

Mining Unfollow Behavior in Large-Scale Online Social Networks via Spatial-Temporal Interaction

Haozhe Wu,^{1*} Zhiyuan Hu,^{1*} Jia Jia,^{1†} Yaohua Bu,² Xiangnan He,³ Tat-Seng Chua,⁴

¹Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China

Beijing National Research Center for Information Science and Technology (BNRist)

The Institute for Artificial Intelligence, Tsinghua University

²Academy of Arts & Design, Tsinghua University, Beijing, China

³School of Information Science and Technology, University of Science and Technology of China

⁴School of Computing, National University of Singapore

wuhz19@mails.tsinghua.edu.cn, zy-hu16@mails.tsinghua.edu.cn, jjia@.tsinghua.edu.cn

Abstract

Online Social Networks (OSNs) evolve through two pervasive behaviors: *follow* and *unfollow*, which respectively signify relationship creation and relationship dissolution. Researches on social network evolution mainly focus on the follow behavior, while the unfollow behavior has largely been ignored. Mining unfollow behavior is challenging because user's decision on unfollow is not only affected by the simple combination of user's attributes like informativeness and reciprocity, but also affected by the complex interaction among them. Meanwhile, prior datasets seldom contain sufficient records for inferring such complex interaction. To address these issues, we first construct a large-scale real-world Weibo¹ dataset, which records detailed post content and relationship dynamics of 1.8 million Chinese users. Next, we define user's attributes as two categories: spatial attributes (*e.g.*, social role of user) and temporal attributes (*e.g.*, post content of user). Leveraging the constructed dataset, we systematically study how the interaction effects between user's spatial and temporal attributes contribute to the unfollow behavior. Afterwards, we propose a novel unified model with heterogeneous information (UMHI) for unfollow prediction. Specifically, our UMHI model: 1) captures user's spatial attributes through social network structure; 2) infers user's temporal attributes through user-posted content and unfollow history; and 3) models the interaction between spatial and temporal attributes by the nonlinear MLP layers. Comprehensive evaluations on the constructed dataset demonstrate that the proposed UMHI model outperforms baseline methods by 16.44% on average in terms of precision. In addition, factor analyses verify that both spatial attributes and temporal attributes are essential for mining unfollow behavior.

1 Introduction

The popularization of the Internet greatly facilitates the development of Online Social Networks (OSNs). Statistics² show that more than 3 billion people around the world now use OSNs each month. The fast development of OSNs is

* Authors contributed equally to this work

† Corresponding author: J. Jia (jjia@tsinghua.edu.cn)

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<https://www.weibo.com/>

²<https://digitalreport.wearesocial.com/>



Figure 1: An example of unfollow behavior in online social network: User1 posted too much spam messages, resulting the unfollow behaviors of his/her followers.

tightly coupled with the evolution of online social relationships. There are two basic actions for users to manage their social relationships: *follow* (relationship creation) and *unfollow* (relationship dissolution), of which an example is shown in Figure 1. Previous research efforts paid much attention to the follow behavior (Liben-Nowell and Kleinberg 2007; Bliss *et al.* 2014; Quercia *et al.* 2012a), while unfollow behavior mining has largely been ignored. Statistics in a real-world Weibo dataset show that almost 40% of users unfollow others at least once a month. The frequent occurrence of unfollow behavior leads to an interesting question: why people unfollow others? Taking a step forward, can we predict the unfollow behavior in OSNs?

Previous research efforts have shown the rationality of analyzing and predicting unfollow behavior. They mainly focus on mining unfollow behavior through user-posted content and social network structures, and have revealed several attributes that are closely related to the unfollow behavior, such as the informativeness of the followee and the reciprocity of relationships. Leveraging the found attributes, these methods predict unfollow behavior by defining handcrafted features (Maity *et al.* 2018; Kwak *et al.* 2012; Xu *et al.* 2013). However, merely defining handcrafted features from unfollow-related attributes cannot generalize well to large-scale setting, because user's decision on unfollow is not only affected by the simple combination of user's attributes, but also by the complex interaction effects among them. Take the community and post content of users as an example, user's personal habits shown in post content would influence the community of user, while the community would also affect user's preference on post content;

thus the intertwined nature of community and post content would result in a highly complicated unfollow decision mechanism. Meanwhile, modeling such interaction effects needs abundant records of online social networks, while prior datasets seldom satisfy such requirement.

