



# Professional Network Matters: Connections Empower Person-Job Fit

Hao Chen\*  
haochen@mail.bnu.edu.cn  
Beijing Normal University  
Beijing, China

Qiang Fu  
qifu@microsoft.com  
Microsoft Research  
Beijing, China

Yanbin Kang  
ybkang@linkedin.com  
LinkedIn Corporation  
Beijing, China

Lun Du†  
lun.du@microsoft.com  
Microsoft Research  
Beijing, China

Xu Chen  
xu.chen@microsoft.com  
Microsoft Research  
Beijing, China

Guangming Lu  
glu@linkedin.com  
LinkedIn Corporation  
Beijing, China

Yuxuan Lu\*  
lu.yuxuan@northeastern.edu  
Northeastern University  
Boston, USA

Shi Han  
shihan@microsoft.com  
Microsoft Research  
Beijing, China

Zi Li  
zili@linkedin.com  
LinkedIn Corporation  
Beijing, China

## ABSTRACT

Online recruitment platforms typically employ Person-Job Fit models in the core service that automatically match suitable job seekers with appropriate job positions. While existing works leverage historical or contextual information, they often disregard a crucial aspect: job seekers' social relationships in professional networks. This paper emphasizes the importance of incorporating professional networks into the Person-Job Fit model. Our innovative approach consists of two stages: (1) defining a Workplace Heterogeneous Information Network (WHIN) to capture heterogeneous knowledge, including professional connections and pre-training representations of various entities using a heterogeneous graph neural network; (2) designing a Contextual Social Attention Graph Neural Network (CSAGNN) that supplements users' missing information with professional connections' contextual information. We introduce a job-specific attention mechanism in CSAGNN to handle noisy professional networks, leveraging pre-trained entity representations from WHIN. We demonstrate the effectiveness of our approach through experimental evaluations conducted across three real-world recruitment datasets from LinkedIn, showing superior performance compared to baseline models.

## CCS CONCEPTS

- **Computing methodologies** → **Machine learning approaches**;
- **Information systems** → **Recommender systems**.

\*This work was done when they were interns at LinkedIn.

†Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

WSDM '24, March 4–8, 2024, Merida, Mexico

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0371-3/24/03...\$15.00

<https://doi.org/10.1145/3616855.3635852>

## KEYWORDS

Person-Job Fit, Heterogeneous Information Network, Graph Neural Network

### ACM Reference Format:

Hao Chen, Lun Du, Yuxuan Lu, Qiang Fu, Xu Chen, Shi Han, Yanbin Kang, Guangming Lu, and Zi Li. 2024. Professional Network Matters: Connections Empower Person-Job Fit. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining (WSDM '24)*, March 4–8, 2024, Merida, Mexico. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3616855.3635852>

## 1 INTRODUCTION

With the rapid development of the Internet, online recruitment platforms (e.g., LinkedIn<sup>1</sup>, Indeed<sup>2</sup>, and ZipRecruiter<sup>3</sup>) are becoming essential for recruiting and job seeking. A considerable number of talent profiles and job descriptions are posted on these platforms. Taking LinkedIn as an example, more than 900 million members have registered, and 90 jobs were posted every second by the first quarter of 2023<sup>4</sup>. Considering such a large number of options, Person-Job Fit (PJF) [26] has become a critical research topic for improving the efficiency of recruitment and job seeking. Person-Job Fit aims to automatically link the right talents to the right job positions according to talent competencies and job requirements.

Previous works on Person-Job Fit mainly focus on leveraging two types of information, namely, (1) historical job application information and (2) textual information in profiles and job descriptions. Collaborative filtering-based methods [3, 30] are applied to capture co-apply relations between job seekers and co-applied relations between job positions in the historical application information. Manually-engineered textual features and deep language models have been widely adopted to leverage the textual information [4, 25, 28, 36].

<sup>1</sup><https://www.linkedin.com>

<sup>2</sup><https://www.indeed.com>

<sup>3</sup><https://www.ziprecruiter.com>

<sup>4</sup><https://news.linkedin.com/about-us#Statistics>

However, workplace social connections among members, commonly called professional networks, have been overlooked as a pivotal source of information. A survey reported on the Official LinkedIn Blog<sup>5</sup> showed that out of more than 15,000 LinkedIn members, 80% believed that professional networking could help find new job opportunities, an even more striking 70% gained job opportunities directly through their connections. As an important source for achieving employment relationships, professional networks have great potential for helping with Person-Job Fit. Incorporating professional networks into the models can offer two main benefits: (1) Professional networks directly improve Person-Job Fit by bridging the gap between job seekers and potential job opportunities. A Person-Job Fit model enriched with professional network information can suggest job opportunities to users based on their professional connections. (2) Professional networks can also help complete job seekers' profiles. Often job seekers' online profiles lack comprehensive details, yet one's professional experience and skills can be, to some extent, discernible through their professional connections.


While professional networks offer advantages, they often contain a lot of noise. This noise includes connections that aren't relevant and information from relevant connections that don't help improve Person-Job Fit. To illustrate, irrelevant connections might involve job seekers' former classmates or recruiters from unrelated industries. Even when the connections are relevant to the job seeker, they could still have information unrelated to the job seeker's aspirations for the role. To tackle these noises, a promising strategy involves using heterogeneous knowledge, such as job seekers' skills, work experience, and educational background, to gauge professional connections' relevance and extract pertinent information to elevate Person-Job Fit.

In this paper, we propose a graph neural network-based framework that utilizes heterogeneous knowledge to integrate professional networks into Person-Job Fit. We address the challenge of social noise in professional networks by designing a job-specific attention mechanism. Initially, we define a Workplace Heterogeneous Information Network (WHIN) that captures heterogeneous knowledge, including professional connections. We employ a heterogeneous graph pre-training technique to learn the representations of various entities in the WHIN. Subsequently, we introduce the Contextual Social Attention Graph Neural Network (CSAGNN), designed to supplement users' lacking information with contextual information from their professional connections. To tackle the social noise in the workplace social network, we infuse a job-specific attention mechanism into CSAGNN, capitalizing on the pre-trained entity representations from WHIN.

