Multi-Relational Graph Attention Network for Social Relationship Inference from Human Mobility Data

Guangming Qin, **Jianpeng Qi**, **Bin Wang**, **Guiyuan Jiang**, **Yanwei Yu*** and **Junyu Dong**
Ocean University of China
qinguangming@stu.ouc.edu.cn, {qijianpeng, wangbin9545, jiangguiyuan, yuyanwei, dongjunyu}@ouc.edu.cn

Abstract

Inferring social relationships from human mobility data holds significant value in real-life spatio-temporal applications, inspiring the development of a series of graph-based methods for deriving such relationships. However, despite their noted effectiveness, we argue that previous methods either rely solely on direct relations between users, neglecting valuable user mobility patterns, or have not fully harnessed the indirect interactions, thereby struggling to capture users' mobility preferences. To address these issues, in this work, we propose the Multi-Relational Graph Attention Network (MR-GAN), a novel graph attention network, which is able to explicitly model indirect relations and effectively capture their different impact. Specifically, we first extract a multi-relational graph from heterogeneous mobility graph to explicitly model the direct and indirect relations, and then utilize influence attention and cross-relation attention to further capture the different influence between users, and different importance of relations for each user. Comprehensive experiments on three real-world mobile datasets demonstrate that the proposed model significantly outperforms state-of-the-art models in predicting social relationships between users. The source code of our model is available at https://github.com/qinguangming1999/MRGAN_IJCAI.

1 Introduction

Location-based social services, such as Foursquare, Yelp, Facebook, and Twitter, are increasingly affecting human daily life and have generated a considerable amount of user spatio-temporal mobility data that record their daily behaviors [Li *et al.*, 2023a]. Revealing social relationships from these mobility trajectory data is crucial to understanding human dynamics, which is widely used to study the impact of human mobility and social bonds on each other [Cho *et al.*, 2011; Yang *et al.*, 2019]. As demonstrated in [Wang *et al.*, 2014; Pham *et al.*, 2016], user mobility data can indeed serve as a strong predictor for inferring social ties. Specifically, estimating social relationships based on users' check-in mobility data is beneficial to various real-world recommendation systems, such as friend suggestion [Yu *et al.*, 2020; Yu *et al.*, 2021], location recommendations [Dai *et al.*, 2022; Li *et al.*, 2024; Wu *et al.*, 2023a], and targeted advertising [Sharma *et al.*, 2021].

Recent years have brought advances in leveraging network representation learning [Liu *et al.*, 2021; Cui *et al.*, 2018; Sun *et al.*, 2019] to infer mobility social relationships, such as walk2friends [Backes *et al.*, 2017], emb-cat [Yu *et al.*, 2018], vec2link [Zhou *et al.*, 2018], Heter-GCN [Wu *et al.*, 2019], MVMN [Zhang *et al.*, 2020], and HMGCL [Li *et al.*, 2023b]. These methods often model users and check-in locations into homogeneous or heterogeneous networks (*e.g.*, user meeting graph), use network embedding technologies to learn vector representations of user nodes, and use node representation similarities to predict mobility social relationships. To utilize check-in semantic information, a complex heterogeneous hypergraph [Xia *et al.*, 2022] is recently proposed [Yang *et al.*, 2019; Yang *et al.*, 2022; Huynh *et al.*, 2022; Li *et al.*, 2022]. These models treat check-ins as hyperedges, and learn user embedding from hypergraph via preserving high-order proximity between users and check-ins.

However, most previous studies either infer potential links between users from direct relations between users (such as user meetings) or unified learn corrections between users from complex interactions between users and check-ins (with contextual semantics). Most recently, SRINet [Qin *et al.*, 2023], a model based on graph structure learning, has been proposed, which claims that by removing noise edges in the user meeting network, graph neural networks can be used to learn effective user representations for inferring friendships. Nonetheless, these methods have the following two major limitations: *First*, in user mobility data, especially commonly used meeting relation or co-locations between users, although they can reflect users' social activities to some extent, they are only partially visible (because many real meetings may not be recorded through check-ins). *Solely relying on this partially visible direct relation to infer social links between users is obviously not sufficient. Second*, various indirect relations reflecting social activities between users have not been fully utilized, and *their effect and impact on social relationship in-*

*Corresponding author: Yanwei Yu.

ference have not been explicitly modeled and effectively captured. For example, the impact of users' preferences for different locations on friend prediction has not been explored in previous user representation learning.

To address the aforementioned challenges, we propose a Multi-Relational Graph Attention Network (MRGAN) for inferring mobility social relationships. Specially, we first build a heterogeneous mobility graph by combining user check-in relation and partially known friendships, as well as user meeting relations and location co-occurrences mined from user mobility data. In addition to the direct relations, *e.g.*, user meeting and friendship, which have been proven to be strongly related to social ties, MRGAN also aims to capture the impact of users' personalized preferences for locations and users' common preferences for location co-occurrences on social relationship inference, respectively. To this end, MRGAN constructs a multi-relational graph derived from a heterogeneous mobility graph. This graph encompasses not only direct relations such as friendships and meetings for the preservation of social connections but also explicitly models indirect relations, including personalized preferences and common preferences, for exploring user mobility preferences. To learn the different impacts of these relations on social relationship inference, MRGAN first separately learns users' relation-specific embedding and utilizes an attention mechanism to learn their different influence on each other, and eventually fuse the relation-specific user embedding with their importance to infer social ties among users. Extensive experiments on three real-life mobility datasets demonstrate that the proposed model significantly outperforms all state-of-the-art baselines, with an average improvement of 4.28% and 5.51% in terms of PRAUC and ROCAUC.

To summarize, we make the following contributions:

- To improve the performance of mobility social relationship inference, we explicitly model the impact of users' personalized preferences for locations and users' common preferences for location co-occurrence on social relationship inference.
- We propose a novel multi-relational graph attention network that effectively captures and integrates multiple relations between users that reflect users' social relationships, *i.e.*, direct known friendships and meeting relations, and indirect personalized location preference relations and common location co-occurrence preference relations.
- We perform extensive experiments on three real-life mobility datasets, and the experimental results verify the superiority of the proposed model compared to all baseline methods in terms of all metrics.