To address these issues, in this work, we first construct a large-scale benchmark dataset on Sina Weibo, which contains 1.8 million Chinese users, 400 million social relationships and 10 million records of unfollow. We record the timeline, content, and upvotes of each user’s microblogs and track the unfollow actions of these users in a month. Then, inspired by previous researches, we define user’s attributes as two categories: spatial attributes (*e.g.*, social role of user) and temporal attributes (*e.g.*, post content of user). Based on the constructed dataset, we systematically study how the interaction between spatial and temporal attributes contribute to the unfollow behavior and conduct exhaustive data observations. Next, for the unfollow prediction task, we propose a novel unified model with heterogeneous information (UMHI) to learn the highly complex interaction. The main idea of UMHI is to model the user’s spatial attributes through social network structure and user’s temporal attributes through user-posted content and short-term unfollow history (a user’s unfollowed-people list). First, information of social network structure is extracted by network embedding; Second, we adopt hierarchical attention network (HAN) (Yang *et al.* 2016b) to learn representations from user-posted content. Third, the matrix factorization (MF) based collaborative filtering (Koren *et al.* 2009) is employed to reduce user’s short-term unfollow history into low dimensional feature vectors. Finally, a unified heterogeneous information fusion network is trained to model the interaction between spatial and temporal attributes. Figure 2 summarizes the workflow of our framework.

Experiments demonstrate that our model outperforms the baseline methods by 16.44% on average in terms of precision. In addition, factor analyses show that both spatial and temporal attributes are essential for mining unfollow behavior. To conclude, we summarize our contributions as follows:

- We construct a real-world benchmark dataset on Sina Weibo with 1.8 million Chinese users and 400 million social relationships. It records user’s post content and relationship dynamics for a whole month. Such large-scale dataset is not only useful for unfollow prediction, but also beneficial for further research like depression detection and rumor detection. Dataset is publicly available at <https://github.com/wuhaozhe/Unfollow-Prediction>.
- We systematically study how the spatial and temporal attributes contribute to the unfollow behavior, unveiling the interaction effects between these two categories of attributes.
- We propose a novel UMHI model, which predicts unfollow behavior by learning spatial and temporal attributes through user’s footprint on OSN. The proposed method outperforms baseline methods by a large margin.

2 Related Work

Previous researches on social network evolution mainly focus on the follow behavior (link prediction and friend recommendation), while the unfollow behavior has received less scrutiny.

For mining follow actions, several researches focus on predicting relationship from social network structures (Tang *et al.* 2015; Perozzi *et al.* 2014; Grover and Leskovec 2016). Perozzi *et al.* proposed Deepwalk algorithm (Perozzi *et al.* 2014) to represent the network structure as low dimensional embeddings through random walks on social networks, then the similarity between embeddings reflects the possibility of establishing relations. The node2vec algorithm (Grover and Leskovec 2016) extends the depth first random walk strategy into a biased random walk procedure, making the learned representations scalable. The LINE algorithm (Tang *et al.* 2015) embeds node into low dimensional representations by optimizing the first order and second order proximity. All of these approaches only capture spatial factors, while unfollow behavior is caused by intricate interactions of spatial and temporal factors. Therefore the aforementioned algorithms all suffer from inferior performance in terms of unfollow prediction.

Compared with these link prediction methods, researches on unfollow behavior mostly resort to rule based methods. Kwak *et al.* firstly researched on the unfollow dynamics in Twitter, they found some unfollow factors, including informativeness, reciprocity and relationship stabilization (Kwak *et al.* 2011). Later, they built a logistic regression model based on structure properties and behavioral properties (Kwak *et al.* 2012). The same group then adopted actor-oriented model (SIENA) to examine the impacts of reciprocity, status, embeddedness, homophily, and informativeness on tie dissolution (Xu *et al.* 2013). Kivran-Swaine *et al.* explored how network structures alone influence unfollow behavior (Kivran-Swaine *et al.* 2011). Quercia *et al.* researched on whether user’s demographics such as age, gender will influence unfollow behavior in facebook (Quercia *et al.* 2012b). Maity *et al.* analyzed the content of the posts made by the Twitter users who lose followers consistently and extracted various behavioral features from followee’s post content to make prediction. (Maity *et al.* 2018). However, these rule-based methods can hardly represent the interaction between the spatial and temporal attributes, therefore can’t generalize well to large-scale settings.

3 Problem Formulation

The problem setting of unfollow prediction is to predict a user’s future unfollow behavior from raw online social networks (OSNs) data. OSNs contain several attributes that are predictive for unfollow behavior, we define them as two categories: spatial attributes and temporal attributes. Formally:

- **Spatial Attributes:** user’s attributes that remain unchanged in time interval $[t_{start}, t_{end}]$ (*e.g.*, user’s social role).
- **Temporal Attributes:** user’s attributes that dynamically change in time interval $[t_{start}, t_{end}]$ (*e.g.*, user-posted content).

