ACTPred: Activity Prediction in Mobile Social Networks

Jibing Gong, Jie Tang*, and A. C. M. Fong

Abstract: A current trend for online social networks is to turn mobile. Mobile social networks directly reflect our real social life, and therefore are an important source to analyze and understand the underlying dynamics of human behaviors (activities). In this paper, we study the problem of activity prediction in mobile social networks. We present a series of observations in two real mobile social networks and then propose a method, ACTPred, based on a dynamic factor-graph model for modeling and predicting users' activities. An approximate algorithm based on mean fields is presented to efficiently learn the proposed method. We deploy a real system to collect users' mobility behaviors and validate the proposed method on two collected mobile datasets. Experimental results show that the proposed ACTPred model can achieve better performance than baseline methods.

Key words: social prediction; activity prediction; user modeling; social networks

1 Introduction

Mobile social networks offer the unique advantage of allowing users to find and connect at real time via mobile phones. A current trend for Internet social networks such as Facebook and Twitter is to turn mobile. In parallel, native mobile social networks such as Foursquare and Gowalla have also been created. The natural characteristics of the mobile social networks make it very different from traditional web-based social networks. First, all users in the mobile social networks use their real names. Second, in the mobile social networks, the relationships between users are the same as they are in the reality. Third, mobile users' behaviors (e.g., SMS, calling log, location, etc.) are

all related to real-world behaviors. This provides an unprecedented opportunity for us to understand the underlying dynamics of users' behaviors in the physical social network.

It is well-recognized that users' activities in a mobile social networks are influenced by various complex and subtle factors[1, 2]. In this work, we aim to answer an interesting question: i.e., can we predict a user's activities based on his/her historic behavior log and mobile social network information, and how to distinguish the effects of different social factors that determine the users' activities?

Recently, considerable related studies have been conducted, for example, activity recognition[3-8], dynamic emotion analysis[9-13], dynamic social network analysis[14-18], and social influence analysis[19-23]. Emotion analysis is to study how an individual's emotional state (e.g., happiness and loneliness) propagates through social relationships[9-11]. Dynamic social network analysis is to model how friendships drift over time using a dynamic model[17] or to investigate how different pre-processing decisions and different network forces such as selection and influence affect the modeling of dynamic networks[18]. Social influence analysis either aims to verify the existence of social influence[19, 24],

- Jibing Gong is with School of Information Science and Engineering, Yanshan University, Qinhuangdao 066004, China. E-mail: gongjibing@gmail.com. The work was done when Jibing Gong was visiting Tsinghua University.
- Jie Tang is with Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China. E-mail: jietang@tsinghua.edu.cn.
- A. C. M. Fong is with University of Glasgow in Singapore, Singapore. E-mail: acmfong@gmail.com.
- *To whom correspondence should be addressed.
  Manuscript received: 2014-05-19; accepted: 2014-05-20

quantify the influence strength[21], or model the influence diffusion process[20]. In the social activity prediction problem we are going to address, we try to model the various factors that may influence users' dynamic activities into a unified model and to predict users' future activities. The social activity prediction problem addressed in this paper is very different from those addressed in these works. Most existing research on activity recognition focuses on classifying the users activity based on his attributes (such as location). However, few works consider the activity prediction problem in (mobile) social networks.

To clearly demonstrate this problem, we give an example in Fig. 1. The left part shows the social network of John and his friends. Each user is associated with his/her historic behavior log at time t. Our goal is to predict each user's activity at time $t + 1$, for example, what will John do tomorrow? In the prediction, we need to not only consider the user's personal attributes (e.g., location, call, SMS message, emotion, etc.), but also consider the user's friends' activities (social influence) and the user's activities at previous time (temporal dependence). Another technical challenge here is that, as shown in Fig. 1, the change of each user's behavior (or attribute) may happen on a continuous time scale. Existing methods that partition the dynamic data into different timestamps would not work well.

In this paper, we formally formulate the problem of activity prediction in mobile social networks and perform a series of observations in two mobile social networks. Based on the observations, we propose a dynamic factor-graph model named ACTPred for predicting users' mobility activities. ACTPred learns a discriminative model for predicting users' activities at time $t + 1$ by incorporating different types of

Fig. 1 Problem illustration of social activity prediction.

information (personal attributes, social influence, and temporal dependence) before time t into a unified model. We develop an approximate algorithm using mean fields to efficiently learn the proposed method.