The main contributions of this study are as follows:

- To our knowledge, we are the first to define a heterogeneous information network that incorporates heterogeneous knowledge in the Person-Job Fit scenario.
- We systematically utilize professional networks in a two-stage approach, WHIN pre-training and CSAGNN, to address the Person-Job Fit task.

<sup>5</sup><https://blog.linkedin.com/2017/june/22/the-best-way-to-land-your-next-job-opportunity>



**Figure 1: The metagraph of Workplace Heterogeneous Information Network. It encompasses not only members (M) and jobs (J), which are crucial for Person-Job Fit, but also entities such as skills (S), companies (C), and schools (H).**

- We present a novel Contextual Social Attention Graph Neural Network (CSAGNN) specifically designed to handle noisy professional networks, effectively mitigating the impact of irrelevant information while focusing on related professional connections and contexts for Person-Job Fit.
- We evaluate our approach on three real-world datasets across diverse industries. Experimental results show that our model outperforms baseline models.

## 2 RELATED WORKS

The related work of our study can be grouped into two main categories, namely, *Person-Job Fit* and *Heterogeneous Information Network-based Recommendation*.


### 2.1 Person-Job Fit

As a core function of the online recruitment platform, Person-Job Fit [26] has received widespread attention.

Mainstream works view Person-Job Fit as a text-matching problem between member profiles and job descriptions to fully use the rich contextual knowledge. For example, JLMIA [31] is a latent variable model to model job descriptions and member profiles jointly. PJFNN [41] encodes member profiles and job descriptions by hierarchical CNN. BPJFNN [28] leverage BiLSTM to get the semantic representation of each word in member profiles and job descriptions. APJFNN [28] automatically weighs abilities mentioned in textual information based on historical recruitment results. A transferable deep global match network [4] is proposed to solve the domain adaptation problem in three levels for Person-Job Fit. JRMPM [36] proposes a matching network with preference modeled to explore the latent preference given the history of the matching process.

Some works take structure knowledge into consideration. For example, MV-CoN [3] adopts a co-teaching mechanism to capture semantic and structure knowledge at the same time. DPGNN [37] explicitly models the two-way selection preference for PJF using GCN. KGRM [39] model members and jobs as two graphs and fuse prior external knowledge, e.g., skill knowledge graph, into the graph representations.

Previous studies have made good use of textual information, skill entity, and direct interaction between members and jobs to



**Figure 2: Steps of WHIN pre-training. (a) Workplace heterogeneous graph with metapath. (b) Subgraph sampling for mini-batch pre-training. (c) A pre-training model with encoder-decoder architecture using Link-level pre-training task.**

help with Person-Job Fit. However, these studies often overlook the importance of workplace social relations and the incorporation of diverse, heterogeneous knowledge.

## 2.2 Heterogeneous Information Network-based Recommendation

Recommender systems have been widely deployed on the Internet to alleviate information overload [15]. Due to the excellent ability to model complex auxiliary information, Heterogeneous Information Network (i.e., HIN) has become one of the mainstream approaches in recommendation tasks. Depending on the training method, HIN-based models can be divided into *two-stage training-based models* and *end-to-end training-based models* [24].

Two-stage training-based models learn low-dimensional representations of nodes and graph structure using unsupervised tasks and use these representations on various downstream tasks [1, 2, 7, 14]. Inspired by DeepWalk [27] and node2Vec [11], metapath2vec [6] leverages metapath-based random walks to construct a heterogeneous neighborhood of nodes and leverages skip-gram model to generate node representations. Similarly, HIN2Vec [9] defines an unsupervised metapath prediction task and jointly learns metapath predictors and node representations. Many models are proposed based on the graph representation methods such as those above. For example, HERec generates node sequences with metapath-based random walks, uses node2vec to learn node representations, and completes recommendations by representation similarity. MAGNN [10] leverages node content features and information of intermediate nodes along the metapath by node content transformation, intra- and inter-metapath aggregation. Besides, there are some methods not relying on manual metapath, such as HRLHG [17], NREP [38], and ECHCDR [20].

Compared with two-stage training-based models, end-to-end-based models can use supervision signals directly while training. Thus, they are more customized to specific downstream tasks. For HINs with rich relations, relation-aware graph neural networks can achieve great results. For example, RGCN [29] assigns a weight matrix to each type of relation, thus extending GCN [35] to multi-relation graphs. DisenHAN [34] projects nodes into different subspaces using type-specific transformation and uses intra- and inter-aggregation to learn node representations. HAN [33] uses a dual attention mechanism to aggregate information from different metapaths. HPN [16] designs a semantic fusion mechanism for learning the importance of metapath and fusing them judiciously. By using

only its immediate connections as input, HGT [13] can learn and extract relevant metapaths for various tasks through its automatic and implicit attention mechanism without requiring manual metapath design.

In this paper, we adopt the two-stage training-based paradigm. During the pre-training stage, we integrate heterogeneous knowledge specific to the Person-Job Fit scenario to obtain diverse entity representations that encompass both structural and textual information. In the downstream stage, we design a novel model that leverages the information from the pre-training stage to filter out noise within professional networks. As a point of comparison, we have also selected two end-to-end-based models as baselines.

## 3 METHODOLOGY

Here we introduce our two-stage approach to improve the performance of the Person-Job Fit task by incorporating professional networks. The task and notation are formally defined in Section 3.1. In the first stage, we pre-train on the Workplace Heterogeneous Information Network (WHIN), which is detailed in Section 3.2. We introduce the pre-training approach on WHIN in Section 3.3. To overcome the noise in professional networks, we propose the contextual social attention graph neural network (CSAGNN) in the second stage, as described in Section 3.4.