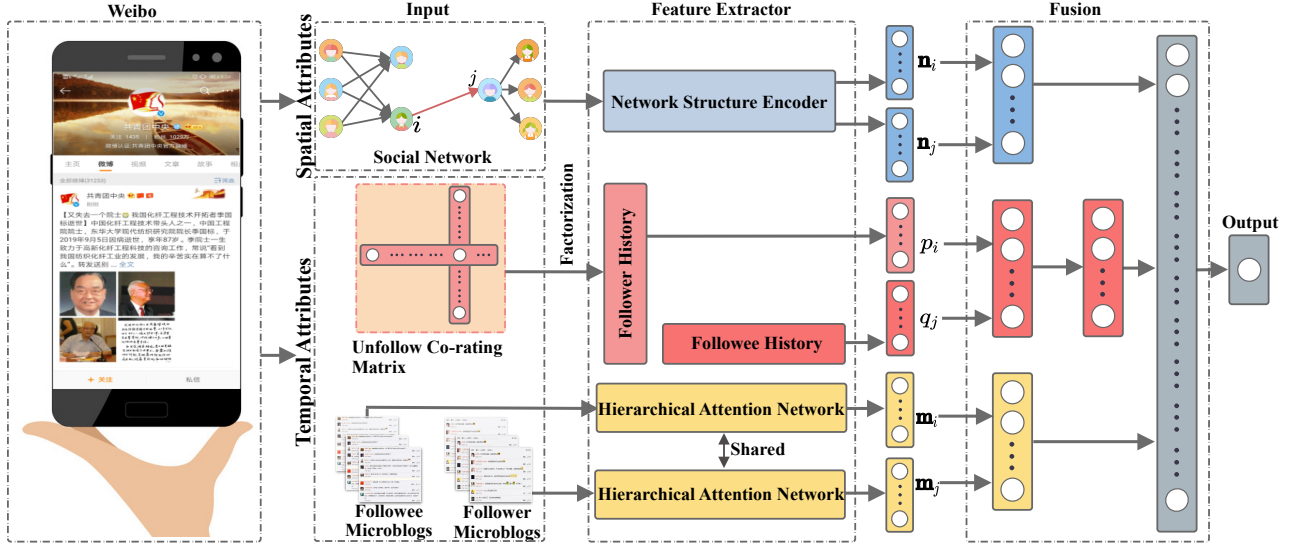


Figure 2: The workflow of UHMI: We predict unfollow behavior by fusing two kinds of information: (1) Spatial attributes, here we utilize user-located network structure to represent such attributes. (2) Temporal attributes, we incorporate user’s temporal post content and unfollow history as the input to represent such attributes. The interaction effects between spatial and temporal attributes is modeled through the fusion stage in UMHI. Details of hierarchical attention network is shown in Figure 4.

Given a set of users V , we use a binary matrix $R \in \mathbb{R}^{|V| \times |V|}$ to denote the link dynamics between users in time interval $[t_{start}, t_{end}]$. Specifically, each entry r_{ij} denotes whether user i has unfollowed user j in $[t_{start}, t_{end}]$, which is defined as:

$$r_{ij} = \begin{cases} 1, & \text{if user } i \text{ unfollowed user } j \text{ in } [t_{start}, t_{end}]; \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

To conduct training-test scheme, given test edges E_{test} , we mask E_{test} from binary matrix R . Specifically, for any $r_{ij} \in E_{test}$, we enforce $r_{ij} = 0$ to get the training matrix $R_{train} = R \setminus E_{test}$, here we call matrix R_{train} to be the unfollow history matrix, r_i to be the unfollow history of user i , and r_j to be the unfollowed history of user j .

Problem. The target of unfollow prediction is to predict $r_{ij} \in E_{test}$. We incorporate social network structure information $\mathbf{n}_i, \mathbf{n}_j$ as spatial attributes, incorporate posted content $\mathbf{m}_i, \mathbf{m}_j$ and unfollow history r_i, r_j as temporal attributes. Then, our objective is to learn a function $y_{ij} = f(\mathbf{n}_i, \mathbf{n}_j, \mathbf{m}_i, \mathbf{m}_j, r_i, r_j)$, which estimates the probability that user i would unfollow user j .

4 Dataset Observation

In this section, we conduct exhaustive data observations to analyze how spatial attributes and temporal attributes interact with each other. In order to visualize the observation results, we define several statistics that respectively represent spatial attributes and temporal attributes. We take user’s social role as an example of spatial attributes. Specifically, following the prior work of Yang (Yang *et al.* 2016a), we divide users into three groups: 5% of users with the highest PageRank score (Page *et al.* 1999) are considered to be the

opinion leader (OpnLdr); 5% of users with the lowest Burts Constrains score (Burt 2017) are considered to be the bridges between disconnected communities in social network, *a.k.a.*, the structure hole (StrHole); and the rest of users are considered to be ordinary users (OrdUsr). Then for the temporal attributes, we define the following two attributes:

- Similarity: the tf-idf similarity between follower and followees’ post content within a month.
- Exposure: the number of microblogs posted by the follower within a month.