To evaluate the proposed method, we deployed a real system to collect users' mobility behavior records, including location, calling logs, and SMS text messages. We also asked the user to annotate their daily emotional status and activities. We collected two mobile datasets. Our experimental results on the two datasets show that the proposed ACTPred model can achieve better performance (10%-30% in terms of F1-measure, with p-test $\ll 0.01$) for social activity prediction than several alternative models.

2 Problem Formulation

In this section, we first give several preliminary definitions and then present a formal definition of the problem.

A static mobile social network can be represented as $G = (V, E)$, where V is the set of $|V| = N$ users and $E \subset V \times V$ is the set of undirected links between users. Each user has various activities in the mobile social network.

Definition 1 Activity A user v_i's activity status at time t is represented as $y_i^t \in \mathcal{Y}$, where $\mathcal{Y}$ is the space of the activity status. We denote the historic log of all users' activity status up to time t as $Y = \{y_i^t\}_{i,t}$.

In general, the activity status can be defined as a set of events. For example, in our mobile social networks, we define eight different events: shopping, work, play, study, sleep, walk, eat, and meeting. In addition, each user v_i can be associated with a number of attributes x_i (e.g., SMS message and emotion). All the attributes would change over time; hence we associate the value of the j-th attribute x_{ij} of user v_i with the time t; i.e., x_{ij}^t (please note that t is defined on a continuous scale). Given this, we can define the input of our problem as a dynamic continuous network.

Definition 2 t-Dynamic continuous network The dynamic continuous network (from time 0 to time t) is denoted as $G^t = (V, E^t, X^t, Y^t)$, where V is the set of users, $e_{ij}^t \in E^t$ is the edge between users v_i and v_j up to time t, and X^t is the continuous time-evolving attributes of all users in the network, and Y^t represents the set of activities of all users in the social network.

We use the superscript t to denote that the dynamic continuous information in the network G^t is up to time
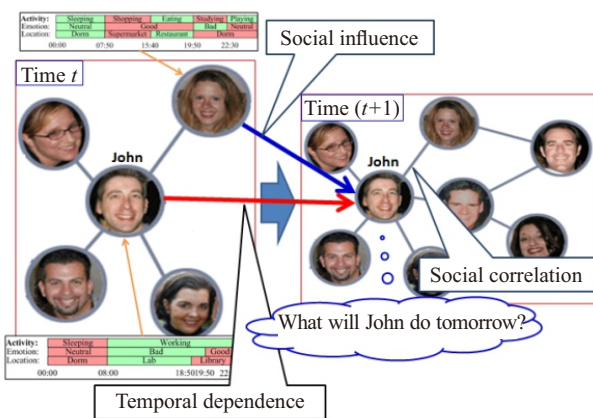
t; that is, all edges E, attribute changes X, and activity status changes Y are recorded until time t. Based on these concepts, we can define the tasks of activity prediction in mobile social networks.

Problem 1 Social activity prediction Given a dynamic continuous network G^t, the goal is to learn a predictive function f to predict the users' activities $Y^{(t+1)}$ at a future time $(t+1)$. Formally, we have

$$f(X^{(t+1)}, V, E^{(t+1)}|G^t) \to Y^{(t+1)}.$$

Our formulation of the social-activity prediction problem addressed in this paper is very different from those in existing works. Most existing works on activity recognition do not consider the problem in the mobile social networks. A few works[9, 10, 25] studied how an individual's emotional state (e.g., happiness and loneliness) influences friends in their social network. However, emotion is very different from activity. The underlying influence patterns are quite different and, more importantly, besides social influence, what are the other fundamental factors that affect users' activities?

3 Data and Observations

Before proposing our method, we first engage in some investigation of the degree to which a user's activity correlates with other social patterns, since a major motivation of our work is to discover the underlying factors that influence the activity dynamics.

3.1 Data sets

We deployed a system in the author's university to collect the personal mobility data. Specifically, sponsored by a large mobile company, we developed software and installed it on mobile phones, under an agreement with the mobile user. The software automatically collected and uploaded the user's behavior data onto a server. For each user, we collected his calling log, SMS text messages, location, etc. Also we invite all users to annotate their activities and emotional status. This data collection resulted in two datasets: One is from May 13th to July 30th, 2010 (Dataset1), and the other is from November 1st to December 31st, 2010 (Dataset2). The two mobile social networks data sets represent more than 79 200 hours of continuous data on human activities and emotional status. In total, there are about 25 114 human annotations of the activity status. The two data sets are from two different groups of users (about 130 users in total). All the users are students in a university. The first

group has 30 users, who have 163 friend relationships with each other (5.5 on average), and in total we have 9756 activity annotations. The second group is comprised of 100 users and 280 social relationships (9.2 on average); in total, these users made 61 324 activity annotations. Statistics of the two data sets are shown in Table 1.