2 Related Work

In the literature, several studies delve into the inference of social relationships through the analysis of human mobility trajectories. Eagle *et al.* [2009] aim to unveil the correlation between meeting events and social relationships by leveraging mobile phone data. Li *et al.* [2011] find that the timing of meetings could serve as an indicator for distinguishing between various types of relationships. An entropy-based model (EBM), proposed by [Pham *et al.*, 2013], considers the diversity of meeting events and location entropy in friendship inference. Wang *et al.* [2014] devise a unified framework, denoted as PGT, that takes into account personal background, global background, and temporal factors. Pham *et al.* [2016] further extend EBM via integrating location semantic information and stay duration.

The advancements in representation learning and graph deep learning [Cai *et al.*, 2023; Wu *et al.*, 2023b; Dai *et al.*, 2023; Wang *et al.*, 2023; Antelmi *et al.*, 2024] have led to the proposal of more sophisticated methods for social relationship inference. Backes *et al.* [2017] apply random-walk based embedding technique on a user-location bipartite graph to predict friendship. Yu *et al.* [2018] construct a user meeting graph and learning embedding on it via hierarchical walk sampling to infer mobility relationship. Inspired by [Schlichtkrull *et al.*, 2018], Wu *et al.* [2019] first employ GCN on heterogeneous graph to infer relationship. Yang *et al.* [2019] present an LBSN hypergraph comprising user, location, time, and semantic nodes, investigating the mutual influence of mobility and social ties on each other. Yang *et al.* [2022] further extend [Yang *et al.*, 2019] to a heterogeneous hypergraph to consider the heterogeneous nature of LBSN data. Zhang *et al.* [2020] develop a multi-view matching network designed to acquire three view-specific representations and integrate them for the ultimate relationship inference. Li *et al.* [2022] model user trajectories and check-in records as hyperedges in a heterogeneous hypergrah, and learn user embedding using attention and contrastive learning for friend recommendation. Huynh *et al.* [2022] develop an embedding technique that encompasses high-order, dynamic, and multi-role contexts in LBSN hypergraph data, catering to both friend suggestion and POI recommendation.

Recently, Li *et al.* [2023b] construct a mobility social multi-graph with multi types of edges and apply supervised constractive learning and attention to learn user embedding for friend recommendation. To solve the noise issue in mobility data, Qin *et al.* [2023] propose a social relationship inference framework, which weeds out noise edge on a multiplex graph via graph structure learning. *All the aforementioned methods either predict friendship from direct relations between users (e.g., meeting and friendship in [Qin* et al.*, 2023]) or learn implicit connections between users from complex interactions between users and check-ins.* For example, in [Yang *et al.*, 2022], in addition to friendship, users can only achieve indirect interaction through intermediate nodes (*e.g.*, POI, time nodes). *Diverging from prior studies, our MRGAN not only leverages explicit and direct relations, such as meetings, but also has the capability to unveil implicit and indirect relations like personalized preferences, modeling them explicitly. Moreover, both direct and indirect relations are fused by taking into consideration their impact and importance.*

3 Problem Definition

We use $U = \{u_1, u_2, \ldots, u_N\}$ to represent the set of all users and $P = \{p_1, p_2, \ldots, p_M\}$ to represent the set of all POIs.

Definition 1 (Check-in Record)**.** *A check-in record, denoted*

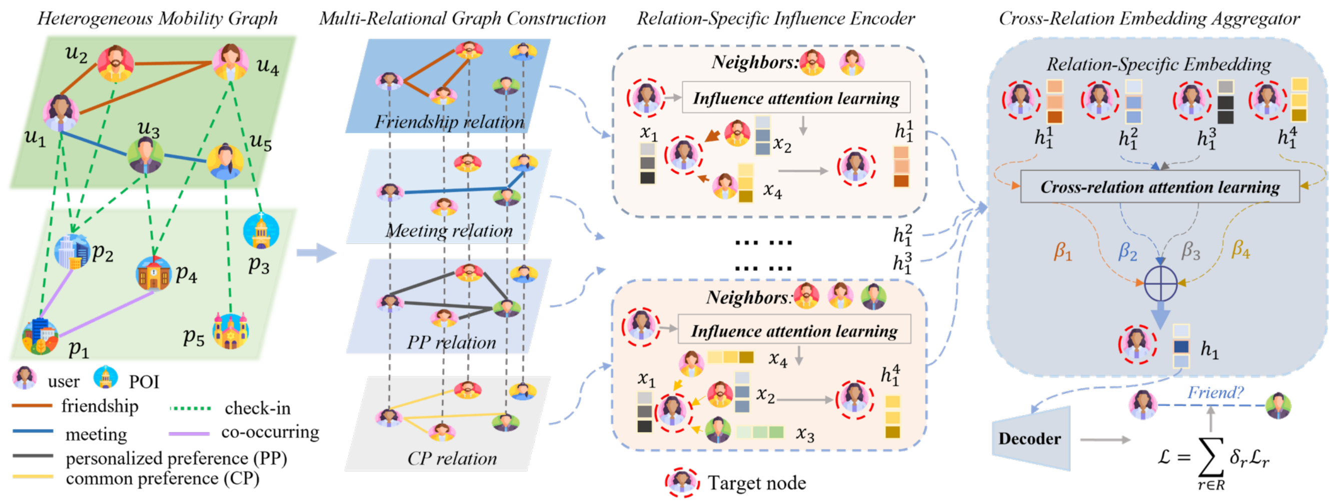

Figure 1: The overall architecture of proposed MRGAN

as $\langle u, t, p \rangle$, *signifies the visit of user* u *to point-of-interest (POI)* p *at time* t. *Here,* p *is a distinct identified venue in the form of* $\langle p_{id}, \ell \rangle$, *where* p_{id} *indicates the POI identifier and* ℓ *denotes the the geographical coordinates of the POI (i.e., longitude and latitude).*

Definition 2 (Mobility Trajectory)**.** *The mobility trajectory, represented as* Tr_u, *is a chronological-ordered sequence of check-in records generated by user* u, *denoted as* $(\langle u, t_1, p_1 \rangle, \langle u, t_2, p_2 \rangle, \ldots, \langle u, t_i, p_i \rangle)$.