### 3.1 Formal Task Description and Notations

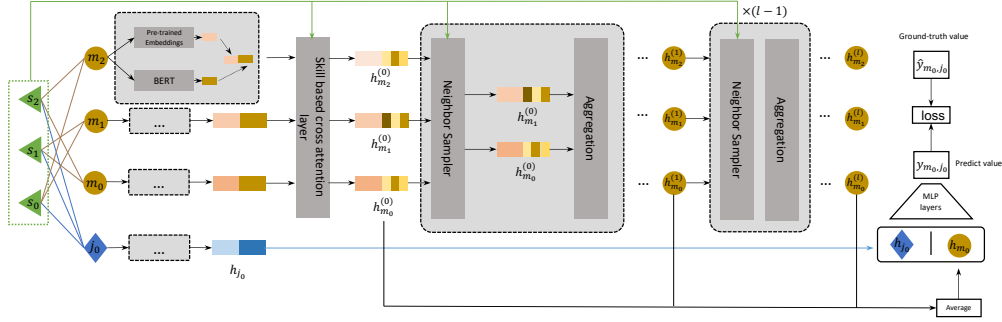
Person-Job Fit aims to match a set of job seekers with a set of job opportunities. Assume that there is a set of members  $\mathcal{M} = \{m_1, m_2, \dots, m_n\}$  who may be seeking job opportunities and a set of jobs  $\mathcal{J} = \{j_1, j_2, \dots, j_m\}$ . Formally, given a candidate pair  $(m_i, j_k)$ , the model is required to learn a function  $\mathcal{F}$  to predict the probability of whether the member applied for the job.

$$\mathcal{F}(m_i, j_k) = \begin{cases} 1 & m_i \text{ applied for } j_k \\ 0 & m_i \text{ didn't apply for } j_k \end{cases}$$

The set of candidate pairs can be defined as

$$T = \{(m_i, j_k) | m_i \in \mathcal{M}, j_k \in \mathcal{J}\}$$

In addition to the sets of members ( $\mathcal{M}$ ) and jobs ( $\mathcal{J}$ ), there exist three other sets of entities that are strongly related to members and jobs: skills ( $\mathcal{S}$ ), companies ( $\mathcal{C}$ ), and schools ( $\mathcal{H}$ ). These entities are referred as auxiliary information and can be leveraged to improve



**Figure 3: Architecture of CSAGNN.** When determining the match between job  $j_0$  and member  $m_0$ , the information from  $m_0$ 's professional connections,  $m_1$  and  $m_2$ , is simultaneously acquired. The initial representations of both the member and job are formed by concatenating the WHIN pre-training embedding with the representation obtained after processing the text information through BERT. All member text information is re-weighted through an attention mechanism based on the skills  $s_i$  required by  $j_0$ . A neighbor sampler module, operating based on the similarity between different members and the required skills, can filter out professional connections that are not relevant to the job.

the model's accuracy. Given an entity  $e$  and a relation  $r$ , an interaction map  $\mathcal{A}_e^r$  is used to specify the destination entities related to  $e$  through  $r$ . For example,  $\mathcal{A}_{m_i}^{apply}$  represents the set of jobs applied for by  $m_i$ . If  $m_i$  has not applied for any jobs,  $\mathcal{A}_{m_i}^{apply} = \emptyset$ .

### 3.2 The Construction of Workplace Heterogeneous Information Network

The Person-Job Fit scenario contains a wide range of heterogeneous knowledge, including various entities such as skills, companies, and schools, as well as relationships such as professional connections and skill requirements. We define a Workplace Heterogeneous Information Network (WHIN) to model this knowledge. As shown in Figure 1, WHIN is a heterogeneous information network consisting of five types of entities, namely Member, Job, Skill, Company, and School. It also contains nine types of relations, which are introduced to capture the different types of connections between the entities.

The Workplace Heterogeneous Information Network primarily consists of natural relations among diverse entities. Besides, two additional metapaths have been artificially constructed. The first metapath, co-apply ( $\mathcal{M}\text{-}\mathcal{J}\text{-}\mathcal{M}$ ), connects members who have applied for the same job. The second metapath, co-applied ( $\mathcal{J}\text{-}\mathcal{M}\text{-}\mathcal{J}$ ), connects multiple jobs that have been applied for by the same member. These metapaths were constructed based on prior knowledge and have proven to be useful in enhancing the performance of the WHIN pre-training process [10].

### 3.3 Pre-training on Workplace Heterogeneous Information Network

**3.3.1 link-level pre-training task.** We utilize a heterogeneous graph neural network with a link-level pre-training task to acquire node representations with rich structural knowledge. To improve the scalability of the pre-training process, we create mini-batches by sampling subgraphs from the entire Workplace Heterogeneous Information Network (WHIN), as illustrated in Figure 2 (a) and (b).

Specifically, a batch of candidate pairs is selected as the source nodes. From these source nodes, their k-hop neighbors are sampled

to construct a subgraph based on each direct relation or metapath. These identified relations and metapaths serve as positive samples. In addition, to introduce negative samples, we randomly select pairs of entities that do not exhibit the expected relationships. The link-level pre-training task revolves around determining the existence of the given samples. We employ an encoder-decoder structure, as illustrated in Figure 2 (c). The encoder learns node representations within each subgraph, while the decoder predicts link existence.

**3.3.2 pre-training model structure.** During the pre-training phase, we select a fixed number of tokens from the textual content of each entity type. These selected tokens are then processed through the BERT model [5]. The resulting BERT outputs are averaged to create initial representations for each entity.

In a formal context, when working with a textual representation associated with an entity denoted as  $e_i \in \mathcal{E}$ , we construct a corresponding initial embedding referred to as  $z_i^{(0)}$  by utilizing BERT. Specifically, We leverage member profiles, job descriptions, skill names, and other textual descriptions to initialize entity representations, thereby enhancing the initial embeddings for each entity type.

After obtaining an initialized representation with semantic information, we utilize an encoder, RGCN [29], to integrate the information between entities. The message-passing process in the RGCN encoder is represented by the following equation:

$$z_i^{(l+1)} = \sigma \left( \sum_{r \in R} \sum_{j \in \mathcal{A}_{e_i}^r} \frac{1}{c_{i,r}} W_r^{(l)} z_j^{(l)} + W_0^{(l)} z_i^{(l)} \right), \quad (1)$$

where  $W_r^{(l)}$  is a parameter matrix at the  $l$ -th layer of relation  $r$ ,  $W_0^{(l)}$  is used for self-loop at the  $l$ -th layer and  $c_{i,r} = |\mathcal{A}_{e_i}^r|$  is a normalization constant.