3.2 Correlation pattern analysis

We conducted a series of analyses on the two mobile social networks. In the analyses, we focused on the following aspects: (1) **attributes correlation:** how a user's attributes correlate with his activity status; (2) **temporal correlation:** how a user's current activity correlates with his activity in the recent past; (3) **social correlation:** how a user's activity correlates with activities of his friends.

3.2.1 Attributes correlation

Each user in the social networks has her/his own characteristics, which may partly determine his activity at a specific time. In our mobile social networks, the main attributes include location, different time (hour) of a day, emotional status (e.g., "happy" or "sad;), calling log, SMS text, etc. We first analyze the correlation between each attribute and the user's activity.

Figure 2a shows the correlation between user's emotion and his activity. It can be easily seen that when the emotion of a user stays at "Positive" ("good" and "wonderful"), she is inclined to go shopping or playing, and vice versa.

On the other hand, "meeting" seems to have a very negative effect: When a user's emotion stays at "bad" and "terrible", he/she is very likely to be in a meeting. Users' activities are also highly time-dependent. Some activities have a clear periodicity. Figure 2b shows the likelihood of each activity being performed at different times of day. Generally speaking, the users are more active in the afternoon. Comparing with "work", "study" mainly happens in the morning and night. "Location" is another important factor to influence users' activities. Figure 2c shows the correlation of user's location with her/his current activity. We see that the users' activities

Table 1 Statistics of the two MSN data sets.

	Number of users	Number of links	Number of avg. links	Number of label
Dataset1	30	164	10.8	9756
Dataset2	100	4632	9.2	61 324

(a) Emotion correlation

(b) Time (hour) of a day

(c) Activity correlation

Fig. 2   Correlation analysis. (a) The user's activities correlate with his/her emotion status. (b) The user's likelihood of performing an activity at different times (hour) of a day. (c) Location correlation. The user's locations correlate with his/her current activity. Different icons represent "shopping", "work", "study", "sleep", "walk", "eat", and "meeting".

imply some places for "eat", some places for "study", etc.

3.2.2   Temporal correlation

The user's activity is highly correlated to her/his behavior in the recent past. Figure 3 confirms this temporal correlation of users' activities in mobile social networks. To most of the activities, a user's activity status has a higher probability to remain the same, except with some activities, such as "eat" and "walk".

3.2.3   Social correlation

In this analysis, we try to find whether and how friends' activities influence each other. To gain a fine-grained understanding influence patterns, we categorize the relationship between users into four classes: "stranger", "know each other", "friend", and "good friend". Figure 4 shows how the user's activity status is influenced by her/his friends' activities. We see that, on average, the user has a higher likelihood to do the same as her/his friends than strangers. We also find an interesting phenomenon that "study" and "shopping"

Fig. 3   Temporal correlation. The user's previous activity's influence on his current activity. *Average* is the likelihood of a user with activity *y* and *Dependent* is the likelihood given that the user does the same at previous time.
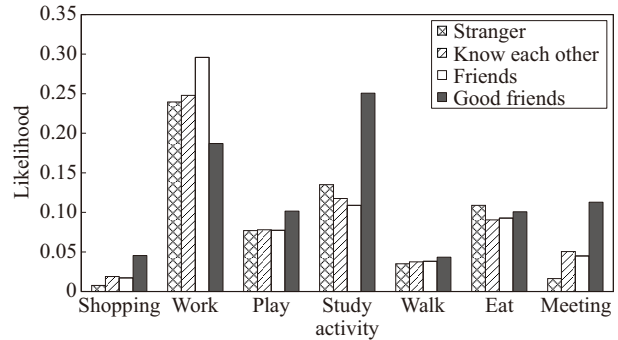

Fig. 4   Social correlation. The influence of friends' activities on one's current activity. It represents the likelihood of a user doing a activity when his/her good friends, friends, people he/she know, strangers do the same.

seem to be more infectious in the network. When a user starts to "study" or "shopping", the possibility that her/his good friends are also studying and shopping is twice as high as strangers. Some other activities such as "play", "eat", and "walk" seem only have influence effects between good friends.

4   Our Approach

To summarize, for modeling and predicting users' activities, we have the following intuitions:

  (1) Attribute correlation: Personal attributes have a strong influence on users' activities.

  (2) Temporal correlation: One's activity is highly time-dependent.

  (3) Social correlation: One's activity is influenced by her/his friends' activities.

