proc. 13th Int. Conf. Web Search Data Mining*, 2020, pp. 699–707.
- [41] M. Niepert, M. Ahmed, and K. Kutzkov, "Learning convolutional neural networks for graphs," in *Proc. 33rd Int. Conf. Mach. Learn.*, 2016, pp. 2014–2023.
- [42] C. Yang, M. Liu, V. W. Zheng, and J. Han, "Node, motif and subgraph: Leveraging network functional blocks through structural convolution," in *Proc. IEEE/ACM Int. Conf. Advances Soc. Netw. Anal. Mining*, 2018, pp. 47–52.
- [43] Z. Ying, J. You, C. Morris, X. Ren, W. Hamilton, and J. Leskovec, "Hierarchical graph representation learning with differentiable pooling," in *Proc. 32nd Int. Conf. Neural Inf. Process. Syst.*, 2018, pp. 4805–4815.
- [44] C. Morris *et al.*, "Weisfeiler and leman go neural: Higher-order graph neural networks," in *Proc. AAAI Conf. Artif. Intell.*, 2019, pp. 4602–4609.
- [45] S. Cao, W. Lu, and Q. Xu, "GraRep: Learning graph representations with global structural information," in *Proc. 24th ACM Int. Conf. Inf. Knowl. Manage.*, 2015, pp. 891–900.
- [46] D. Wang, P. Cui, and W. Zhu, "Structural deep network embedding," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2016, pp. 1225–1234.
- [47] Z. Zhang, P. Cui, X. Wang, J. Pei, X. Yao, and W. Zhu, "Arbitrary-order proximity preserved network embedding," in *Proc. 24th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2018, pp. 2778–2786.
- [48] J. Huang, X. Liu, and Y. Song, "Hyper-path-based representation learning for hyper-networks," in *Proc. 28th ACM Int. Conf. Inf. Knowl. Manage.*, 2019, pp. 449–458.
- [49] L. F. Ribeiro, P. H. Saverese, and D. R. Figueiredo, "struc2vec: Learning node representations from structural identity," in *Proc. 23rd ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2017, pp. 385–394.
- [50] M. Zhang and Y. Chen, "Weisfeiler-lehman neural machine for link prediction," in *Proc. 23rd ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2017, pp. 575–583.
- [51] T. Lyu, Y. Zhang, and Y. Zhang, "Enhancing the network embedding quality with structural similarity," in *Proc. ACM Conf. Inf. Knowl. Manage.*, 2017, pp. 147–156.
- [52] C. Donnat, M. Zitnik, D. Hallac, and J. Leskovec, "Learning structural node embeddings via diffusion wavelets," in *Proc. 24th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2018, pp. 1320–1329.
- [53] B. Schölkopf and A. J. Smola, *Learning With Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. Cambridge, MA, USA: MIT Press, 2001.
- [54] A. Liaw *et al.*, "Classification and regression by randomforest," *R News*, vol. 2, no. 3, pp. 18–22, 2002.
- [55] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2016, pp. 785–794.
- [56] S. Chang, W. Han, J. Tang, G.-J. Qi, C. C. Aggarwal, and T. S. Huang, "Heterogeneous network embedding via deep architectures," in *Proc. 21th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2015, pp. 119–128.
- [57] F. Wu *et al.*, "Learning of multimodal representations with random walks on the click graph," *IEEE Trans. Image Process.*, vol. 25, no. 2, pp. 630–642, Feb. 2016.
- [58] Y. Zhao, Z. Liu, and M. Sun, "Representation learning for measuring entity relatedness with rich information," in *Proc. 24th Int. Conf. Artif. Intell.*, 2015, pp. 1412–1418.
- [59] Y. Dong, N. V. Chawla, and A. Swami, "metapath2vec: Scalable representation learning for heterogeneous networks," in *Proc. 23rd ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2017, pp. 135–144.
- [60] Y. Shi, H. Gui, Q. Zhu, L. Kaplan, and J. Han, "AspEm: Embedding learning by aspects in heterogeneous information networks," in *Proc. SIAM Int. Conf. Data Mining*, 2018, pp. 144–152.
- [61] Y. Shi, Q. Zhu, F. Guo, C. Zhang, and J. Han, "Easing embedding learning by comprehensive transcription of heterogeneous information networks," in *Proc. 24th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2018, pp. 2190–2199.