We employ  $L$  encoder layers for Link Prediction, utilizing the final representation  $z_i^{(L)}$  for entity  $e_i$ . We apply a multi-layer perceptron (MLP) to score each relation or metapath  $r_i$ , represented by vector  $M_{r_i}$ . The score for a particular link  $(s, r, d)$  is determined as follows:

$$y_{(s,r,d)} = \text{MLP}(z_s || M_r || z_d) \quad (2)$$

We employ cross entropy as the loss function for optimizing the pre-training model. Notably, the described pre-training process incorporates professional network information into entity embeddings.

$$\mathcal{L}_{pretrain} = -\left(y_{(s,r,d)} \log y'_{(s,r,d)} + (1 - y_{(s,r,d)}) \log (1 - y'_{(s,r,d)})\right) \quad (3)$$

### 3.4 Contextual Social Attention Graph Neural Network

As discussed earlier, professional networks are inherently noisy. To capture job-specific insights from the profiles of  $m_i$ 's professional connections, a multi-head attention (MHA) mechanism [32] is employed. As shown in Eq. 4,  $c_i$  is the profile representation of member  $m_i$  obtained from a pre-trained BERT model. The MHA mechanism uses each required skill embedding  $z_{s_k}$  as the query, and  $c_i$  as both the key and value. It computes attention weights between  $c_i$  and the required skill embeddings. The weighted sum averages these attention outputs for all required skills. This results in the job-specific contextual feature  $F_{m_i}^c(j_k)$  that captures relevant information from  $m_i$ 's connections for job  $j_k$ .

$$F_{m_i}^c(j_k) = \frac{1}{|\mathcal{A}_{j_k}^{require}|} \sum_{s_k \in \mathcal{A}_{j_k}^{require}} \text{MHA}(Q = z_{s_k}, K = c_i, V = c_i) \quad (4)$$

As shown in Eq. 5, contextual feature  $F_{m_i}^c(j_k)$  is concatenated with the WHIN pre-trained feature  $F_{m_i}^s$  to form member's  $j_k$ -specific initial feature.

$$h_{m_i}^{(0)}(j_k) = (F_{m_i}^c(j_k) || F_{m_i}^s) \quad (5)$$

The relevance degree of professional connections to the job  $j_k$  needs to be considered in the messaging process. Here, we use the distance between WHIN pre-trained representations of skills mastered by members and the skills required by jobs as the weights during message passing.

$$z'(\hat{S}) = \sum_{s \in \hat{S}} z_s / |\hat{S}| \quad (6)$$

$$d_{m_i}(j_k) = \frac{z'(\mathcal{A}_{m_i}^{master}) \cdot z'(\mathcal{A}_{j_k}^{require})}{\sqrt{(z'(\mathcal{A}_{m_i}^{master}))^2} \sqrt{(z'(\mathcal{A}_{j_k}^{require}))^2}} \quad (7)$$

Given a set of skills  $\hat{S}$ , average embedding can be described in Eq. 6 and distance between members' and jobs' skills  $d_{m_i}(j_k)$  can be described in Eq. 7. Further, due to some members or jobs may have hundreds of related skills, we sampled the skill set like  $\mathcal{A}_{m_i}^{master}$  and  $\mathcal{A}_{j_k}^{require}$  in Eq. 4 and Eq. 7. The number of sampled skills,  $n_s$ , is a hyperparameter of CSAGNN. We will discuss the influence of  $n_s$  in Section 4.3.

$$h_{m_i}^{(l+1)}(j_k) = \sigma \left( W_2^{(l)} \sum_{m' \in \mathcal{A}_{m_i}^{connect}} \alpha_{m_j}(j_k) h_{m_j}^{(l)}(j_k) + W_1^{(l)} h_{m_i}^{(l)} \right) \quad (8)$$

$$\alpha_{m_j}(j_k) = \frac{d_{m_j}(j_k)}{\sum_{m' \in \mathcal{A}_{m_i}^{connect}} d_{m'}(j_k)} \quad (9)$$

As shown in Eq. 8, we can get the messaging function of professional network enhanced GNN by text attention and skill distance attention where  $\alpha_{m_j}(j_k)$  is normalized message aggregation weight described in Eq. 9. We use the average of each layer's representation as the final representation of  $m_i$ :

$$h_{m_i}(j_k) = \sum_{l=0}^{L+1} h_i^{(l)} j_k / L \quad (10)$$

Compared with member profiles, job descriptions provided by recruiters often have a high quality. Thus, we directly use contextual features and WHIN pre-trained features as the final representation for all job entities.

$$h_{j_k} = (F_{j_k}^c || F_{j_k}^s) \quad (11)$$

$h_{j_k}(m_i)$  and  $h_{m_i}(j_k)$  will be concatenated and passed through a vanilla multi-layer Perceptron (MLP) to get the predicted value  $y_{m_i,j_k}$ , and the cross entropy loss could be formulated as Eq. 13, where  $y'_{m_i,j_k}$  is the ground truth:

$$y_{m_i,j_k} = \text{MLP}(h_{m_i}(j_k) || h_{j_k}) \quad (12)$$

$$\mathcal{L}_{csa} = -\left(y_{m_i,j_k} \log y'_{m_i,j_k} + (1 - y_{m_i,j_k}) \log (1 - y'_{m_i,j_k})\right) \quad (13)$$

## 4 EXPERIMENTS

In this section, we will validate the performance of the CSAGNN model on three real datasets and answer the following research questions:

- **RQ1:** Can CSAGNN outperform existing models on real-world datasets?
- **RQ2:** What are the effects of different components in our model?
- **RQ3:** Are professional networks helpful for the Person-Job Fit task?
- **RQ4:** Does the WHIN pre-training provide meaningful representations of skills for CSAGNN?