Before analyzing the results of data observations, we first elaborate the details of dataset construction. We build a large-scale benchmark dataset on Sina Weibo, which contains 1.8 million Chinese users and 400 million social relationships. We record each user’s microblog post content and relationship dynamics from September 28, 2012 to October 29, 2012. Each post is recorded with its time, content and upvotes. During the month we observe, we found that 10,705,319 (2.53%) edges have been broken at least once, and 714,945 (40.00%) users have unfollowed others at least once, verifying that unfollow is a pervasive behavior.

Although the unfollow behavior is pervasive, the ratio between unfollow relationships and hold relationships is still unbalanced (2.53% in our dataset). Therefore, we build a balanced sub-dataset E_{test} for fair data observation and further training-test scheme. Table 1 shows the composition of E_{test} , here we filter out the edges in which either follower or followee post no microblog content.

To measure how the interaction effects between spatial and temporal attributes affect user’s decision on unfollow, we define metric $rou(c)$ (ratio of unfollow under condition c) to be the ratio between the number of unfollow edges and

Table 1: Sum of edges in the E_{test} among different social roles and relations statuses. OrdUsr, OpnLdr and StrHole are the shorthands of ordinary user, opinion leader and structure hole, meaning that the followee of the edge is OrdUsr/OpnLdr/StrHole.

Social role	OrdUsr	OpnLdr	StrHole	Sum
hold	1887	2551	1364	5802
unfollow	2391	2539	1860	6790
Sum	4287	5090	3224	12592

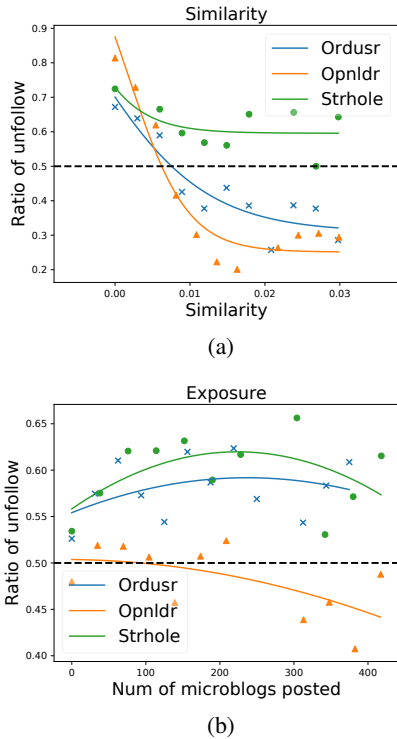


Figure 3: The results of exploratory analysis, (a) (b) respectively plots the distribution of rou under the interaction between similarity / exposure and social role.

the number of all edges. Formally:

$$rou(c) = \frac{N_{un}(c)}{N_{ho}(c) + N_{un}(c)}, \quad (2)$$

where $N_{un}(c)$ is the number of edges that break relationships under condition c within the observed month, $N_{ho}(c)$ is the number of edges that hold relationships. Edges with higher $rou(c)$ are more likely to break, while with lower $rou(c)$ are more stable. In the following analysis, we leverage E_{test} to estimate the distribution of $rou(c)$ under different combinations of user’s spatial attributes and temporal attributes.

Interaction between similarity and social role. Similarity reveals user’s homogeneity in short-term post content. Two users with higher homogeneity are more likely to have

stable relationship. Figure 3a plots the relationship between similarity and rou under different social roles. We observe that as similarity increases, the rou quickly decreases to the saturation point. Figure 3a also shows that the interaction effects between social role and similarity is significant: opinion leader is most sensitive to similarity, since it varies the most as similarity increases. On the other hand, for structure holes, similarity seems to be a weak factor.

Interaction between exposure and social role. Exposure evaluates the activeness of followee in recent time. From Figure 3b, we observe that with the increase of exposure, rou shows different tendency under different social roles. The rou of opinion leaders is naturally lower than that of ordinary users and structure holes. Meanwhile, the rou of ordinary users and structure holes fluctuates with the increase of exposure, while the rou of opinion leaders monotonically decreases. Such phenomenon has a simple interpretation: followers are concerned with the messages from opinion leaders, and would not curtly break up relationships with them.

Summary. We reveal that the spatial and temporal attributes interplay with each other and result in an intricate mechanism of unfollow behavior. Therefore, for unfollow prediction, it is necessary to employ the nonlinearity of neural network to model such interactions. In next section we will introduce how the UMHI model leverages neural network to extract discriminative features for unfollow prediction.