By leveraging these intuitions, we formulate the problem of activity prediction in a dynamic continuous factor-graph model and propose a method, ACTPred, for predicting users' activity. The basic idea in this
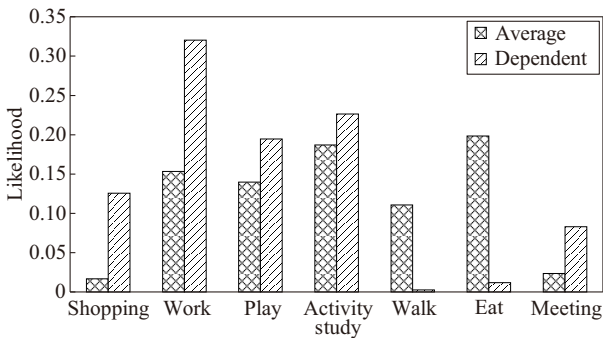
method is to define the above intuitions using factor functions and then combine all the factor functions in a probabilistic model. An objective function is defined by the joint probability of the factor functions, and training the model is to find a parameter configuration that maximizes the joint probability. For model-parameter estimation, we present a Mean-Field-based algorithm. Prediction of the user's activity is to find an activity status configuration that can maximize the probability according to the training parameters.

4.1 The proposed model

In particular, we define three kinds of factor functions:

- **Attribute correlation factor function** $\{f(x_{ik}^t, y_i^t)\}_k$. It denotes the attribute value associated with each user v_i at time t.
- **Friend influence factor function** $g(y_i^t, y_j^{t'})$, $t' < t$, represents the influence of user v_j's activity at time t' on user v_i's activity at time t. This function only appears when user i and user j are friends, and we define the function as $g_{ij}(y_i^t, y_j^{t'}) = \beta_{ij}(t - t')g(y_i^t, y_j^{t'})$.
- **Temporal correlation factor function** $h(y_i^{t'}, y_i^t)$, $t' < t$, represents the dependency of one activity status at time t on his activity at the recent past time. Specifically, we define function h as $h_i(y_i^t, y_i^{t'}) = \lambda_{ij}(t - t')h(y_i^{t'}, y_i^t)$.

We define the factor functions $f(x_{ik}^t, y_i^t)$, $g(y_i^t, y_j^{t'})$, and $h(y_i^{t'}, y_i^t)$ as binary functions, i.e., if x_{ik}^t and y_i^t fit the data, the function value is 1, otherwise 0. The three factors can be instantiated in different ways, reflecting our prior knowledge for different applications. Here, we use the mobile social networks as the example to explain how we define the factor functions. Based on the mobile social networks data observation, we define the following attribute-factor functions:

Location This feature represents the location of the user. In mobile social networks, we use GSM data to locate the user. The location is usually denoted as the longitude and the latitude values. To reduce noise in the data, we group the location points by area code (one attribute of the GSM data).

Emotion This feature represents the user's emotional status. In our dataset, all users are required to annotate their emotional status when they perform the activities.

Time (hour) of a day This feature represents the time when the user performed an activity. For example, if Nancy went shopping at 10 o'clock in the morning,

then the value of this feature is 10.

Finally, by combining all the factor functions together, we can define the following objective function:

$$p(Y|G^t) = \frac{1}{Z}\exp\{\sum_{v_i \in V}\sum_{x_{ik}^t \in X}\alpha_k f_k(x_{ik}^t, y_i^t) +$$
$$\sum_{v_j \in \mathrm{NB}(v_i)}\sum_{(y_i^t, y_j^{t'}) \in Y^t}\beta_{ji}(t - t')g(y_i^t, y_j^{t'}) +$$
$$\sum_{v_i \in V}\sum_{(y_i^t, y_i^{t'}) \in Y^t}\lambda_i(t - t')h(y_i^{t'}, y_i^t)\} \qquad (1)$$

where Z is a normalization factor; $\mathrm{NB}(v_i)$ denotes the set of v_i's neighbors in the network; $(y_i^t, y_j^{t'})$ indicates a pair of activities between v_i at time t and v_j at a recent past time t'.