### 4.1 Experiment Setup

**4.1.1 Datasets.** We validate our model on three large real-world datasets from LinkedIn, a well-known workplace social platform. Datasets are divided by the industry to which the members and jobs belong. The first two datasets contain members and jobs from the technology and finance industries. Due to the singularity of the industry, social noise is limited. The third dataset, which includes members and jobs from various industries such as healthcare, education, and semiconductors, is intended to test the ability of CSAGNN



**Table 1: Performance comparison of all baselines and our models. The best and second-best results are shown in bold and underlined, respectively. All values are multiplied by 100.**

Dataset Metric		Models							
		LightGCN	PJFNN	BPJFNN	APJFNN	MV-CoN	HAN	HGT	CSAGNN
Tech	AUC	63.94 ± 2.08	64.40 ± 1.85	64.62 ± 2.23	64.59 ± 2.02	65.53 ± 3.69	66.02 ± 3.52	<u>67.63 ± 2.63</u>	<b>68.57 ± 3.46</b>
	ACC	63.38 ± 3.21	63.45 ± 1.65	63.71 ± 2.91	63.64 ± 3.38	64.37 ± 4.26	64.19 ± 2.57	<u>66.81 ± 2.21</u>	<b>68.76 ± 1.91</b>
	F <sub>1</sub>	61.73 ± 1.38	62.22 ± 3.48	62.51 ± 3.19	62.14 ± 1.76	63.96 ± 3.44	63.72 ± 1.83	<u>64.97 ± 1.54</u>	<b>65.22 ± 2.67</b>
	AP	57.38 ± 1.79	58.19 ± 2.17	58.38 ± 2.46	57.91 ± 3.34	58.73 ± 3.55	59.10 ± 3.90	<u>60.99 ± 2.85</u>	<b>61.59 ± 3.46</b>
Finance	AUC	64.18 ± 1.71	64.60 ± 1.79	65.81 ± 2.81	65.59 ± 3.32	66.24 ± 2.90	66.17 ± 2.71	<u>68.43 ± 3.87</u>	<b>69.60 ± 1.83</b>
	ACC	62.89 ± 2.65	63.80 ± 2.29	64.60 ± 3.48	64.28 ± 3.24	64.47 ± 3.73	<u>67.87 ± 3.98</u>	67.83 ± 2.06	<b>69.69 ± 3.23</b>
	F <sub>1</sub>	61.91 ± 1.86	62.78 ± 1.57	63.43 ± 1.73	63.41 ± 1.81	64.26 ± 2.98	64.56 ± 3.96	<u>65.59 ± 2.71</u>	<b>67.49 ± 1.84</b>
	AP	58.96 ± 1.25	59.37 ± 2.23	60.37 ± 2.32	60.14 ± 3.27	60.58 ± 3.17	60.38 ± 3.31	<u>61.06 ± 2.65</u>	<b>62.70 ± 3.35</b>
Hybrid	AUC	66.98 ± 1.23	67.27 ± 2.15	68.79 ± 1.99	68.81 ± 1.84	69.52 ± 2.91	69.39 ± 3.90	<u>70.60 ± 1.54</u>	<b>72.27 ± 1.64</b>
	ACC	65.93 ± 1.57	67.24 ± 3.06	68.46 ± 2.75	68.73 ± 2.32	69.33 ± 3.46	69.42 ± 2.88	<u>69.71 ± 3.98</u>	<b>72.37 ± 1.92</b>
	F <sub>1</sub>	64.62 ± 1.84	65.25 ± 1.98	65.35 ± 2.47	65.36 ± 3.86	64.80 ± 3.82	64.03 ± 2.44	<u>65.92 ± 1.91</u>	<b>69.58 ± 2.05</b>
	AP	60.52 ± 1.67	61.57 ± 1.16	62.31 ± 1.72	62.26 ± 1.05	63.59 ± 3.69	63.33 ± 3.96	<u>64.02 ± 3.28</u>	<b>64.49 ± 2.33</b>

to filter noise in professional networks. This hybrid dataset implies more complex social relationships and a wider range of job opportunities. The WHINs constructed from these three datasets all contain about 200,000 entities and 10 million links. Table 2 displays the statistics of the datasets. The skills, schools, and companies in the dataset all come from explicit annotations by members or jobs on LinkedIn. We cannot provide further detailed information about the datasets to protect the users' privacy. Each candidate pair within the datasets is labeled as either positive or negative based on the feedback provided by members.

**Table 2: Statistics of datasets. M refers to Member, J refers to Job, S refers to Skill, CP refers to Candidate Pair, and PC refers to Professional Connection.**

Industry	#M	#J	#S	#CP	#PC
Tech	33000+	62000+	27000+	136000+	1922000+
Finance	20000+	27000+	23000+	36000+	615000+
Hybrid	83000+	120000+	33000+	200000+	2768000+

**4.1.2 Baseline Models.** We compare our model with the following baseline models.

- **LightGCN [12]** is a simplified Graph Convolutional Network model for collaborative filtering with competitive performance and less complexity.
- **PJFNN [41]** is a method based on a convolutional neural network (CNN). Hierarchical CNN encodes Resumes and job descriptions independently, and the matching degree is calculated by cosine similarity.
- **BPJFNN [28]** leverages bidirectional LSTM to learn the representations of resumes and job descriptions.

- **APJFNN [28]** leverages bidirectional LSTM and hierarchical attention mechanism to learn the representations of resumes and job descriptions.
- **MV-CoN [3]** combines text matching model and RGCN to learn representations of resumes and job descriptions.
- **HAN [33]** uses a dual attention mechanism to aggregate neighbor information via different metapaths.
- **HGT [13]** designs node- and edge-type dependent parameters to characterize the heterogeneous attention over each edge.

These baselines can be divided into three groups: (1) context-based models that treat Person-Job Fit as a text match problem and use contextual knowledge from members' profiles and jobs' descriptions: PJFNN, BPJFNN, APJFNN, and MV-CoN. In particular, MV-CoN additionally introduces structural information to enhance model performance. (2) collaborative filtering-based model that uses direct interactions between members and jobs: LightGCN. (3) end-to-end heterogeneous graph neural network models: HAN, HGT.

**4.1.3 Evaluation and Implementation Details.** We use four widely used metrics to evaluate the ranking performance: AUC, accuracy (ACC), F<sub>1</sub>, and average precision (AP).

The Person-Job Fit models, namely PJFNN, BPJFNN, and APJFNN, are implemented using RecBole—an established open-source recommendation library [40]. The MV-CoN model leverages the original code provided in their respective paper [3]. Other models are implemented with PyTorch Geometric [8]. We have employed BERT-Tiny<sup>6</sup> to reduce computational costs.