5 Proposed Model

Based on the exploratory analysis in Section 4, we propose a novel UMHI model which incorporates heterogeneous information to predict the unfollow behavior. Our model simultaneously takes the spatial attributes and temporal attributes as input. More specifically, (1) we capture spatial attributes through social network structure, and utilize network embedding to compress the graph structured data into feature vectors. (2) We infer user’s temporal attributes from user-posted content and unfollow history, the hierarchical attention network (HAN) and matrix factorization (MF) are respectively leveraged to learn the feature vectors from post content and unfollow history. (3) We employ the nonlinearity of MLP layers to model the interaction effects between spatial attributes and temporal attributes. The overall architecture of UMHI is presented in Figure 2.

The remainder of this section is organized to elaborate each component of UMHI framework.

5.1 Network Structure Encodes

The user-located social network structure is closely related to the unfollow behavior. As revealed in Section 4, users with higher pagerank score have lower probability to be unfollowed. Except for the social role which affects followee’s decision on unfollow, prior work by Quericia *et al.* (2012b) argues that common friends between two users also greatly affects the stability of relationship.

To comprehensively encode network structure into low dimensional feature vectors, we leverage the network embedding method. Different network embedding methods have

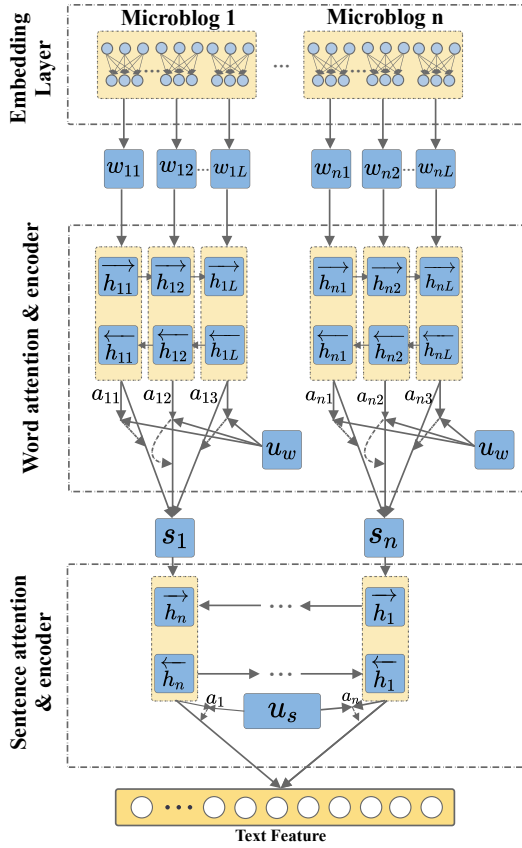


Figure 4: An illustration of hierarchical attention network. The first layer of HAN incorporates the word embeddings of each post as input, then outputs a dimension fixed post content representation, the second layer of HAN incorporates post content representations of all posts as input, and aggregates representations into a dimension fixed feature vector.

different capacities for encoding different kinds of structure information (Dalmia *et al.* 2018). Because of different optimization strategies, LINE algorithm has better capacity on capturing local information, while Deepwalk and Node2vec prefer to encode global information.

For unfollow prediction, experiment results in Section 6 demonstrate that local information plays a more important role than global. Therefore, we employ the LINE algorithm to represent social network structure. LINE has two different objectives: first-order proximity and second-order proximity. The first-order proximity refers to the local pairwise proximity between vertices in the network, while the second-order proximity is the similarity in the neighborhood network structures between two nodes. We denote LINE with first-order proximity as LINE1, and LINE with second-order proximity as LINE2. The UMHI framework incorporates both LINE1 and LINE2 as input, here we respectively denote the node embeddings of follower i and followee j as \mathbf{n}_i and \mathbf{n}_j .

5.2 Post Content Encodes

Users' post content reflects their temporal preferences. By observing the post content of follower and followee, we may guess whether the follower loses interest in the followee's posts at present. Since that both the words in each post and all posts of one user are serialized sequences, we employ the hierarchical attention network (HAN) to extract the discriminative features. The hierarchical attention network, as shown in Figure 4, contains two bidirectional attention LSTM layers. The first LSTM layer encodes each post into dimension fixed representation, the second layer aggregates representations from all posts into a dimension fixed feature vector. Finally the attention mechanism (Bahdanau *et al.* 2014) assign different weights to different words and posts.

More specifically, we firstly embed each word of post content into word embeddings $w_{ij} \in \mathbb{R}^N$ (i is microblog ID, j is word ID). For a microblog with T words, the bidirectional attention LSTM f_1 takes $w_{i1} \dots w_{iT}$ as input, and outputs the sentence representation s_i , formally as:

$$s_i = f_1(w_{i1} \dots w_{iT}). \quad (3)$$

After obtaining representation s_i of each post, for one user with L microblogs, the second LSTM layer f_2 aggregates the learned representation into a dimension fixed feature vector \mathbf{m} , formally:

$$\mathbf{m} = f_2(s_1 \dots s_L). \quad (4)$$

The overall UMHI framework simultaneously extracts the post content feature vector of follower i and followee j , we respectively denote the two feature vectors as \mathbf{m}_i and \mathbf{m}_j .