- [62] H. Gui, J. Liu, F. Tao, M. Jiang, B. Norick, and J. Han, "Large-scale embedding learning in heterogeneous event data," in *Proc. IEEE 16th Int. Conf. Data Mining*, 2016, pp. 907–912.
- [63] T.-Y. Fu, W.-C. Lee, and Z. Lei, "HIN2Vec: Explore meta-paths in heterogeneous information networks for representation learning," in *Proc. ACM Conf. Inf. Knowl. Manage.*, 2017, pp. 1797–1806.
- [64] C. Yang, Y. Feng, P. Li, Y. Shi, and J. Han, "Meta-graph based HIN spectral embedding: Methods, analyses, and insights," in *Proc. IEEE Int. Conf. Data Mining*, 2018, pp. 657–666.
- [65] R. Hussein, D. Yang, and P. Cudr é-Mauroux, "Are meta-paths necessary? Revisiting heterogeneous graph embeddings," in *Proc. 27th ACM Int. Conf. Inf. Knowl. Manage.*, 2018, pp. 437–446.
- [66] Y. Zhang, Y. Xiong, X. Kong, S. Li, J. Mi, and Y. Zhu, "Deep collective classification in heterogeneous information networks," in *Proc. World Wide Web Conf.*, 2018, pp. 399–408.
- [67] M. Schlichtkrull, T. N. Kipf, P. Bloem, R. Van Den Berg, I. Titov, and M. Welling, "Modeling relational data with graph convolutional networks," in *Proc. Eur. Semantic Web Conf.*, 2018, pp. 593–607.

- [68] M. Qu, J. Tang, and J. Han, "Curriculum learning for heterogeneous star network embedding via deep reinforcement learning," in *Proc. 11th ACM Int. Conf. Web Search Data Mining*, 2018, pp. 468–476.
- [69] C. Zhang, A. Swami, and N. V. Chawla, "SHNE: Representation learning for semantic-associated heterogeneous networks," in *Proc. 12th ACM Int. Conf. Web Search Data Mining*, 2019, pp. 690–698.
- [70] Y. Cen, X. Zou, J. Zhang, H. Yang, J. Zhou, and J. Tang, "Representation learning for attributed multiplex heterogeneous network," in *Proc. 25th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2019, pp. 1358–1368.
- [71] C. Zhang, D. Song, C. Huang, A. Swami, and N. V. Chawla, "Heterogeneous graph neural network," in *Proc. 25th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2019, pp. 793–803.
- [72] B. Hu, Y. Fang, and C. Shi, "Adversarial learning on heterogeneous information networks," in *Proc. 25th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2019, pp. 120–129.
- [73] X. Wang, Y. Zhang, and C. Shi, "Hyperbolic heterogeneous information network embedding," in *Proc. 33rd AAAI Conf. Artif. Intell.*, 2019, pp. 5337–5344.
- [74] C. Yang, J. Zhang, and J. Han, "Neural embedding propagation on heterogeneous networks," in *Proc. IEEE Int. Conf. Data Mining*, 2019, pp. 698–707.
- [75] X. Wang et al., "Heterogeneous graph attention network," in *Proc. World Wide Web Conf.*, 2019, pp. 2022–2032.
- [76] Y. Lu, C. Shi, L. Hu, and Z. Liu, "Relation structure-aware heterogeneous information network embedding," in *Proc. AAAI Conf. Artif. Intell.*, 2019, pp. 4456–4463.
- [77] C. Shi, B. Hu, W. X. Zhao, and S. Y. Philip, "Heterogeneous information network embedding for recommendation," *IEEE Trans. Knowl. Data Eng.*, vol. 31, no. 2, pp. 357–370, Feb. 2019.
- [78] Y. Cao, X. Wang, X. He, Z. Hu, and T.-S. Chua, "Unifying knowledge graph learning and recommendation: Towards a better understanding of user preferences," in *Proc. World Wide Web Conf.*, 2019, pp. 151–161.
- [79] Q. Wang, Z. Mao, B. Wang, and L. Guo, "Knowledge graph embedding: A survey of approaches and applications," *IEEE Trans. Knowl. Data Eng.*, vol. 29, no. 12, pp. 2724–2743, Dec. 2017.