The embedding dimensions for all models are standardized at 32. As discussed in Section 3.3.1, subgraph sampling for mini-batch

<sup>6</sup><https://huggingface.co/prajjwal1/bert-tiny>


WHIN pre-training is executed using the PyTorch Geometric sub-graph sampler [8]. Subgraph construction starts from nodes within a batch of candidate pairs, and each hop samples five neighbors via relations (or metapaths). Reconnecting all relations (and metapaths) between the sampled nodes follows this. We set the number of sampled hops to 3. The number of sampled hops is fixed at 3. In CSAGNN, the count of sampled skills ( $n_s$ ) is designated as 10, and CSAGNN comprises two layers. The first 128 tokens from all text inputs are captured for all models. All models are optimized using the Adam optimizer [19], with a learning rate adjusted from 0.01 to 0.0001. For the purpose of evaluation, the dataset utilized across all models is randomly partitioned into a ratio of 8:1:1 for training, validation, and testing respectively. Three independent experiments are conducted, each one repeated, to ascertain consistent and reliable outcomes.

## 4.2 The Overall Comparison (RQ1)


Table 1 shows the performance of all baseline models and our model, CSAGNN. The results indicate better performance on the hybrid dataset than on individual industry datasets, possibly because jobs within the same industry are more similar and thus harder to distinguish. End-to-end heterogeneous graph neural network models have shown superior performance on all three datasets compared to context-based and CF-based models. HGT, utilizing heterogeneous knowledge, has improved AUC scores by 3.20%, 3.31%, and 1.55% compared to the best baseline MV-CoN, which employs homogeneous knowledge. CSAGNN further improved AUC scores by 1.39%, 1.71%, and 2.37% compared to HGT.

**Table 3: Ablation studies conducted on all datasets, with all values multiplied by 100.**

Industry	Model	AUC	ACC	F1	AP
Tech	CSAGNN	<b>68.57</b>	<b>68.76</b>	<b>65.22</b>	<b>61.59</b>
	w/o S	68.26	68.42	64.41	61.34
	w/o A	68.23	68.40	64.38	61.29
	w/o CSA	68.06	68.22	64.26	61.21
	w/o CSA&H	62.84	62.78	61.35	57.53
Finance	CSAGNN	<b>69.60</b>	<b>69.69</b>	<b>67.55</b>	<b>62.77</b>
	w/o S	69.39	69.48	66.68	62.21
	w/o A	69.22	69.32	66.44	62.32
	w/o CSA	68.98	69.04	66.15	62.06
	w/o CSA&H	63.75	62.78	61.49	58.74
Hybrid	CSAGNN	<b>72.27</b>	<b>72.37</b>	<b>69.58</b>	<b>64.49</b>
	w/o S	70.81	71.67	68.68	64.67
	w/o A	70.03	71.47	68.18	64.06
	w/o CSA	68.37	69.39	64.84	62.65
	w/o CSA&H	65.85	64.32	63.26	59.49



(a) Model performance varying the number of sampled skills while fixing CSAGNN layers to 1.



(b) Model performance varying the number of CSAGNN layers while fixing the number of sampled skills to 10.

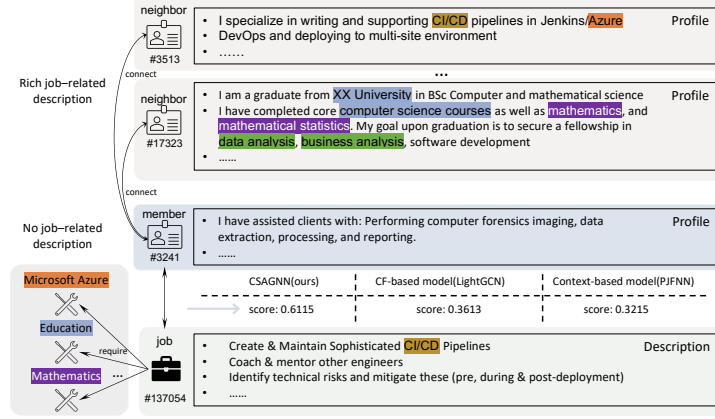
**Figure 4: Hyperparameter tuning experiments to investigate the specific effects of social relations and job-specific attention mechanism.**

## 4.3 Ablation Studies and Parameter Tuning (RQ2)


Our approach employs two main techniques: WHIN pre-training, which integrates heterogeneous knowledge, including professional networks, and CSAGNN, which incorporates professional networks with an attention mechanism. We conducted ablation studies to analyze the effectiveness of different techniques, where we considered the following variants of CSAGNN: (1) CSAGNN w/o S removes messages from professional connections but retains the job-specific attention mechanism for members themselves. (2) CSAGNN w/o A removes the job-specific attention mechanism, using only structural knowledge from the WHIN pre-train approach. (3) CSAGNN w/o CSA removes both professional connections and job-specific attention mechanisms, using only structural knowledge from the WHIN pre-train approach. (4) CSAGNN w/o CSA&H removes WHIN pre-trained embeddings and CSA mechanism, only using text information as input and leverages MLP to predict.

As detailed in Table 3, within single industry datasets with low social noise, the job-specific attention mechanism and professional connections' messages yielded limited benefits. However, WHIN pre-training embeddings were crucial for enhancing performance. On hybrid datasets with high social noise, where candidate pairs may span different industries, the job-specific attention mechanism, along with the WHIN pre-training embedding, significantly improved the model.

To investigate the specific roles of professional networks and job-specific attention mechanisms in our model, we performed hyperparameter tuning experiments by fixing one module and varying the other to observe its impact on performance. Specifically, we first fixed the number of CSAGNN layers to 1 and tested the model's performance with different numbers of sampled skills ( $n_s$ ), as described in Section 3.4. We then fixed  $n_s$  to 10 and tested the model with different numbers of CSAGNN layers. The results on all datasets showed that, compared to the neighbor sampling range, the gain



**Figure 5: A case where professional connections improve performance on Person-Job Fit. CSAGNN can improve the performance of Person Job Fit by filtering and aggregating information from professional networks. For privacy protection reasons, we have rewritten the statement in the example while ensuring that the semantic information remains unchanged.**



**Figure 6: Visualization of WHIN Pre-trained skill embeddings showing clear distinction between Technology and Health-Related Skills.**

in model performance from the attention mechanism saturates when the number of sampled skills is relatively small. This observation led us to explore the balance between model performance and efficiency by selecting a smaller value for  $n_s$ .