5.3 Unfollow History Matrix Factorization

Unfollow history records user's unfollowed-people list in recent time. Users with similar unfollow history show homogeneity in unfollow behavior. Specifically, if two users unfollowed similar followees recently, they may make similar unfollow decisions in the future. However, the unfollow behavior is sparse, statistics in our constructed dataset show that each user would only unfollow 6 users on average in a month. Therefore, it is hard to directly predict unfollow action by a user's unfollowed-people list.

Inspired by the matrix factorization (MF) based collaborative filtering algorithm, we construct a unfollow history matrix R_{train} as defined in Section 3. By matrix factorization we compress the unfollow history of follower i to latent vector p_i , and also map the unfollowed history of followee j to latent vector q_j . The dot product of p_i and q_j produces \hat{r}_{ij} , which estimates the value of r_{ij} in R_{train} . Formally:

$$\hat{r}_{ij} = p_i q_j^T, \quad (5)$$

where p_i and q_j are learned by optimizing the following regularized squared error:

$$\min_{p^*, q^*} \sum_{\langle i, j \rangle \in |V| \times |V|} (r_{ij} - p_i q_j^T)^2 + \lambda (\|p_i\|^2 + \|q_j\|^2). \quad (6)$$

The regularized squared error is minimized by stochastic gradient method (SGD) without negative sampling. The UMHI framework then incorporates both p_i and q_j as input. Experiment results demonstrate that feature vectors p_i and q_j are effective for the final prediction.

5.4 Output and Objective Function

Until now, given follower i and followee j , we’ve obtained network embeddings $\mathbf{n}_i, \mathbf{n}_j$, post content feature vectors $\mathbf{m}_i, \mathbf{m}_j$, and unfollow history feature vectors p_i, q_j . Among which \mathbf{n}_i and \mathbf{n}_j are spatial attributes, while $\mathbf{m}_i, \mathbf{m}_j, p_i$ and q_j are temporal attributes. To model the complicated interactions between spatial attributes and temporal attributes, we employ MLP layers to incorporate the unified representation as input. The nonlinearity of MLP layers can adequately represent the interaction mechanism between spatial attributes and temporal attributes. Finally, the output layer estimates the unfollow probability y_{ij} between follower i and followee j . Formally:

$$\mathbf{d} = \mathbf{n}_i \oplus \mathbf{n}_j \oplus \mathbf{m}_i \oplus \mathbf{m}_j \oplus p_i \oplus q_j, \quad (7)$$

$$y_{ij} = \text{Sigmoid}(\text{MLP}(\mathbf{d})). \quad (8)$$

The objective function to be minimized is defined as:

$$O = \sum_{(i,j) \in |V| \times |V|} (r_{ij} \log(y_{ij}) + (1-r_{ij}) \log(1-y_{ij})), \quad (9)$$

where r_{ij} is the element of R_{train} . The objective function is a format of cross-entropy. Since positive instances and negative instances are unbalanced, we sample equal number of positive and negative instances in each mini-batch.

6 Experiments

In this section, we firstly introduce dataset construction, evaluation metrics and UMHI implementation details, then we demonstrate the effectiveness of UMHI framework through comprehensive experiments. We show quantitative results on unfollow prediction, and respectively analyze spatial attributes, temporal attributes and the interaction effects through experimental results. Finally, we confirm the robustness of UMHI framework under different train-test split.

6.1 Dataset

Since prior researches did not publish their datasets, we only conduct experiments on the dataset we built. Section 4 has introduced how we construct the sub dataset E_{test} . To evaluate the model performance and prevent information leakage, we conduct five-fold cross validation on E_{test} . Specifically, we randomly holdout 20% of E_{test} for test, and combine the remaining data with $R_{train} = R \setminus E_{test}$ for training.

6.2 Metrics

We adopt the following three popular metrics to evaluate the performance of unfollow prediction:

- **Precision:** It measures the probability that a predicted positive instance would be the true positive.
- **Recall:** It measures the probability that the true positive would be predicted to be positive instance.
- **AUC:** It measures the probability that a classifier will rank a randomly chosen positive instance higher than a randomly chosen negative one.

6.3 Implementation Details

We first use jieba³ to cut microblog posts into separated words, use gensim⁴’s word2vec model to embed words into embeddings, and implement our UMHI model with Keras.

There are two stages for training our UMHI framework. In the first stage, we pretrain each component of UMHI, and in the second stage, we combine the three components of UMHI and fine tune the fusion MLP layer.