Based on the aforementioned definitions, we formally outline the problem studied in this paper as follows:

Problem (Mobility Social Relationship Inference)**.** *Given a set of users* U, *our aim is to learn an inference model* $\mathcal{F}(u_i, u_j) \rightarrow \hat{y}$ *for each pair of users based on their corresponding mobility trajectories, and partial social relationship (friendship), where* $\hat{y} \in \{0, 1\}$, *and* $\hat{y} = 1$ *indicates that they are friends, otherwise, they are not.*

Typically, mobility social relationships are manifested in the spatio-temporal interactions among users and are commonly defined as *Meeting Event* in the literature [Yu *et al.*, 2018; Qin *et al.*, 2023]:

Definition 3 (Meeting Event)**.** *Users* u_i *and* u_j *are deemed to have a meeting event if they both checked in at the same location within a time threshold* τ, *i.e.,* $\exists \langle u_i, t, p \rangle \in Tr_{u_i}, \langle u_j, t', p' \rangle \in Tr_{u_j}$ *such that* $p = p'$ *and* $|t - t'| \leq \tau$.

Definition 4 (Meeting Frequency)**.** *The number of all meeting events between users* u_i *and* u_j *is defined as the meeting frequency between them, denoted by* $m_{i,j}$.

4 The Proposed Model

In this section, we elaborate our MRGAN. The overall architecture of the proposed MRGAN is shown in Figure 1, which consists of four key components: (1) *heterogeneous mobility graph*, (2) *multi-relational graph construction*, (3) *relation-specific influence encoder*, and (4) *cross-relation embedding aggregator*. First, we introduce a heterogeneous mobility graph to model user-specific mobility preference and spatio-temporal interactions among users and POIs. Second, to effectively and explicitly model direct and indirect connections (*e.g.*, shared mobility preference), we further construct the multi-relational graph which derives a multiplex graph with four types of relations from the original heterogeneous mobility graph. Third, for each relation, we design a relation-specific influence encoder with the influence attention mechanism to explore distinct social influences among different users. Finally, the cross-relation embedding aggregator based on relation-level attention is involved to aggregate user embedding across all relations. With the aggregated user embedding, we sample positive and negative instances on the multi-relational graph and design a optimization objective specially.

4.1 Heterogeneous Mobility Graph

To better predict social ties from mobility data, we need to effectively model user mobility data to capture users' movement patterns and interaction behaviors, *e.g.*, meeting events, and check-ins. Motivated by [Wu *et al.*, 2019; Li *et al.*, 2023b], we construct a heterogeneous mobility graph illustrated in Figure 1. There are two types of node, namely, user and POI, and four types of relations (edges): *friendship* (*i.e.*, social relationship), *check-in*, *meeting*, and *co-occurring*. The *friendship* relation between user pairs indicates they are friends, *i.e.*, the known friend relationships. The *check-in* relation between a user and a POI represents that the user have check-in records at the POI, and this relation is weighted by number of check-in records. The *meeting* relation among users denotes that they have meeting events, which is weighted by meeting frequency. The *co-occurring* relation denotes the link between POI pairs that co-occur frequently in user mobility trajectories.

However, the multiple types of relations we described above in the heterogeneous mobility graph could only capture the direct connectivity among nodes, which may neglect useful indirect spatio-temporal interactions. Existing inference models based on heterogeneous graphs only focus on learning user embedding from direct relations, making it challenging to uncover implicit connections between users, such co-visiting at same POI between users, which reflects shared mobility preferences.

4.2 Multi-Relational Graph Construction

To explicitly model the indirect relations that reflect users' mobility preferences while retaining the previous direct relations, we then construct a multi-relational graph based on the direct and indirect relations derived from the original heterogeneous mobility graph. Specifically, we adopt two types of direct relations, *i.e.*, *friendship* and *meeting*, and two types of indirect relations, *i.e.*, *personalized preference* and *common preference*. The direct relations (*friendship* and *meeting*) have proven to be quite effective in capturing social ties in prior studies [Yu *et al.*, 2018; Wu *et al.*, 2019; Qin *et al.*, 2023]. *Personalized preference* is a relation between users with the POI acting as the intermediate node (user-POI-user), for example, in Figure 1, user u_1 would be linked to user u_2 through the POI p_2. *Personalized preference* preserves the connections between users who share same individualized mobility habits. *Common preference* connects users with POI pairs linked indirectly by *co-occurring* relation serving as the medium (*i.e.*, user-POI-POI-user), *e.g.*, user u_1 and u_3 can be connected implicitly via the linked POI pair $< p_1, p_2 >$ or $< p_1, p_4 >$. *Common preference* is employed to explore the impact of common preference among all users on visited user pairs, because co-occurring between POIs is selected from trajectories of all users. To explicitly model the indirect relations, in multi-relational graph, we straightforwardly link the users that are connected by indirect relations in original heterogeneous mobility graph.

We denote the multi-relational graph as $\mathcal{G} = \{\mathcal{U}, \{\mathcal{E}_F, \mathcal{E}_M, \mathcal{E}_{PP}, \mathcal{E}_{CP}\}\}$, where $\mathcal{U}$ is the set of all users, $\mathcal{E}_F$, $\mathcal{E}_M$, $\mathcal{E}_{PP}$ and $\mathcal{E}_{CP}$ represent the edges of *friendship*, *meeting*, *personalized preference* and *common preference* respectively. Edges in $\mathcal{E}_{PP}$ and $\mathcal{E}_{CP}$ are also weighted by the number of indirect links between user pairs in the original graph. Notice that we use $\mathcal{G}_r$ to represent the r-th layer subgraph in the multi-relational graph.