#### 4.4 Case Study (RQ3)

The value of utilizing professional networks is illustrated in the example shown in Figure 5. In this case, we have rephrased the information to preserve the users' privacy while retaining its essential meaning. Even though the member's profile contains minimal information, the CSAGNN model adeptly predicts their classification by harnessing job-related insights from the member's professional connections. This contrasts sharply with the incorrect predictions made by the collaborative filtering-based LightGCN and context-based PJFNN models. This specific example accentuates the vital role of professional connections in the classification process and underlines the CSAGNN model's distinctive capability to leverage such relationships for precise predictions.

#### 4.5 Visualization (RQ4)

To demonstrate the effectiveness of WHIN pre-training in capturing meaningful skill representations for CSAGNN, we analyzed the

pre-trained embeddings of selected skills in the embedding space. Utilizing Principal Component Analysis (PCA) as our method of dimensional reduction [18], we were able to visualize the WHIN pre-trained embeddings for various skills. Figure 6 reveals that skills closely related to programming, such as C++ and Python, are grouped together in the embedding space. In contrast, skills associated with healthcare form a distinct cluster. It is noteworthy that these classifications were achieved using only the skill names as initialization. Yet, our WHIN pre-training method successfully distinguishes between skill categories by learning from heterogeneous knowledge, highlighting the capability of WHIN to recognize and differentiate skills across multiple domains.

## 5 CONCLUSION

This paper introduces a novel two-stage approach for leveraging professional networks in Person-Job Fit, including the formation of the Workplace Heterogeneous Information Network (WHIN). WHIN encompasses various entities such as members, jobs, skills, companies, and schools, with an emphasis on the professional connections among members. By employing heterogeneous graph pre-training techniques, the approach acquires representations that integrate professional connections and other diverse information for different entities. These representations are subsequently applied to the CSAGNN model, helping to filter out social noise.

Experimental results show that professional connections provide valuable job-specific insights. The WHIN pre-training method is also promising for applications like skill completion and professional connection recommendations, opening new research avenues. Furthermore, for large-scale applications of CSAGNN, a key area of future work involves reducing the model's computational overhead [21–23].

## ACKNOWLEDGMENTS

Thanks to Yihan Cao and Yushu Du for their insightful discussion. Thanks to Sriram Vasudevan and Peide Zhong for their comprehensive review feedback.



## ETHICAL CONSIDERATIONS

The integration of professional networks into Person-Job Fit models offers significant potential for enriching recommendations. However, we must recognize and address two vital ethical considerations that may lead to adverse societal implications:

Firstly, the core of our approach involves accessing members' professional connections and corresponding profiles, which raises privacy concerns. In our research, we ensured that the data collection and processing respected privacy by adhering to proper consent mechanisms and limiting access to pertinent information. Future users should follow suit, being mindful of the need for clear and voluntary consent from members and carefully controlling access to connection information.

Additionally, connections often exhibit demographic clustering, which could introduce biases favoring certain groups within networks. The models might infer sensitive attributes like race, gender, or age. Future users should conduct audits to ensure demographic equity and prohibit direct utilization of protected class information.

In summary, while the enhancement of Person-Job Fit models through professional network data brings advancements, it also introduces risks concerning privacy and fairness. These ethical challenges call for a concerted effort from platforms to resolve, ensuring that the innovations foster an environment that is both respectful of individual rights and free from discriminatory biases.