Stage I: Pretrain of Each Component. During the pre-train stage, we respectively pretrain LINE, HAN, and Matrix Factorization. We set the embedding size of LINE to 100, and train LINE for 100 epochs. When pretraining HAN, we set the size of LSTM cell to 100 and train for 10 epochs. The network is optimized by the adam optimizer (Kingma and Ba 2014), with the learning rate 0.001, and $\beta_1 = 0.1, \beta_2 = 0.001$. We optimize Matrix Factorization with latent size of 64, learning rate of 0.01, and 100 epochs.

Stage II: Global Fine-Tuning. In global fine-tuning stage, we fix the parameters of LINE, HAN and Matrix Factorization, and train the MLP layers with adam optimizer for 10 epochs, the learning rate is set to 0.001. We choose the performance of models when the precision, recall and AUC of test set achieves biggest value during training.

6.4 Comparison Methods

To justify the effectiveness of the proposed model, we compare the performance of our model with two kinds of baselines. Firstly, we compare UMHI with prior unfollow prediction methods. Prior unfollow prediction methods usually extract rule based features, and conduct training-test scheme on dataset with limited size, therefore showing inferior performance in our large-scale real-world setting. Secondly, we compare LINE with other network embedding methods so as to verify that local structure information plays an more important role than global information for unfollow prediction. The compared methods are listed as follows:

- **Doc2vec & Action Features+LR (DA + LR).** Referring to Maity’s method (Maity *et al.* 2018), we extract Doc2vec features and action features, then predict edge status by logistic regression.
- **Structural & Action Features + LR (SA + LR).** Referring to Kwak’s method (Kwak *et al.* 2012), we extract structural features and action features of users, then predict edge status by logistic regression.
- **Deepwalk.** Deepwalk (Perozzi *et al.* 2014) is a network embedding method, which leverages truncated random walks to obtain the structural information of each vertex.
- **Node2vec.** Node2vec (Grover and Leskovec 2016) is a network embedding method which designs a biased random walk procedure. In experiment, we use grid search to choose random walk strategies with the best performance.
- **The Proposed Method.** To demonstrate the effectiveness of different parts in UMHI model, we assemble LINE1, LINE2, HAN and matrix factorization (MF) into different combinations and compare their performance.

³<https://github.com/fxsjy/jieba>

⁴<https://radimrehurek.com/gensim/>

6.5 Experimental Results

Table 2: Results of unfollow prediction

Method	Precision	Recall	AUC
DA + LR	0.6292	0.5786	0.6840
SA + LR	0.6348	0.5270	0.6932
Deepwalk	0.6281	0.8123	0.6755
Node2vec	0.6288	0.8114	0.6707
LINE1	0.6316	0.7893	0.6845
LINE2	0.6413	0.7553	0.6924
LINE1 + LINE2	0.6458	0.7561	0.7059
HAN	0.6766	0.8292	0.7441
MF	0.7701	0.8686	0.8136
LINE1 + LINE2 + HAN	0.7034	0.8516	0.7718
UMHI	0.7868	0.8131	0.8673


Table 2 displays the performance across different models, from this table, we have the following analysis.

Overall Performance. To verify the validness of interaction effects between spatial and temporal attributes, we make the following comparisons, as shown in Figure 5a. Firstly, for the three methods (DA + LR, SA + LR, and LINE1 + LINE2 + HAN) that take same sources of input, DA + LR and SA + LR are two handcrafted methods that can hardly represent interaction effects, while LINE1 + LINE2 + HAN can represent highly nonlinear interaction mechanism. Experiment results show that LINE1 + LINE2 + HAN outperforms DA + LR by 7.42% and outperforms SA + LR by 6.86% in terms of precision, verifying that the interaction effects are powerful.

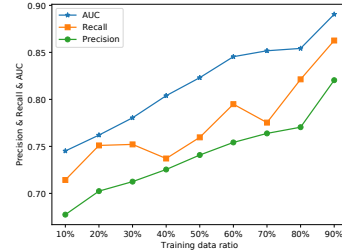
Meanwhile, from the comparison among LINE1 + LINE2, HAN, and LINE1 + LINE2 + HAN, we can observe that LINE1 + LINE2 + HAN significantly outperforms the first two models, confirming the effectiveness of interaction. Additionally, the comparison between LINE1 + LINE2 + HAN and UMHI demonstrate that feeding predictive unfollow history information would further boost prediction performance. Also, we notice that UMHI compromises on the recall value, that’s because unfollow is a sparse behavior and UMHI tends to be conservative.