- [80] Y. Lin, Z. Liu, M. Sun, Y. Liu, and X. Zhu, "Learning entity and relation embeddings for knowledge graph completion," in *Proc. 29th AAAI Conf. Artif. Intell.*, 2015, pp. 2181–2187.
- [81] B. Yang, W.-T. Yih, X. He, J. Gao, and L. Deng, "Embedding entities and relations for learning and inference in knowledge bases," in *Proc. Int. Conf. Learn. Representations*, 2015.
- [82] T. Dettmers, P. Minervini, P. Stenetorp, and S. Riedel, "Convolutional 2D knowledge graph embeddings," in *Proc. 32nd AAAI Conf. Artif. Intell.*, 2018.
- [83] K. Wang, Y. Liu, X. Xu, and D. Lin, "Knowledge graph embedding with entity neighbors and deep memory network," 2018, *arXiv: 1808.03752*.
- [84] C. Shang, Y. Tang, J. Huang, J. Bi, X. He, and B. Zhou, "End-to-end structure-aware convolutional networks for knowledge base completion," in *Proc. AAAI Conf. Artif. Intell.*, 2019, pp. 3060–3067.
- [85] Z. Hu, Y. Dong, K. Wang, and Y. Sun, "Heterogeneous graph transformer," in *Proc. Web Conf.*, 2020, pp. 2704–2710.
- [86] Y. Ren, B. Liu, C. Huang, P. Dai, L. Bo, and J. Zhang, "Heterogeneous deep graph infomax," 2019, *arXiv: 1911.08538*.
- [87] Y. Dong, Z. Hu, K. Wang, Y. Sun, and J. Tang, "Heterogeneous network representation learning," in *Proc. 29th Int. Joint Conf. Artif. Intell.*, 2020, pp. 4861–4867.
- [88] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, "A comprehensive survey on graph neural networks," *IEEE Trans. Neural Netw. Learn. Syst.*, to be published, doi: [10.1109/TNNLS.2020.2978386](https://doi.org/10.1109/TNNLS.2020.2978386).
- [89] V. P. Dwivedi, C. K. Joshi, T. Laurent, Y. Bengio, and X. Bresson, "Benchmarking graph neural networks," 2020, *arXiv: 2003.00982*.
- [90] X. Yue et al., "Graph embedding on biomedical networks: Methods, applications and evaluations," *Bioinformatics*, vol. 36, no. 4, pp. 1241–1251, 2020.
- [91] H. Jiang, Y. Song, C. Wang, M. Zhang, and Y. Sun, "Semi-supervised learning over heterogeneous information networks by ensemble of meta-graph guided random walks," in *Proc. 26th Int. Joint Conf. Artif. Intell.*, 2017, pp. 1944–1950.
- [92] M. McPherson, L. Smith-Lovin, and J. M. Cook, "Birds of a feather: Homophily in social networks," *Annu. Rev. Sociol.*, vol. 27, no. 1, pp. 415–444, 2001.
- [93] M. Belkin and P. Niyogi, "Laplacian eigenmaps and spectral techniques for embedding and clustering," in *Proc. 14th Int. Conf. Neural Inf. Process. Syst.*, 2002, pp. 585–591.
- [94] D. Zhou, O. Bousquet, T. N. Lal, J. Weston, and B. Schölkopf, "Learning with local and global consistency," in *Proc. 16th Int. Conf. Neural Inf. Process. Syst.*, 2004, pp. 321–328.
- [95] X. Zhu et al., "Semi-supervised learning using gaussian fields and harmonic functions," in *Proc. 20th Int. Conf. Mach. Learn.*, 2003, pp. 912–919.
- [96] J. B. Tenenbaum, V. De Silva, and J. C. Langford, "A global geometric framework for nonlinear dimensionality reduction," *Science*, vol. 290, no. 5500, pp. 2319–2323, 2000.
- [97] S. T. Roweis and L. K. Saul, "Nonlinear dimensionality reduction by locally linear embedding," *Science*, vol. 290, no. 5500, pp. 2323–2326, 2000.
- [98] 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 *Proc. 11th ACM Int. Conf. Web Search Data Mining*, 2018, pp. 459–467.
- [99] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *Proc. 26th Int. Conf. Neural Inf. Process. Syst.*, 2013, pp. 3111–3119.