## REFERENCES

- [1] Wendong Bi, Lun Du, Qiang Fu, Yanlin Wang, Shi Han, and Dongmei Zhang. 2022. Make heterophily graphs better fit gnn: A graph rewiring approach. *arXiv preprint arXiv:2209.08264* (2022).
- [2] Wendong Bi, Lun Du, Qiang Fu, Yanlin Wang, Shi Han, and Dongmei Zhang. 2023. MM-GNN: Mix-Moment Graph Neural Network towards Modeling Neighborhood Feature Distribution. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*. 132–140.
- [3] Shuqing Bian, Xu Chen, Wayne Xin Zhao, Kun Zhou, Yupeng Hou, Yang Song, Tao Zhang, and Ji-Rong Wen. 2020. Learning to match jobs with resumes from sparse interaction data using multi-view co-teaching network. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 65–74.
- [4] Shuqing Bian, Wayne Xin Zhao, Yang Song, Tao Zhang, and Ji-Rong Wen. 2019. Domain adaptation for person-job fit with transferable deep global match network. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*. 4810–4820.
- [5] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. <https://doi.org/10.18653/v1/N19-1423>
- [6] Yuxiao Dong, Nitesh V Chawla, and Ananthram Swami. 2017. metapath2vec: Scalable representation learning for heterogeneous networks. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*. 135–144.
- [7] Lun Du, Xu Chen, Fei Gao, Qiang Fu, Kunqing Xie, Shi Han, and Dongmei Zhang. 2022. Understanding and Improvement of Adversarial Training for Network Embedding from an Optimization Perspective. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*. 230–240.
- [8] Matthias Fey and Jan E. Lenssen. 2019. Fast Graph Representation Learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*.
- [9] Tao-yang Fu, Wang-Chien Lee, and Zhen Lei. 2017. Hin2vec: Explore meta-paths in heterogeneous information networks for representation learning. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. 1797–1806.
- [10] Xinyu Fu, Jiani Zhang, Ziqiao Meng, and Irwin King. 2020. Magnn: Metapath aggregated graph neural network for heterogeneous graph embedding. In *Proceedings of The Web Conference 2020*. 2331–2341.
- [11] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*. 855–864.
- [12] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*. 639–648.
- [13] Ziniu Hu, Yuxiao Dong, Kuansan Wang, and Yizhou Sun. 2020. Heterogeneous graph transformer. In *Proceedings of The Web Conference 2020*. 2704–2710.
- [14] Jincheng Huang, Lun Du, Xu Chen, Qiang Fu, Shi Han, and Dongmei Zhang. 2023. Robust Mid-Pass Filtering Graph Convolutional Networks. In *Proceedings of the ACM Web Conference 2023*. 328–338.
- [15] Jacob Jacoby. 1984. Perspectives on information overload. *Journal of consumer research* 10, 4 (1984), 432–435.
- [16] Houye Ji, Xiao Wang, Chuan Shi, Bai Wang, and Philip Yu. 2021. Heterogeneous graph propagation network. *IEEE Transactions on Knowledge and Data Engineering* (2021).
- [17] Zhuoren Jiang, Yue Yin, Liangcai Gao, Yao Lu, and Xiaozhong Liu. 2018. Cross-language citation recommendation via hierarchical representation learning on heterogeneous graph. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. 635–644.
- [18] Ian T Jolliffe and Jorge Cadima. 2016. Principal component analysis: a review and recent developments. *Philosophical transactions of the royal society A: Mathematical, Physical and Engineering Sciences* 374, 2065 (2016), 20150202.
- [19] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [20] Jin Li, Zhaohui Peng, Senzhang Wang, Xiaokang Xu, Philip S Yu, and Zhenyun Hao. 2020. Heterogeneous Graph Embedding for Cross-Domain Recommendation Through Adversarial Learning. In *International Conference on Database Systems for Advanced Applications*. Springer, 507–522.
- [21] Tianyi Li, Lu Chen, Christian S Jensen, and Torben Bach Pedersen. 2021. TRACE: Real-time compression of streaming trajectories in road networks. *Proceedings of the VLDB Endowment* 14, 7 (2021), 1175–1187.
- [22] Tianyi Li, Ruikai Huang, Lu Chen, Christian S Jensen, and Torben Bach Pedersen. 2020. Compression of uncertain trajectories in road networks. *Proceedings of the VLDB Endowment* 13, 7 (2020), 1050–1063.
- [23] Zhiqi Lin, Cheng Li, Youshan Miao, Yunxin Liu, and Yinlong Xu. 2020. Pagraph: Scaling gnn training on large graphs via computation-aware caching. In *Proceedings of the 11th ACM Symposium on Cloud Computing*. 401–415.
- [24] Jiawei Liu, Chuan Shi, Cheng Yang, Zhiyuan Lu, and S Yu Philip. 2022. A survey on heterogeneous information network based recommender systems: Concepts, methods, applications and resources. *AI Open* (2022).
- [25] Yao Lu, Sandy El Helou, and Denis Gillet. 2013. A recommender system for job seeking and recruiting website. In *Proceedings of the 22nd International Conference on World Wide Web*. 963–966.
- [26] Jochen Malinowski, Tobias Keim, Oliver Wendt, and Tim Weitzel. 2006. Matching people and jobs: A bilateral recommendation approach. In *Proceedings of the 39th Annual Hawaii International Conference on System Sciences (HICSS'06)*, Vol. 6. IEEE, 137c–137c.
- [27] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 701–710.
- [28] Chuan Qin, Hengshu Zhu, Tong Xu, Chen Zhu, Liang Jiang, Enhong Chen, and Hui Xiong. 2018. Enhancing person-job fit for talent recruitment: An ability-aware neural network approach. In *The 41st international ACM SIGIR conference on research & development in information retrieval*. 25–34.
- [29] Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. 2018. Modeling relational data with graph convolutional networks. In *European semantic web conference*. Springer, 593–607.
- [30] Walid Shalaby, BahaaEddin AlAila, Mohammed Korayem, Layla Pournajaf, Khalifeh AlJadda, Shannon Quinn, and Wlodek Zadrozny. 2017. Help me find a job: A graph-based approach for job recommendation at scale. In *2017 IEEE international conference on big data (big data)*. IEEE, 1544–1553.
- [31] Dazhong Shen, Hengshu Zhu, Chen Zhu, Tong Xu, Chao Ma, and Hui Xiong. 2018. A joint learning approach to intelligent job interview assessment. In *IJCAI*, Vol. 18. 3542–3548.
- [32] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).
- [33] Xiao Wang, Houye Ji, Chuan Shi, Bai Wang, Yanfang Ye, Peng Cui, and Philip S Yu. 2019. Heterogeneous graph attention network. In *The world wide web conference*. 2022–2032.
- [34] Yifan Wang, Suyao Tang, Yuntong Lei, Weiping Song, Sheng Wang, and Ming Zhang. 2020. Disentangle: Disentangled heterogeneous graph attention network for recommendation. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 1605–1614.
- [35] Max Welling and Thomas N Kipf. 2016. Semi-supervised classification with graph convolutional networks. In *J. International Conference on Learning Representations (ICLR 2017)*.

- [36] Rui Yan, Ran Le, Yang Song, Tao Zhang, Xiangliang Zhang, and Dongyan Zhao. 2019. Interview choice reveals your preference on the market: To improve job-resume matching through profiling memories. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 914–922.
- [37] Chen Yang, Yupeng Hou, Yang Song, Tao Zhang, Ji-Rong Wen, and Wayne Xin Zhao. 2022. Modeling Two-Way Selection Preference for Person-Job Fit. In *Proceedings of the 16th ACM Conference on Recommender Systems*. 102–112.
- [38] Libin Yang, Zeqing Zhang, Xiaoyan Cai, and Lantian Guo. 2019. Citation recommendation as edge prediction in heterogeneous bibliographic network: a network representation approach. *IEEE Access* 7 (2019), 23232–23239.
- [39] Kaichun Yao, Jingshuai Zhang, Chuan Qin, Peng Wang, Hengshu Zhu, and Hui Xiong. 2022. Knowledge Enhanced Person-Job Fit for Talent Recruitment. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*. IEEE, 3467–3480.
- [40] Wayne Xin Zhao, Shanlei Mu, Yupeng Hou, Zihan Lin, Yushuo Chen, Xingyu Pan, Kaiyuan Li, Yujie Lu, Hui Wang, Changxin Tian, et al. 2021. Recbole: Towards a unified, comprehensive and efficient framework for recommendation algorithms. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. 4653–4664.
- [41] Chen Zhu, Hengshu Zhu, Hui Xiong, Chao Ma, Fang Xie, Pengliang Ding, and Pan Li. 2018. Person-job fit: Adapting the right talent for the right job with joint representation learning. *ACM Transactions on Management Information Systems (TMIS)* 9, 3 (2018), 1–17.