4.3 Relation-Specific Influence Encoder

Given multi-relational graph $\mathcal{G}$, here, we aim to learn relation-specific user embedding for each relation (*i.e.*, layer). In each layer, different neighbors of user show different social influence and contribute distinctly to its embedding during learning process. The social influence expresses the social pattern of users, *e.g.*, socially active users may be assigned a higher weight, because they are more inclined to make friends. In the spirit of [Veličković *et al.*, 2017; Wang *et al.*, 2021; Tian *et al.*, 2023], we propose a relation-specific influence encoder, where we employ attention mechanism [Vaswani *et al.*, 2017] to distinguish different social impact of neighbors and aggregate information from meaningful neighbors to derive node embedding. Specifically, let $\mathcal{N}_i^r$ denote the neighbors of node u_i in the r-th layer $\mathcal{G}_r$, we fuse the neighbors of node u_i for relation type r via influence attention:

$$h_i^r = \sigma(\sum_{j \in \mathcal{N}_i^r} \alpha_{i,j}^r \cdot x_j), \tag{1}$$

where $x \in \mathbb{R}^{d \times 1}$ represents feature of node, h_i^r is the learned embedding of node u_i with relation type r, and σ is nonlinear activation. $\alpha_{i,j}^r$ is the attention value (*i.e.*, social influence) of node u_j to node u_i in relation layer $\mathcal{G}_r$, which is formulated as follows:

$$\alpha_{i,j}^r = \frac{\exp(LeakyReLU(\mathbf{a}_r^\intercal \cdot [x_i || x_j]))}{\sum_{k \in \mathcal{N}_i^r} \exp(LeakyReLU(\mathbf{a}_r^\intercal \cdot [x_i || x_k]))} \tag{2}$$

where $\mathbf{a}_r \in \mathbb{R}^{2d \times 1}$ is the influence attention vector for type r and $||$ denotes concatenate operator. To enable the training process more stable, we extend the influence attention to multi-head attention:

$$h_i^r = ||_{k=1}^K \sigma(\sum_{j \in \mathcal{N}_i^r} \alpha_{i,j}^r \cdot x_j). \tag{3}$$

With carefully preserving the social influence of users, we could obtain four groups of relation-specific node embeddings for all users, denoted as $\{\mathbf{H}^F, \mathbf{H}^M, \mathbf{H}^P, \mathbf{H}^C\}$ respectively.

4.4 Cross-Relation Embedding Aggregator

Relation-specific embeddings of user nodes only capture the single mobility interaction pattern of users. To obtain the comprehensive node embedding, we further fuse all relation-specific embedding of node with cross-relation attention. Given relation-specific embedding set $\{h_i^F, h_i^M, h_i^P, h_i^C\}$ of node u_i, we compute the weight of each relation type as follows:

$$w_r = \frac{1}{|\mathcal{U}|} \sum_{i \in \mathcal{U}} \mathbf{a}_{RA}^\intercal \cdot \tanh(\mathbf{W}_{RA} h_i^r + \mathbf{b}_{RA}), \tag{4}$$

$$\beta_r = \frac{\exp(w_r)}{\sum_{r \in \mathcal{R}} \exp(w_r)}, \tag{5}$$

where $\mathcal{U}$ is the set of all users, $\mathbf{a}_{RA}$ represents relation-level attention vector, $\mathbf{W}_{RA}$ and $\mathbf{b}_{RA}$ are learnable parameters, β_r is the importance for relation r to node u_i, and $\mathcal{R} = \{F, M, P, C\}$ denotes the set of relations. We thus fuse relation-specific embedding with weighted sum to get the final cross-relation embedding h_i of node u_i:

$$h_i = \sum_{r \in \mathcal{R}} \beta_r \cdot h_i^r. \tag{6}$$

Given relation-specific embedding, we finally get final cross-relation embedding $\mathbf{H}$, which effectively captures multiple mobility interaction patterns.

4.5 Model Optimization

In this section, we present the optimization objective of MR-GAN. Specifically, we employ binary cross entropy loss function for each relation through negative sampling to optimize

Dataset	Gowalla	Brightkite	Foursquare
City	AUSTIN	SF and LA	NEW YORK
#users	7,355	6,393	13,692
#check-ins	207,278	223,549	251,323
#POIs	5,155	20,596	25,395
#friends	35,696	20,660	21,431

Table 1: Statistics of datasets

model parameters:

$$\mathcal{L}_r = -\sum_{(i,j)\in\Theta_r} \log\sigma(<\mathbf{H}_i^\top, \mathbf{H}_j>) - \sum_{(i',j')\in\Theta_r^-} \log\sigma(-<\mathbf{H}_{i'}^\top, \mathbf{H}_{j'}>), \tag{7}$$

where $\mathcal{L}_r$ denotes the loss for the relation r, $\mathbf{H}_i$ is the embedding of user u_i. $<,>$ is a cosine similarity measure function, which acts as the decoder. Θ is the set of positive user pair instances, to learn model parameters in a more effective and stable way, we sample the positive instances via random walk [Perozzi *et al.*, 2014] on each relation layer. Θ^- is the set of negative user pairs sampled from all pairs of users who have not connected in each relation layer.

Given the multi-relational graph and the four relations $\mathcal{R}$, the final loss function can be formulated as:

$$\mathcal{L} = \sum_{r\in\mathcal{R}} \delta_r \mathcal{L}_r, \tag{8}$$

where δ_r is the hyper-parameter which controls weight of $\mathcal{L}_r$.

5 Experiments

In this section, we perform model evaluation to investigate the effectiveness of our MRGAN and baseline methods on three real-world datasets. Our experiments aim to answer the research questions as follows:

- **RQ1**: What is the performance of our MRGAN as compared to various state-of-the-art social relationship inference methods?
- **RQ2**: How do the key components contribute to the performance?
- **RQ3**: How do the key hyperparameters influence the performance of the proposed MRGAN?

5.1 Datasets

In experiments, three real-world check-in datasets, Gowalla, Brightkite [Cho *et al.*, 2011] and Foursquare [Yang *et al.*, 2019], are used to evaluate the performance of methods. Notice that, we do not discard any user mobility data as in [Qin *et al.*, 2023], while previous works [Pham *et al.*, 2016; Yang *et al.*, 2022; Huynh *et al.*, 2022] discard inactive users with few check-ins. We present the data statistics in Table 1.