4.2 Model learning

Training the factor graph model is to estimate a parameter configuration $\theta = (\{\alpha_k\}, \{\beta_{ji}\}, \{\lambda_i\})$ from a given historic attribute-value log X^t, which maximizes the log-likelihood objective function $\mathcal{L}(\theta) = \log p_\theta(Y|G^t)$, i.e.,

$$\theta^\star = \arg\max_\theta \log p(Y = y|x, \theta) \qquad (2)$$

It is usually intractable to do exact inference in such a graphical probabilistic model. The intrinsic difficulty is to calculate the normalization factor Z, which sums up all possible configurations of Y that is exponentially proportional to the number of nodes in the graph. In this paper, we employ a Mean-Field algorithm[26] to make an approximate inference on the graphical model. Details of the learning algorithm are summarized in Algorithm 1, where $b_i(y)$ is the marginal distribution of user i doing activity y. The general process is to first calculate the marginal probability using mean field and then use the gradient learning to update each parameter $\theta = (\{\alpha_k\}, \{\beta_{ji}\}, \{\lambda_i\})$.

4.3 Activity prediction

Based on the learned parameter θ, we can predict the users' future activities. Instead of constructing a graph model again and applying the learned parameter to make an inference on the graph, in this paper, for simplicity, we utilize an ICA[27] for predicting users' activities. This algorithm can be also viewed as a "hard" version of the Mean-Field algorithm. In particular, it performs a preliminary prediction using only local attributes, and then propagates the prediction probability in the social network to update the

---

Algorithm 1: Model learning.

Input: number of iterations, parameters $\boldsymbol{\theta} = (\{\alpha_k\}, \{\beta_{ji}\}, \{\lambda_i\})$, and graph G, learning rate η;

Output: learned parameters $\boldsymbol{\Theta} = (\{\alpha_k\}, \{\beta_{ji}\}, \{\lambda_i\})$;

Initialize parameters;

repeat

 for each $Y_i \in \vec{Y}$ **do**

  | $b_i(y) \leftarrow 1$

 end

 repeat

  for each $Y_i \in \vec{Y}$ **do**

   | $b_i(y) \leftarrow \alpha\phi_i(x^t, y^t) \cdot$

   | $\exp\{\sum_{(Y_i, Y_j) \in E} \sum_{y'} \beta_{ij} g(y_i{}^t, y'_j) b_j(y'_j)\}$

   | where α is a normalizer.

  end

 until all $b_i(y)$ stop changing;

 for each $\theta_i \in \boldsymbol{\Theta}$ **do**

  | Calculate $\boldsymbol{\Delta}_i$;

  | $\theta_i^{\text{new}} = \theta_i^{\text{old}} + \eta\boldsymbol{\Delta}_i^{\text{old}}$;

 end

until convergence;

---

prediction results. The update scheme is similar to PageRank. Such a strategy has a theoretical convergence property, and was also used in Ref. [28].

5 Experimental Results

In this section, we present the experimental results of the proposed ACTPred, and then evaluate its effectiveness and efficiency.

We use the datasets introduced in Section 3 in our evaluation. We consider two alternative methods for social activity prediction.

- **SVM-Simple**. The method only uses users' attributes (i.e., user attribute factor functions) as features to train a classification model, and then use the trained model to predict the activity of users. For SVM, we use LIBSVM (http://www.csie.ntu.edu.tw/ cjlin/libsvm).
- **SVM-Net**. Besides using users' attributes, the method also includes the network information (i.e., social correlation) as features.

We evaluate the performance of activity prediction by different methods in terms of Precision, Recall, F1-Measure, and average Accuracy. The learning algorithm for ACTPred is implemented in C++ and all experiments are conducted on the server with Windows Server 2003, Intel Xeon(TM) CPU 3.20 GHz, and 4 GB memory.

5.1 Prediction performance

On both datasets, we use the historical data (from time 0 to $t-1$) of the users as the training data. Then we predict a user's activity status at time t given the attributes of the user at time t. In particular, we choose the data in the last 10 days as the test data and the rest (about 50 days' worth) as the training data. Table 2 lists the two approaches' average prediction performance on the two data sets. Table 3 further provides the detailed performance of different approaches on each activity. From the results, we have the following observations.

First, these results show that our method significantly outperforms the comparison methods on both data sets, with an average of 6% (by F1-Measure) improvement compared with the SVM-Simple and SVM-Net methods.

Second, for some activities, such as shopping, the precision and recall of SVM-Simple and SVM-Net are both 0.0, which means the two methods do not predict any "shopping" activity in the test data. By comparison, our method, by taking advantage of the social influence and temporal correlation, reaches a performance of 12.05% (by F1-Measure).

Third, from Tables 2 and 3, we can also see that simply combining all the features (social correlation, temporal correlations) together (as in SVM-Net) will yield an improvement (2.7% by F1) over SVM-Simple. However, the performance is still unsatisfactory. Our method, by leveraging of social influence information and the temporal correlation information, achieves a further improvement over SVM-Net.