- [100] H. Zhang, L. Qiu, L. Yi, and Y. Song, "Scalable multiplex network embedding," in *Proc. 27th Int. Joint Conf. Artif. Intell.*, 2018, pp. 3082–3088.
- [101] Y. He, Y. Song, J. Li, C. Ji, J. Peng, and H. Peng, "HeteSpaceyWalk: A heterogeneous spacey random walk for heterogeneous information network embedding," in *Proc. 28th ACM Int. Conf. Inf. Knowl. Manage.*, 2019, pp. 639–648.
- [102] C. Park, D. Kim, Q. Zhu, J. Han, and H. Yu, "Task-guided pair embedding in heterogeneous network," in *Proc. 28th ACM Int. Conf. Inf. Knowl. Manage.*, 2019, pp. 489–498.
- [103] J. Tang, M. Qu, and Q. Mei, "PTE: Predictive text embedding through large-scale heterogeneous text networks," in *Proc. 21th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2015, pp. 1165–1174.
- [104] M. Qu, J. Tang, J. Shang, X. Ren, M. Zhang, and J. Han, "An attention-based collaboration framework for multi-view network representation learning," in *Proc. ACM Conf. Inf. Knowl. Manage.*, 2017, pp. 1767–1776.
- [105] Z. Jiang, Z. Gao, J. Lan, H. Yang, Y. Lu, and X. Liu, "Task-oriented genetic activation for large-scale complex heterogeneous graph embedding," in *Proc. Web Conf.*, 2020, pp. 1581–1591.
- [106] W. Zhang, Y. Fang, Z. Liu, M. Wu, and X. Zhang, "mg2vec: Learning relationship-preserving heterogeneous graph representations via metagraph embedding," *IEEE Trans. Knowl. Data Eng.*, to be published, doi: [10.1109/TKDE.2020.2992500](https://doi.org/10.1109/TKDE.2020.2992500).
- [107] H. Wang, F. Zhang, M. Hou, X. Xie, M. Guo, and Q. Liu, "SHINE: Signed heterogeneous information network embedding for sentiment link prediction," in *Proc. 11th ACM Int. Conf. Web Search Data Mining*, 2018, pp. 592–600.
- [108] A. M. Fard, E. Bagheri, and K. Wang, "Relationship prediction in dynamic heterogeneous information networks," in *Proc. Eur. Conf. Inf. Retrieval*, 2019, pp. 19–34.
- [109] S. Vashishth, S. Sanyal, V. Nitin, and P. Talukdar, "Composition-based multi-relational graph convolutional networks," in *Proc. Int. Conf. Learn. Representations*, 2020.
- [110] V. W. Zheng et al., "Heterogeneous embedding propagation for large-scale e-commerce user alignment," in *Proc. IEEE Int. Conf. Data Mining*, 2018, pp. 1434–1439.
- [111] K. Zhao et al., "Deep adversarial completion for sparse heterogeneous information network embedding," in *Proc. Web Conf.*, 2020, pp. 508–518.
- [112] X. Fu, J. Zhang, Z. Meng, and I. King, "Magnn: Metapath aggregated graph neural network for heterogeneous graph embedding," in *Proc. Web Conf.*, 2020, pp. 2331–2341.
- [113] H. Tong, C. Faloutsos, and J.-Y. Pan, "Fast random walk with restart and its applications," in *Proc. 6th Int. Conf. Data Mining*, 2006, pp. 613–622.
- [114] Z. Sun, Z. Deng, J. Nie, and J. Tang, "ROTATE: Knowledge graph embedding by relational rotation in complex space," in *Proc. Int. Conf. Learn. Representations*, 2019.
- [115] A. Vaswani et al., "Attention is all you need," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 6000–6010.

- [116] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics*, 2019, pp. 4171–4186.
- [117] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [118] Z. Wang, J. Zhang, J. Feng, and Z. Chen, "Knowledge graph embedding by translating on hyperplanes," in *Proc. 28th AAAI Conf. Artif. Intell.*, 2014, pp. 1112–1119.
- [119] M. Nickel, V. Tresp, and H.-P. Kriegel, "A three-way model for collective learning on multi-relational data," in *Proc. 28th Int. Conf. Mach. Learn.*, 2011, pp. 809–816.
- [120] M. Nickel, L. Rosasco, and T. Poggio, "Holographic embeddings of knowledge graphs," in *Proc. 30th AAAI Conf. Artif. Intell.*, 2016, pp. 1955–1961.