5.2 Baselines

We compare our MRGAN against with three categories of baselines:

Classic graph embedding models:

- **DeepWalk** [Perozzi *et al.*, 2014] – applies random walks to gain sequences of nodes, and uses Skip-gram to learn embedding.
- **node2vec** [Grover and Leskovec, 2016] – conducts second-order random walks to embed nodes using Skip-Gram model.
- **GCN** [Kipf and Welling, 2017] – is a semi-supervised graph convolutional network that designed for the homogeneous graphs.

Heterogeneous graph embedding models:

- **R-GCN** [Schlichtkrull *et al.*, 2018] – assigns different weight matrices for the propagation on different kinds of edges.
- **HAN** [Wang *et al.*, 2019] – is a semi-supervised heterogeneous graph neural network.
- **HGMAE** [Tian *et al.*, 2023] – is a heterogeneous graph masked autoencoder.

Social relationship inference models:

- **emb-cat** [Yu *et al.*, 2018] – applies skip-gram based model on user meeting graph to learn node embeddings.
- **Heter-GCN** [Wu *et al.*, 2019] – applies GCN on user mobility heterogeneous graph for social relationship inference.
- **MVMN** [Zhang *et al.*, 2020] – learns view-specific representations and combines them to infer final social relationships.
- **LBSN2Vec++** [Yang *et al.*, 2019; Yang *et al.*, 2022] – samples hyperedges from LBSN heterogeneous hypergraph and learns node embeddings.
- **MSC-LBSN** [Huynh *et al.*, 2022] – is a multi-context embedding method designed to learn multiple representations for each user.
- **HHGNN** [Li *et al.*, 2022] – is a heterogeneous hypergraph neural network with supervised contrastive learning for the friend recommendation.
- **HMGCL** [Li *et al.*, 2023b] – captures spatio-temporal characteristics of human trajectories on a heterogeneous multigraph.
- **SRINet** [Qin *et al.*, 2023] – learns user embedding for social relationship inference by denoising mobility data.

5.3 Experiment Settings

In our experiments, we sample 25% friendships in social networks on each dataset as training set, which is used to build the social relation in heterogeneous mobility graph. Then we randomly sample another 5% friendships for validating, and use the remaining 70% as testing set.

Our MRGAN is implemented in PyTorch and optimized by Adam [Kingma and Ba, 2014] for parameter learning. In the model implementation, we set the learning rate to 0.001, and employ learning rate scheduler provided by PyTorch to adjust the learning rate based on the number of epochs. We set dimension of node representation to 128, the number of attention head to 4, the dimension of influence and cross-relation

Method	Gowalla		Brightkite		Foursquare	
	ROCAUC	PRAUC	ROCAUC	PRAUC	ROCAUC	PRAUC
DeepWalk (*KDD, 2014*)	0.6035	0.6172	0.5963	0.6130	0.6137	0.6039
node2vec (*KDD, 2016*)	0.6212	0.6159	0.6133	0.6014	0.6257	0.6211
GCN (*ICLR, 2017*)	0.6275	0.6318	0.6167	0.6312	0.6276	0.6153
R-GCN (*ESWC, 2018*)	0.6928	0.7013	0.6842	0.6913	0.6795	0.6923
HAN (*WWW, 2019*)	0.7685	0.7863	0.7512	0.7683	0.7530	0.7616
HGMAE (*AAAI, 2023*)	0.8265	0.8364	0.8106	0.8203	0.8051	0.8103
emb-cat (*Ubicomp, 2018*)	0.6405	0.6337	0.6515	0.6426	0.6359	0.6496
Heter-GCN (*IJCAI, 2019*)	0.7753	0.7867	0.7596	0.7481	0.7562	0.7633
MVMN (*TNNLS, 2020*)	0.7612	0.7831	0.7456	0.7434	0.7573	0.7648
LBSN2Vec++ (*TKDE, 2022*)	0.7521	0.7472	0.7286	0.7188	0.7352	0.7269
MSC-LBSN (*TKDE, 2022*)	0.7892	0.8054	0.7705	0.7628	0.7736	0.7721
HHGNN (*CIKM, 2022*)	0.7963	0.8016	0.7986	0.8103	0.7857	0.7969
HMGCL (*WWWJ, 2023*)	<u>0.8356</u>	0.8409	0.8163	0.8279	<u>0.8301</u>	0.8329
SRINet (*AAAI, 2023*)	0.8315	<u>0.8693</u>	<u>0.8166</u>	<u>0.8362</u>	0.8290	<u>0.8336</u>
MRGAN	**0.8901**	**0.9010**	**0.8693**	**0.8772**	**0.8543**	**0.8695**
Improvement	6.52%	3.64%	7.11%	4.90%	2.91%	4.30%

Table 2: Performance comparison of all models on three real-world datasets. *Improvement* denotes the improvement of the best results compared with second-best results. The best results are shown in **bold** and the best among baselines is <u>underlined</u>.

attention vector to 128. Guided by [Qin *et al.*, 2023], we set time threshold τ = 2 hours. We keep location co-occurrence relations with frequency $\epsilon >= 10$. In multi-relational graph, we retain top-η PP relations and top-γ CP relations by frequency for each user (Default $\eta = 4$ and $\gamma = 10$).

Notice that we can reproduce their results using the source code of SRINet [Qin *et al.*, 2023], so for a fair comparison, we directly adopt the results of SRINet and its baselines in their original paper. For other baselines, we evaluate using the source code released by their authors, and tune the parameters of baselines for optimal performance. We repeat each experiment 10 times on each dataset for all methods and report the results of methods after averaging.

5.4 Evaluation Metrics

Derived from [Yu *et al.*, 2018; Qin *et al.*, 2023], in our evaluation, we employ the widely used Area Under the ROC Curve (ROCAUC) and the Area Under the Precision-Recall Curve (PRAUC) to evaluate the performance of various methods. In the context of social relationship inference, we calculate the pairwise cosine similarity of learned user representations in each method to derive a score indicating the likelihood of two users being friends.