5.2 Analysis and discussion

We conduct the following analysis on the results.

5.2.1 Effect of the number of sampling iterations

We design an experiment to see how the number of the iterations in our learning algorithm affects

Table 2 **Average accuracy and F1-score of activity prediction in the two mobile social networks with different approaches.** (%)

Method	Dataset1		Dataset2	
	Accuracy	F1	Accuracy	F1
ACTPred	**73.65**	**30.52**	**73.26**	**30.43**
SVM-Simple	49.61	23.61	51.63	24.19
SVM-Net	51.77	23.59	54.13	26.83

Table 3 **Performance of activity prediction in the two mobile social networks with different approaches.** **(%)**

Activity	Method	MSN Dataset1			MSN Dataset2		
		Precision	Recall	F1-Measure	Precision	Recall	F1-Measure
Working	ACTPred	4.56	39.09	42.09	44.31	42.22	**43.24**
	SVM-Simple	30.31	83.35	44.45	25.52	13.46	17.62
	SVM-Net	30.66	84.74	**45.03**	29.48	56.98	38.86
Studying	ACTPred	42.71	41.38	**42.03**	45.33	39.55	**42.24**
	SVM-Simple	50.61	18.08	26.65	33.33	4.93	8.60
	SVM-Net	51.15	13.80	21.74	33.14	7.31	11.98
Playing	ACTPred	38.41	42.82	**40.50**	29.89	36.29	**32.78**
	SVM-Simple	50.61	18.08	26.65	33.33	4.93	8.60
	SVM-Net	51.15	13.80	21.74	33.14	7.31	11.98
Eating	ACTPred	8.91	28.18	13.54	12.70	34.83	18.62
	SVM-Simple	39.90	27.73	32.72	27.57	50.86	**35.75**
	SVM-Net	40.28	28.63	**33.47**	32.28	36.69	34.34
Shopping	ACTPred	6.85	50.0	**12.05**	19.57	26.47	**22.50**
	SVM-Simple	0.0	0.0	0.0	0.0	0.0	0.0
	SVM-Net	0.0	0.0	0.0	0.0	0.0	0.0
Walking	ACTPred	10.32	37.14	**16.15**	6.47	30.56	**10.68**
	SVM-Simple	76.47	1.03	2.03	0.0	0.0	0.0
	SVM-Net	93.73	1.17	2.30	10.00	0.08	0.15

the prediction performance. We used the average F1-Measure of all classifiers to measure the overall performance. Since the results on Dataset1 and Dataset2 look almost the same, we only display the experimental result on Dataset2 in Fig. 5. As it illustrates, the ACTPred algorithm using mean field learning converges in less than 20 iterations on both data sets, suggesting a good convergence property.

5.2.2 Factor contribution analysis

In ACTPred, we consider three types of factor functions: social influence (F), temporal-correlation (T), and user attributes. We remove these factors one-by-one to show their impacts on the prediction performance. In particular, we first remove friend influence factor function, denoted by ACTPred-F, then we remove the time-dependent factor function, denoted by ACTPred-FT.

Figure 6 shows the average F1-Measure scores of the different versions of ACTPred. We observe a clear drop on the performance when ignoring some of the factors, which indicates our method works well by combining the different factor functions. From Fig. 6, we can see that the temporal correlation is a very important factor. There is a drop from "ACTPred-F" to "ACTPred-FT"; performances decrease from 29.1% to 21.6% in Dataset1, and from 30.0% to 24.5% in Dataset2, when the temporal factor was removed. We also study how users' most recent activities could influence users' current activities. Figure 7 shows three prediction results: (1) without considering temporal correlation (denoted as −T), (2) only considering the most recent previous activity (denoted as +T1), and (3) considering two recent previous activities (denoted as +T2).

Fig. 5 **The influence of sampling iterations (Dataset 2).**

Fig. 6 **Contribution of different factor functions. ACTPred-F excludes friend influence factor. ACTPred-FT excludes both the friend and temporal-correlation factors.**

Fig. 7 Prediction performance of ACTPred using different temporal factors. "–T", "+T1", and "+T2" respectively stand for the prediction results without considering temporal correlation, with considering the most recent previous activity, and considering two recent previous activities.

5.2.3 Error analysis

We conducted an error analysis on the results of our approach. We observe two major sources of errors.

- **Missing data**. Some mobile context data is missing. For example, "Yin" did not turn on his phone for a half day, thus his mobile context data was lost. The location and the temporal correlation information is very important for building a successful predictive model.