Spatial Attributes Comparison. Compared with Deepwalk and Node2vec, we observe that LINE1 + LINE2 achieves an improvement of 1.73% in terms of precision. Such improvement verifies that local structure is more important than global structure under the unfollow prediction setting. Meanwhile, we discover that LINE2 is notably better than LINE1 for unfollow prediction. This is because LINE1 only considers the relationship between two nodes, while LINE2 considers the common neighbors between the two nodes. Therefore it reveals that the shared environment of two users can reflect the strength of the relationship more accurately than the relationship itself.

Temporal Attributes Comparison. We compare the two temporal attributes: unfollow history and post content. Experiment results show that MF performs 9.35% better than HAN in terms of precision, therefore unfollow history con-



(a) Precision, AUC and Recall of different comparison methods.



(b) Precision and AUC with different training and test data size.

Figure 5: Experimental Results

tributes significantly more than the post content.

Robustness Analysis. To verify the robustness of UMHI framework, we change the proportion of training set and test set and redo the experiments. Results in Figure 5b show that the model is effective under limited training data size. Even with small size of training set (10%-30%), our model can still have an acceptable and steady performance.

7 Conclusion

In this work, we constructed a large-scale social network dataset for unfollow behavior mining. Our dataset contains 1.8 million Chinese users and records relation dynamics of these users in a month. Based on the constructed dataset, we conducted extensive analyses on how users’ spatial attributes and temporal attributes affect their decisions on follow, and revealed the interaction effects between these two categories of attributes. Then, we proposed the UMHI framework to learn users’ spatial attributes and temporal attributes through their footprints in online social networks. The proposed framework outperforms baseline methods by a large margin, and the detailed factor analyses show that each component of UMHI is effective.

For future researches, the constructed dataset still contains ample social dynamics that deserve further exploring. During the recorded month, some followers have launched a burst of unfollow behaviors, and some relationships have experienced several status alterations, detecting such anomalies in online social networks is beneficial for some downstream tasks like depression detection and rumor detection.

8 Acknowledgments

This work is supported by National Key Research and Development Plan (2016YFB1001200), Beijing Academy of Artificial Intelligence (BAAI), and National Natural Science Foundation of China (61521002, 61972372, 61831022).

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- Catherine A Bliss, Morgan R Frank, Christopher M Danforth, and Peter Sheridan Dodds. An evolutionary algorithm approach to link prediction in dynamic social networks. *Journal of Computational Science*, 5(5):750–764, 2014.
- Ronald S Burt. Structural holes versus network closure as social capital. In *Social capital*, pages 31–56. Routledge, 2017.
- Ayushi Dalmia, Manish Gupta, et al. Towards interpretation of node embeddings. In *Companion of the The Web Conference 2018 on The Web Conference 2018*, pages 945–952. International World Wide Web Conferences Steering Committee, 2018.
- 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.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Funda Kivran-Swaine, Priya Govindan, and Mor Naaman. The impact of network structure on breaking ties in online social networks: unfollowing on twitter. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 1101–1104. ACM, 2011.
- Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, (8):30–37, 2009.
- Haewoon Kwak, Hyunwoo Chun, and Sue Moon. Fragile online relationship: a first look at unfollow dynamics in twitter. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1091–1100. ACM, 2011.
- Haewoon Kwak, Sue B Moon, and Wonjae Lee. More of a receiver than a giver: why do people unfollow in twitter? In *ICWSM*, 2012.
- David Liben-Nowell and Jon Kleinberg. The link-prediction problem for social networks. *Journal of the American society for information science and technology*, 58(7):1019–1031, 2007.
- Suman Kalyan Maity, Ramanth Gajula, and Animesh Mukherjee. Why did they# unfollow me?: Early detection of follower loss on twitter. In *Proceedings of the 2018 ACM Conference on Supporting Groupwork*, pages 127–131. ACM, 2018.
- Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab, 1999.
- 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.
- Daniele Quercia, Harry Askham, and Jon Crowcroft. Tweetlda: supervised topic classification and link prediction in twitter. In *Proceedings of the 4th Annual ACM Web Science Conference*, pages 247–250. ACM, 2012.
- Daniele Quercia, Mansoureh Bodaghi, and Jon Crowcroft. Loosing friends on facebook. In *Proceedings of the 4th Annual ACM Web Science Conference*, pages 251–254. ACM, 2012.
- Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. Line: Large-scale information network embedding. In *Proceedings of the 24th International Conference on World Wide Web*, pages 1067–1077. International World Wide Web Conferences Steering Committee, 2015.
- Bo Xu, Yun Huang, Haewoon Kwak, and Noshir Contractor. Structures of broken ties: exploring unfollow behavior on twitter. In *Proceedings of the 2013 conference on Computer supported cooperative work*, pages 871–876. ACM, 2013.
- Yang Yang, Jia Jia, Boya Wu, and Jie Tang. Social role-aware emotion contagion in image social networks. In *AAAI*, pages 65–71, 2016.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. Hierarchical attention networks for document classification. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1480–1489, 2016.