- [121] T. Trouillon, J. Welbl, S. Riedel, É. Gaussier, and G. Bouchard, "Complex embeddings for simple link prediction," in *Proc. 33rd Int. Conf. Mach. Learn.*, 2016, pp. 2071–2080.
- [122] S. M. Kazemi and D. Poole, "Simple embedding for link prediction in knowledge graphs," in *Proc. 32nd Int. Conf. Neural Inf. Process. Syst.*, 2018, pp. 4289–4300.
- [123] I. Balazevic, C. Allen, and T. Hospedales, "TuckER: Tensor factorization for knowledge graph completion," in *Proc. Conf. Empir. Methods Natural Lang. Process.*, 2019, pp. 5185–5194.
- [124] Y. Xiao, Z. Zhang, C. Yang, and C. Zhai, "Non-local attention learning on large heterogeneous information networks," in *Proc. IEEE Int. Conf. Big Data*, 2019, pp. 978–987.
- [125] Z. Zhu, X. Fan, X. Chu, and J. Bi, "HGCN: A heterogeneous graph convolutional network-based deep learning model toward collective classification," in *Proc. 26th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2020, pp. 1161–1171.
- [126] H. Hong *et al.*, "HetETA: Heterogeneous information network embedding for estimating time of arrival," in *Proc. 26th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2020, pp. 2444–2454.
- [127] C. Wang *et al.*, "RelSim: Relation similarity search in schema-rich heterogeneous information networks," in *Proc. SIAM Int. Conf. Data Mining*, 2016, pp. 621–629.
- [128] S. Ji, S. Pan, E. Cambria, P. Marttinen, and P. S. Yu, "A survey on knowledge graphs: Representation, acquisition and applications," 2020, *arXiv:2002.00388*.
- [129] J. Qiu, H. Ma, Y. Dong, K. Wang, and J. Tang, "Revisiting knowledge base embedding as tensor decomposition," 2018. [Online]. Available: <https://openreview.net/pdf?id=S1sRrN-CW>
- [130] J. Shang, J. Liu, M. Jiang, X. Ren, C. R. Voss, and J. Han, "Automated phrase mining from massive text corpora," *IEEE Trans. Knowl. Data Eng.*, vol. 30, no. 10, pp. 1825–1837, Oct. 2018.
- [131] X. Wang, Y. Zhang, Q. Li, X. Ren, J. Shang, and J. Han, "Distantly supervised biomedical named entity recognition with dictionary expansion," in *Proc. IEEE Int. Conf. Bioinf. Biomedicine*, 2019, pp. 496–503.
- [132] Q. Li, X. Wang, Y. Zhang, F. Ling, C. H. Wu, and J. Han, "Pattern discovery for wide-window open information extraction in biomedical literature," in *Proc. IEEE Int. Conf. Bioinf. Biomedicine*, 2018, pp. 420–427.
- [133] F. M. Suchanek, G. Kasneci, and G. Weikum, "Yago: A core of semantic knowledge," in *Proc. 16th Int. Conf. World Wide Web*, 2007, pp. 697–706.
- [134] D. Vrandečić and M. Krötzsch, "Wikidata: A free collaborative knowledgebase," *Commun. ACM*, vol. 57, no. 10, pp. 78–85, 2014.
- [135] F. Zhang *et al.*, "OAG: Toward linking large-scale heterogeneous entity graphs," in *Proc. 25th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2019, pp. 2585–2595.
- [136] Z. Liu *et al.*, "Interactive paths embedding for semantic proximity search on heterogeneous graphs," in *Proc. 24th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2018, pp. 1860–1869.
- [137] Y. Zhang *et al.*, "Your style your identity: Leveraging writing and photography styles for drug trafficker identification in darknet markets over attributed heterogeneous information network," in *Proc. World Wide Web Conf.*, 2019, pp. 3448–3454.
- [138] S. Fan *et al.*, "Metapath-guided heterogeneous graph neural network for intent recommendation," in *Proc. 25th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2019, pp. 2478–2486.
- [139] X. Niu *et al.*, "A dual heterogeneous graph attention network to improve long-tail performance for shop search in e-commerce," in *Proc. 26th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2020, pp. 3405–3415.
- [140] W. Luo *et al.*, "Dynamic heterogeneous graph neural network for real-time event prediction," in *Proc. 26th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2020, pp. 3213–3223.