- **Unexpected labeling**. Sometimes, the users may change their activities "unpredictably". For example, "Wei" continuously worked in the laboratory for two days during June 15th and 16th. We finally found that this is because of an approaching deadline. However, such deadline information is not collected in our model.

6 Related Work

6.1 Activity recognition

Some previous papers have investigated the problem of activity recognition. In Ref. [4], a probabilistic approach referred to as Segmentation-based Activity Recognition (SAR) for activity recognition is proposed. In this approach, instead of the precise location information, a rough idea of the general trends of a user's movements is sufficient for activity recognition. Reference [29] studied the problem of activity prediction from videos, and Ref. [30] further studied collective activity prediction from videos. However, all these works do not consider the social perspective in the prediction. One of the key features of real-world human activities is that multiple goals are often mixed together in sophisticated way. Reference [5] analyzed the MIT PlaceLab House

data and presented a two-level probabilistic framework from observed sensor-reading sequences using a Conditional Random Fields (CRF) model. In Ref. [31], a hybrid approach by combining ontology and statistical inference is proposed. However, none of these takes social influence into consideration. Comparing with these works, in our paper we study the problem of activity prediction in mobile social networks, and propose a continuous dynamic factor graph to solve the problem.

6.2 Dynamic behavior analysis

Quite a few works have been conducted for social dynamic behavior analysis. Shi et al.[32] studied the patterns of participation behavior, and the features that influence such behavior on different real world data sets. Tang and Liu[33] proposed relational learning to address the interdependency among data instances. Backstrom et al.[34] proposed a partitioning on the data that selects for active communities of engaged individuals. Eagle et al.[35] presented a new method for measuring human social behavior based on mobile phone data. Tang et al.[21] built a topical factor graph model to measure the influential strength in the social network. And Tan et al.[36] proposed a noise tolerant model for predicting user's actions in online social networks. Also, some works are conducted in terms of social science, trying to identify the principles underneath human beings' behaviors in a social network. The works of Rosenquist et al.[37] and Fowler and Christaki[38] are among the representative researches in this field.

6.3 Social network analysis

Considerable research has been done for dynamic social network analysis and social influence. Kossinets and Watts[39] provided statistical analysis on a empirical data set, revealing interesting patterns about how nodes' interaction affects the network structure. Onnela et al.[40] studied the correlation between interaction strength and network structure. Ahmed and Xing[41] proposed a random field based model to infer the interaction between nodes with the samples of users' status at different time slots. Crandall et al.[24] further investigated the correlation between social similarity and influence. More recently, La Fond and Neville[42] examined the effects of social influence on people's opinion. Gomez-Rodriguez et al.[43] proposed an effective model to track the flow of information and influence in an online social network. Some other

works also study this problem from the perspective of heteroengeous information network[44].

7 Conclusions

In this paper, we study a novel problem of activity prediction in mobile social networks. We propose a method, called ACTPred, for modeling and predicting users' activities in the social network. We present a series of observation analyses and propose a factor graph model to formalize the discovered intuitions in a unified model. For model learning, we employ a Mean Field algorithm to obtain an approximate solution. Experimental results on two real social networks demonstrate that the proposed approach can accurately predict users' activities and obtains a clear improvement (10%-20% by F1) over the comparison methods.

References

[1] M. Granovetter, The strength of weak ties, *American Journal of Sociology,* vol. 78, no. 6, pp. 1360-1380, 1973.

[2] D. Krackhardt, The strength of strong ties: The importance of philos in networks and organization, in *Networks and Organizations*, N. Nohria and R. G. Eccles, Eds. Cambridge, USA: Harvard Business School Press, 1992, pp. 216-239.

[3] S. Abdullah, N. D. Lane, and T. Choudhury, Towards population scale activity recognition: A framework for handling data diversity, in *Proc. AAAI'12*, Toronto, Canada, 2012, pp. 851-857.

[4] J. Yin, Q. Yang, D. Shen, and Z. N. Li, Activity recognition via user-trace segmentation, *ACM Transactions on Sensor Networks,* vol. 4, no. 4, pp. 19-34, 2008.

[5] D. H. Hu and Q. Yang, Cigar: Concurrent and interleaving goal and activity recognition, in *Proc. AAAI'08*, Chicago, USA, 2008, pp. 1363-1368.

[6] J. R. Kwapisz, G. M. Weiss, and S. A. Moore, Activity recognition using cell phone accelerometers, *SIGKDD Explor. Newsl.,* vol. 12, no. 2, pp. 74-82, Mar. 2011.