5.5 Experiment Results

Overall Performance (RQ1)

As the experimental result shown in Table 2, our MRGAN significantly outperforms all baselines across three datasets in terms of all evaluation metrics. This observation confirms the the superiority of our proposed MRGAN. This may be because our model not only considers the social connections between users reflected in direct relations, but also captures users' mobility preferences implied by indirect relations, which can be used to learn more interest correlations between users, further improving accuracy in social relationship inference. Moreover, benefiting from the influence attention, the social influence from various users can be comprehensively taken into account, besides, with relation-level attention, MRGAN could preserve semantic information of user's different mobility patterns. Specifically, results in Table 2 indicate that our MRGAN achieves average gains of 5.51% ROCAUC and 4.28% PRAUC in comparison to the best-performed baseline across all datasets.

SRINet achieves the second-best performance in most cases, which improves the performance by handling mobility noise data issue. Despite effectiveness, SRINet overlooks the user mobility preference, a crucial aspect in friend inference. Heter-GCN and HMGCL, another two inference models based on heterogeneous graph, take the user mobility pattern into account, but hardly capture the mobility preference by singularly applying embedding learning on direct relations. Although MSC-LBSN, LBSN2Vec++ and HHGNN apply sophisticated embedding learning method on heterogeneous hypergraph, they suffer from the noise issue. In addition, HAN and HGMAE, two heterogeneous graph representation learning models, achieve competitive results even compared with inference models. The possible reason is HAN and HGMAE are both powered by hierarchical attention, which suggests that effectively accounting for the influence of diverse users and various mobility semantics is crucial for friendship inference.

Ablation Study (RQ2)

To analysis the effectiveness of the key components in our model, we conduct the following ablation studies. Two major categories of variants are designed: relation-induced and attention-induced variants.

Relation-induced variants are employed to verify the con-

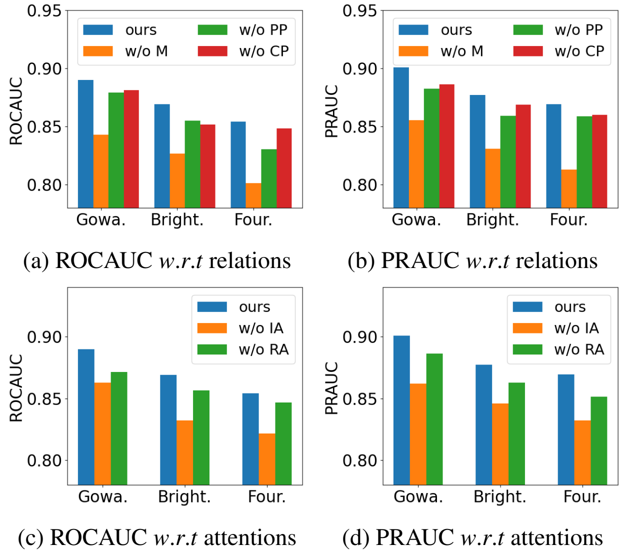

(a) ROCAUC *w.r.t* relations (b) PRAUC *w.r.t* relations

(c) ROCAUC *w.r.t* attentions (d) PRAUC *w.r.t* attentions

Figure 2: The comparison of MRGAN and its variants

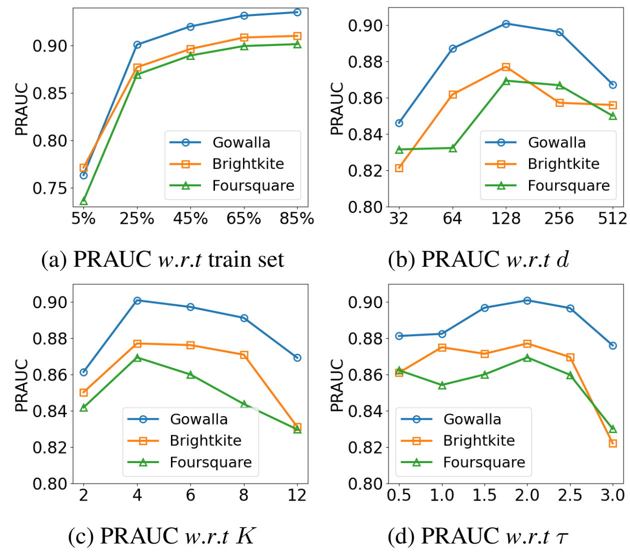

(a) PRAUC *w.r.t* train set (b) PRAUC *w.r.t* d

(c) PRAUC *w.r.t* K (d) PRAUC *w.r.t* τ

Figure 3: Hyper-parameter impact of training set ratio, embedding dimension d, attention head K, time threshold τ (hour)

tribution of different relations to model performance, which consists of:

- ***w/o M*** removes meeting relation and corresponding loss.
- ***w/o PP*** deletes personalized preference relation and corresponding loss.
- ***w/o CP*** excludes common preference relation and corresponding loss.

Attention-induced variants are designed to verify the contribution of attention mechanism :

- ***w/o IA*** removes the node-level attention (*i.e.*, influence attention).
- ***w/o RA*** takes away the relation-level attention in cross-relation embedding aggregator.

The results of ablation studies are summarized in Figure 2. From the results, three conclusions are made as follows: (1) Five variants are consistently worse than MRGAN across all datasets in terms of ROCAUC and PRAUC, indicating the effectiveness and necessity of the five key components. (2) With regard to relation-induced variants, the performance of ***w/o M*** experiences a significant decline compared to that of MRGAN, which demonstrates that the meeting relation is powerful for capturing possible social links. ***w/o PP*** and ***w/o CP*** perform worse than MRGAN in all cases, which suggests the employment of PP and CP relations reveal the shared mobility preference and common mobility pattern of users. (3) Regarding attention-induced variants, the performance of ***w/o IA*** drops significantly compared to MRGAN, which means that considering and learning the importance of social influence among users is effective for relationship inference. The gap between MRGAN and ***w/o RA*** shows that aggregation of different interaction patterns is critical.

Hyper-parameter Analysis (RQ3)

We investigate the sensitivity of four main parameters: training set ratio, user embedding dimension d, the number of attention head K and time threshold τ.