- [141] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin, "LIBLINEAR: A library for large linear classification," in *Proc. Int. Conf. Mach. Learn.*, 2008.
- [142] M. Wang *et al.*, "Deep graph library: Towards efficient and scalable deep learning on graphs," in *Proc. Int. Conf. Learn. Representations Workshops*, 2019.
- [143] M. Fey and J. E. Lenssen, "Fast graph representation learning with PyTorch Geometric," in *Proc. ICLR Workshop Representation Learn. Graphs Manifolds*, 2019.
- [144] X. Han *et al.*, "OpenKE: An open toolkit for knowledge embedding," in *Proc. Conf. Empir. Methods Natural Lang. Process.*, 2018, pp. 139–144.
- [145] J. You, R. Ying, and J. Leskovec, "Position-aware graph neural networks," in *Proc. 36th Int. Conf. Mach. Learn.*, 2019, pp. 7134–7143.
- [146] W. Hu *et al.*, "Strategies for pre-training graph neural networks," in *Proc. Int. Conf. Learn. Representations*, 2020.
- [147] L. Zhou, Y. Yang, X. Ren, F. Wu, and Y. Zhuang, "Dynamic network embedding by modeling triadic closure process," in *Proc. 32nd AAAI Conf. Artif. Intell.*, 2018.
- [148] L. Du, Y. Wang, G. Song, Z. Lu, and J. Wang, "Dynamic network embedding: An extended approach for skip-gram based network embedding," in *Proc. 27th Int. Joint Conf. Artif. Intell.*, 2018, pp. 2086–2092.
- [149] D. Zügner, A. Akbarnejad, and S. Günnemann, "Adversarial attacks on neural networks for graph data," in *Proc. 24th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2018, pp. 2847–2856.
- [150] N. Liu, X. Huang, J. Li, and X. Hu, "On interpretation of network embedding via taxonomy induction," in *Proc. 24th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2018, pp. 1812–1820.
- [151] X. Chen, M. Chen, W. Shi, Y. Sun, and C. Zaniolo, "Embedding uncertain knowledge graphs," in *Proc. AAAI Conf. Artif. Intell.*, 2019, pp. 3363–3370.
- [152] T. A. Rahman, B. Surma, M. Backes, and Y. Zhang, "Fairwalk: Towards fair graph embedding," in *Proc. 28th Int. Joint Conf. Artif. Intell.*, 2019, pp. 3289–3295.
- [153] A. J. Bose and W. L. Hamilton, "Compositional fairness constraints for graph embeddings," in *Proc. 36th Int. Conf. Mach. Learn.*, 2019, pp. 715–724.
- [154] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, "How powerful are graph neural networks?," in *Proc. Int. Conf. Learn. Representations*, 2018.
- [155] F.-Y. Sun, J. Hoffman, V. Verma, and J. Tang, "InfoGraph: Unsupervised and semi-supervised graph-level representation learning via mutual information maximization," in *Proc. Int. Conf. Mach. Learn.*, 2019.
- [156] C. Park, J. Han, and H. Yu, "Deep multiplex graph infomax: Attentive multiplex network embedding using global information," in *Proc. AAAI Conf. Artif. Intell.*, 2020.
- [157] H. Wang *et al.*, "MCNE: An end-to-end framework for learning multiple conditional network representations of social network," in *Proc. 25th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2019, pp. 1064–1072.
- [158] C. Yang, A. Pal, A. Zhai, N. Pancho, C. Rosenberg, and J. Leskovec, "MultiSage: Empowering GCN with contextualized multi-embeddings on web-scale multipartite networks," in *Proc. 26th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2020, pp. 2434–2443.
- [159] C. Yang *et al.*, "CubeNet: Multi-facet hierarchical heterogeneous network construction, analysis, and mining," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2019.
- [160] Y. Shi *et al.*, "Discovering hypernymy in text-rich heterogeneous information network by exploiting context granularity," in *Proc. 28th ACM Int. Conf. Inf. Knowl. Manage.*, 2019, pp. 599–608.
- [161] J. Shang, X. Zhang, L. Liu, S. Li, and J. Han, "NetTaxo: Automated topic taxonomy construction from text-rich network," in *Proc. Web Conf.*, 2020, pp. 1908–1919.