As shown in Figure 3a, as the training set ratio increases from 5% to 25%, a substantial improvement in performance is observed. Subsequently, with further increments in the training set, performance exhibits a gradual and incremental increase. The observed result aligns with our expectation, and our model can attain the desired performance with a training set ratio of 25%. In Figure 3b, we observe that model performance initially improves with increasing dimensionlity, peaking at 128. Subsequently, as dimensions continue to increase, performance degrades. Clearly, too small dimensions poses challenges in adequately expressing the comprehensive information captured by our MRGAN. Figure 3c illustrates the performance *w.r.t.* attention head K in influence and cross-relation attentions. We can see that model performance reaches its highest when K is 4. As illustrated in Figure 3d, the model performance reaches its peak at $\tau = 2$. It is noteworthy that before $\tau = 2$, the performance shows a gradual increase or intermittent fluctuations. We interpret this phenomenon as a trade-off between noise and interactive informativeness. The smaller τ is, the less noise is introduced, but a large amount of interactive information is filtered out. On the other hand, as τ increases, the noise increases, but so does the amount of interactive information. However, when τ is too large, noise takes precedence, resulting in a significant performance degradation, *e.g.*, $\tau = 3$.

6 Conclusion

In this paper, we present a multi-relational graph attention network for inferring social relationships from user mobility data. Our proposed model effectively learns user representations for friendship inference by explicitly modeling users' mobility preference and effectively capture the different impact and importance of both direct and indirect relations. Experiments on three real-world mobility datasets show that our model significantly outperforms state-of-the-art baselines.

Acknowledgments

This work is partially supported by the National Natural Science Foundation of China under grant Nos. 62176243, 62372421, and 41927805, the Project funded by China Postdoctoral Science Foundation (No. 2023M743328), the Postdoctoral Fellowship Program of CPSF (No. GZC20232500), and the Postdoctoral Project funded by Qingdao (No. QDBSH20240102093). Guangming Qin and Jianpeng Qi contribute equally to this work.

References

[Antelmi *et al.*, 2024] Alessia Antelmi, Gennaro Cordasco, Mirko Polato, Vittorio Scarano, Carmine Spagnuolo, and Dingqi Yang. A survey on hypergraph representation learning. *ACM Comput. Surv.*, 56(1):24:1–24:38, 2024.

[Backes *et al.*, 2017] Michael Backes, Mathias Humbert, Jun Pang, and Yang Zhang. walk2friends: Inferring social links from mobility profiles. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 1943–1957, 2017.

[Cai *et al.*, 2023] Xuheng Cai, Lianghao Xia, Xubin Ren, and Chao Huang. How expressive are graph neural networks in recommendation? In Ingo Frommholz, Frank Hopfgartner, Mark Lee, Michael Oakes, Mounia Lalmas, Min Zhang, and Rodrygo L. T. Santos, editors, *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, CIKM 2023, Birmingham, United Kingdom, October 21-25, 2023*, pages 173–182. ACM, 2023.

[Cho *et al.*, 2011] Eunjoon Cho, Seth A Myers, and Jure Leskovec. Friendship and mobility: user movement in location-based social networks. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1082–1090, 2011.

[Cui *et al.*, 2018] Peng Cui, Xiao Wang, Jian Pei, and Wenwu Zhu. A survey on network embedding. *IEEE transactions on knowledge and data engineering*, 31(5):833–852, 2018.

[Dai *et al.*, 2022] Shaojie Dai, Yanwei Yu, Hao Fan, and Junyu Dong. Spatio-temporal representation learning with social tie for personalized poi recommendation. *Data Science and Engineering*, 7(1):44–56, 2022.

[Dai *et al.*, 2023] Shaojie Dai, Jinshuai Wang, Chao Huang, Yanwei Yu, and Junyu Dong. Dynamic multi-view graph neural networks for citywide traffic inference. *ACM Trans. Knowl. Discov. Data*, 17(4):53:1–53:22, 2023.

[Eagle *et al.*, 2009] Nathan Eagle, Alex Sandy Pentland, and David Lazer. Inferring friendship network structure by using mobile phone data. *PNAS*, 106(36):15274–15278, 2009.

[Grover and Leskovec, 2016] Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 855–864. ACM, 2016.

[Huynh *et al.*, 2022] Thanh Trung Huynh, Vinh Van Tong, Thanh Tam Nguyen, Jun Jo, Hongzhi Yin, and Quoc Viet Hung Nguyen. Learning holistic interactions in lbsns with high-order, dynamic, and multi-role contexts. *IEEE Transactions on Knowledge and Data Engineering*, 2022.

[Kingma and Ba, 2014] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[Kipf and Welling, 2017] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017.

[Li *et al.*, 2011] Zhenhui Li, Cindy Xide Lin, Bolin Ding, and Jiawei Han. Mining significant time intervals for relationship detection. In *International Symposium on Spatial and Temporal Databases*, pages 386–403. Springer, 2011.

[Li *et al.*, 2022] Yongkang Li, Zipei Fan, Jixiao Zhang, Dengheng Shi, Tianqi Xu, Du Yin, Jinliang Deng, and Xuan Song. Heterogeneous hypergraph neural network for friend recommendation with human mobility. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pages 4209–4213, 2022.

[Li *et al.*, 2023a] Shuzhe Li, Wei Chen, Bingqi Yan, Zhen Li, Shunzhi Zhu, and Yanwei Yu. Self-supervised contrastive representation learning for large-scale trajectories. *Future Gener. Comput. Syst.*, 148:357–366, 2023.

[Li *et al.*, 2023b] Yongkang Li, Zipei Fan, Du Yin, Renhe Jiang, Jinliang Deng, and Xuan Song. Hmgcl: Heterogeneous multigraph contrastive learning for lbsn friend recommendation. *World Wide Web*, 26(4):1625–1648, 2023.

[Li *et al.*, 2024] Shuzhe Li, Wei Chen, Bin Wang, Chao Huang, Yanwei Yu, and Junyu Dong. Mcn4rec: Multi-level collaborative neural network for next location recommendation. *ACM Transactions on Information Systems*, 42(4):1–26, 2024.

[Liu *et al.*, 2021] Xiao Liu, Fanjin Zhang, Zhenyu Hou, Li Mian, Zhaoyu Wang, Jing Zhang, and Jie Tang. Self-supervised learning: Generative or contrastive. *IEEE transactions on knowledge and data engineering*, 35(1):857–876, 2021.