- [162] Y. Zhang *et al.*, "HiGitClass: Keyword-driven hierarchical classification of GitHub repositories," in *Proc. IEEE Int. Conf. Data Mining*, 2019, pp. 876–885.
- [163] Y. Zhang, Y. Meng, J. Huang, F. F. Xu, X. Wang, and J. Han, "Minimally supervised categorization of text with metadata," in *Proc. 43rd Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 2020, pp. 1231–1240.



Carl Yang received the BEng degree in computer science and engineering from Zhejiang University, China, in 2014, and the PhD degree under Jiawei Han in computer science from the University of Illinois, Urbana-Champaign, Champaign, Illinois. In his research, he develops data-driven techniques and neural architectures for learning with context-rich network data. His research results have been published in top conferences like KDD, NeurIPS, WWW, ICDE, SIGIR, ICML. He also received the Dissertation Completion Fellowship of UIUC, in 2020, Best Paper Award of ICDM 2020, and has started as an assistant professor in Emory University, Atlanta, Georgia after PhD graduation.



Yuxin Xiao received the BS degree in computer science and BS degree in statistics and mathematics from the University of Illinois at Urbana-Champaign, Champaign, Illinois, in 2020. He is currently working toward the master's degree at Carnegie Mellon University, Pittsburgh, Pennsylvania. His research focuses on machine learning on graph data and he has published first-authored papers in WWW and IEEE BigData. He also received the CRA Outstanding Undergraduate Researcher Award and C.W. Gear Outstanding Undergraduate Award at UIUC, in 2020.



Yu Zhang received the BS degree from Peking University, China, in 2017, and the MS degree from the University of Illinois at Urbana-Champaign, Champaign, Illinois, in 2019. He is currently working toward the PhD degree in computer science at the University of Illinois at Urbana-Champaign, Champaign, Illinois, advised by Prof. Jiawei Han. His research interests include text-rich network mining, text classification, and their applications to bioinformatics. He has published first-authored papers in SIGIR, WSDM, ICDM, and CIKM. He also received WWW Best Poster Award Honorable Mention, in 2018.



Yizhou Sun (Member, IEEE) received the PhD degree in computer science from the University of Illinois at Urbana-Champaign, Champaign, Illinois, in 2012. She is currently an associate professor at the Department of Computer Science, University of California, Los Angeles, Los Angeles, California. Her principal research interest include mining graphs/networks, and more generally in data mining, machine learning, and network science, with a focus on modeling novel problems and proposing scalable algorithms for large-scale, real-world applications. She is a pioneer researcher in mining heterogeneous information network, with a recent focus on deep learning on graphs/networks. She has more than 100 publications in books, journals, and major conferences. Tutorials of her research have been given in many premier conferences. She received 2012 ACM SIGKDD Best Student Paper Award, 2013 ACM SIGKDD Doctoral Dissertation Award, 2013 Yahoo ACE (Academic Career Enhancement) Award, 2015 NSF CAREER Award, 2016 CS@ILLINOIS Distinguished Educator Award, 2018 Amazon Research Award, and 2019 Okawa Foundation Research Grant.



Jiawei Han (Fellow, IEEE) is Michael Aiken chair professor with the Department of Computer Science, University of Illinois at Urbana-Champaign, Champaign, Illinois. He has been researching into data mining, information network analysis, and database systems, and their various applications, with more than 600 publications. He served as the founding editor-in-chief of the *ACM Transactions on Knowledge Discovery from Data (TKDD)* (2007-2012). He has received ACM SIGKDD Innovation Award (2004), IEEE Computer Society Technical Achievement Award (2005), IEEE Computer Society W. Wallace McDowell Award (2009), Daniel C. Drucker Eminent Faculty Award at UIUC (2011), and Japan's Funai Achievement Award (2018). He is a fellow of ACM and served as co-director of KnowEnG, a Center of Excellence in Big Data Computing (2014-2019), funded by NIH Big Data to Knowledge (BD2K) Initiative and as the director of Information Network Academic Research Center (INARC) (2009-2016) supported by the Network Science-Collaborative Technology Alliance (NS-CTA) program of U.S. Army Research Lab. His co-authored textbook "Data Mining: Concepts and Techniques" (Morgan Kaufmann) has been adopted popularly as a textbook worldwide.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.