[Perozzi *et al.*, 2014] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 701–710. ACM, 2014.

[Pham *et al.*, 2013] Huy Pham, Cyrus Shahabi, and Yan Liu. Ebm: an entropy-based model to infer social strength from spatiotemporal data. In *Proceedings of the 2013 ACM SIGMOD*, pages 265–276. ACM, 2013.

[Pham *et al.*, 2016] Huy Pham, Cyrus Shahabi, and Yan Liu. Inferring social strength from spatiotemporal data. *ACM Transactions on Database Systems (TODS)*, 41(1):7, 2016.

[Qin *et al.*, 2023] Guangming Qin, Lexue Song, Yanwei Yu, Chao Huang, Wenzhe Jia, Yuan Cao, and Junyu Dong.

Graph structure learning on user mobility data for social relationship inference. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 4578–4586, 2023.

[Schlichtkrull *et al.*, 2018] Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. Modeling relational data with graph convolutional networks. In *The Semantic Web: 15th International Conference, ESWC 2018, Heraklion, Crete, Greece, June 3–7, 2018, Proceedings 15*, pages 593–607. Springer, 2018.

[Sharma *et al.*, 2021] Anshuman Sharma, Yogesh K Dwivedi, Vikas Arya, and Muhammad Qutubuddin Siddiqui. Does sms advertising still have relevance to increase consumer purchase intention? a hybrid pls-sem-neural network modelling approach. *Computers in Human Behavior*, 124:106919, 2021.

[Sun *et al.*, 2019] Fan-Yun Sun, Jordan Hoffmann, Vikas Verma, and Jian Tang. Infograph: Unsupervised and semi-supervised graph-level representation learning via mutual information maximization. *arXiv preprint arXiv:1908.01000*, 2019.

[Tian *et al.*, 2023] Yijun Tian, Kaiwen Dong, Chunhui Zhang, Chuxu Zhang, and Nitesh V Chawla. Heterogeneous graph masked autoencoders. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 9997–10005, 2023.

[Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[Veličković *et al.*, 2017] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.

[Wang *et al.*, 2014] Hongjian Wang, Zhenhui Li, and Wang-Chien Lee. Pgt: Measuring mobility relationship using personal, global and temporal factors. In *Proceedings of 2014 IEEE International Conference on Data Mining (ICDM)*, pages 570–579. IEEE, 2014.

[Wang *et al.*, 2019] Xiao Wang, Houye Ji, Chuan Shi, Bai Wang, Yanfang Ye, Peng Cui, and Philip S Yu. Heterogeneous graph attention network. In *The world wide web conference*, pages 2022–2032, 2019.

[Wang *et al.*, 2021] Xiao Wang, Nian Liu, Hui Han, and Chuan Shi. Self-supervised heterogeneous graph neural network with co-contrastive learning. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, pages 1726–1736, 2021.

[Wang *et al.*, 2023] Xiao Wang, Deyu Bo, Chuan Shi, Shaohua Fan, Yanfang Ye, and Philip S. Yu. A survey on heterogeneous graph embedding: Methods, techniques, applications and sources. *IEEE Trans. Big Data*, 9(2):415–436, 2023.

[Wu *et al.*, 2019] Yongji Wu, Defu Lian, Shuowei Jin, and Enhong Chen. Graph convolutional networks on user mobility heterogeneous graphs for social relationship inference. In *IJCAI*, pages 3898–3904, 2019.

[Wu *et al.*, 2023a] Junzhuang Wu, Yujing Zhang, Yuhua Li, Yixiong Zou, Ruixuan Li, and Zhenyu Zhang. Sstp: Social and spatial-temporal aware next point-of-interest recommendation. *Data Science and Engineering*, 8(4):329–343, 2023.

[Wu *et al.*, 2023b] Shiwen Wu, Fei Sun, Wentao Zhang, Xu Xie, and Bin Cui. Graph neural networks in recommender systems: A survey. *ACM Comput. Surv.*, 55(5):97:1–97:37, 2023.

[Xia *et al.*, 2022] Lianghao Xia, Chao Huang, Yong Xu, Jiashu Zhao, Dawei Yin, and Jimmy Huang. Hypergraph contrastive collaborative filtering. In *Proceedings of the 45th International ACM SIGIR conference on research and development in information retrieval*, pages 70–79, 2022.

[Yang *et al.*, 2019] Dingqi Yang, Bingqing Qu, Jie Yang, and Philippe Cudre-Mauroux. Revisiting user mobility and social relationships in lbsns: a hypergraph embedding approach. In *The world wide web conference*, pages 2147–2157, 2019.

[Yang *et al.*, 2022] Dingqi Yang, Bingqing Qu, Jie Yang, and Philippe Cudré-Mauroux. Lbsn2vec++: Heterogeneous hypergraph embedding for location-based social networks. *IEEE Transactions on Knowledge and Data Engineering*, 34(4):1843–1855, 2022.

[Yu *et al.*, 2018] Yanwei Yu, Hongjian Wang, and Zhenhui Li. Inferring mobility relationship via graph embedding. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(3):1–21, 2018.

[Yu *et al.*, 2020] Junliang Yu, Hongzhi Yin, Jundong Li, Min Gao, Zi Huang, and Lizhen Cui. Enhancing social recommendation with adversarial graph convolutional networks. *IEEE Transactions on knowledge and data engineering*, 34(8):3727–3739, 2020.

[Yu *et al.*, 2021] Junliang Yu, Hongzhi Yin, Min Gao, Xin Xia, Xiangliang Zhang, and Nguyen Quoc Viet Hung. Socially-aware self-supervised tri-training for recommendation. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, pages 2084–2092, 2021.

[Zhang *et al.*, 2020] Wei Zhang, Xin Lai, and Jianyong Wang. Social link inference via multiview matching network from spatiotemporal trajectories. *IEEE transactions on neural networks and learning systems*, 2020.

[Zhou *et al.*, 2018] Fan Zhou, Bangying Wu, Yi Yang, Goce Trajcevski, Kunpeng Zhang, and Ting Zhong. Vec2link: Unifying heterogeneous data for social link prediction. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 1843–1846, 2018.