[7] E. Kim, S. Helal, and D. J. Cook, Human activity recognition and pattern discovery, *IEEE Pervasive Computing,* vol. 9, no. 1, pp. 48-53, Jan.- Mar. 2010.

[8] Y. M. Zhang, Y. F. Zhang, E. Swears, N. Larios, Z. H. Wang, and Q. Ji, Modeling temporal interactions with interval temporal bayesian networks for complex activity recognition, *IEEE Trans. Pattern Anal. Mach. Intell.,* vol. 35, no. 10, pp. 2468-2483, Oct. 2013.

[9] J. H. Fowler and N. A. Christakis, Dynamic spread of happiness in a large social network: longitudinal analysis over 20 years in the framingham heart study, *British Medical Journal,* vol. 337, no. a2338, pp. 1-9, Dec. 2008.

[10] J. T. Cacioppo, J. H. Fowler, and N. A. Christakis, Alone in the crowd: The structure and spread of loneliness in a large social network, *Journal of Personality and Social Psychology,* vol. 97, no. 6, pp. 977-991, Dec. 2009.

[11] J. Tang, Y. Zhang, J. M. Sun, J. H. Rao, W. J. Yu, Y. R. Chen, and A. C. M. Fong, Quantitative study of individual emotional states in social networks, *IEEE Trans. Affe. Comp.,* vol. 3, no. 2, pp. 132-144, April-June 2012.

[12] J. Jia, S. Wu, X. H. Wang, P. Y. Hu, L. H. Cai, and J. Tang, Can we understand van Gogh's mood? Learning to infer affects from images in social networks, in *Proc. ACM MM'13*, Barcelona, Spain, 2013, pp. 857-860.

[13] Y. Yang, J. Jia, S. M. Zhang, B. Y. Wu, Q. C. Chen, J. Z. Li, C. X. Xing, and J. Tang, How do your friends on social media disclose your emotions? in *Proc. AAAI'14*, Quebéc, Canada, 2014, pp. 1-7.

[14] D. Gruhl, R. Guha, D. Liben-Nowell, and A. Tomkins, Information diffusion through blogspace, in *Proc. WWW'04*, New York, USA, 2004, pp. 491-501.

[15] J. Kleinberg, Temporal dynamics of on-line information streams, in *Data Stream Managemnt: Processing High-Speed Data.* New York, USA: Springer, 2005.

[16] J. Leskovec, L. Backstrom, R. Kumar, and A. Tomkins, Microscopic evolution of social networks, in *Proc. KDD'08*, Las Vegas, USA, 2008, pp. 462-470.

[17] P. Sarkar and A. W. Moore, Dynamic social network analysis using latent space models, *ACM SIGKDD Explorations Newsletter,* vol. 7, no. 2, pp. 31-40, Dec. 2005.

[18] J. Scripps, P. N. Tan, and A. H. Esfahanian, Measuring the effects of preprocessing decisions and network forces in dynamic network analysis, in *Proc. KDD'09*, Paris, France, 2009, pp. 747-756.

[19] A. Anagnostopoulos, R. Kumar, and M. Mahdian, Influence and correlation in social networks, in *Proc. KDD'08*, Las Vegas, USA, 2008, pp. 7-15.

[20] D. Kempe, J. Kleinberg, and É. Tardos, Maximizing the spread of influence through a social network, in *Proc. KDD'03*, Washington, DC, USA, 2003, pp. 137-146.

[21] J. Tang, J. M. Sun, C. Wang, and Z. Yang, Social influence analysis in large-scale networks, in *Proc. KDD'09*, Paris, France, 2009, pp. 807-816.

[22] J. Tang, S. Wu, and J. M. Sun, Confluence: Conformity influence in large social networks, in *Proc. KDD'13*, Chicago, IL, USA, 2013, pp. 347-355.

[23] J. Zhang, J. Tang, H. L. Zhuang, C. W. K. Leung, and J. Z. Li, Role-aware conformity influence modeling and analysis in social networks, in *Proc. AAAI'14*, Quebéc, Canada, 2014, pp. 1-7.

[24] D. Crandall, D. Cosley, D. Huttenlocher, J. Kleinberg, and S. Suri, Feedback effects between similarity and social influence in online communities, in *Proc. KDD'08*, Las Vegas, USA, 2008, pp. 160-168.

[25] Y. Zhang, J. Tang, J. M. Sun, Y. R. Chen, and J. H. Rao, Moodcast: Emotion prediction via dynamic continuous factor graph model, in *Proc. ICDM'10*, Sydney, Australia, 2010, pp. 1193-1198.

[26] E. P. Xing, M. I. Jordan, and S. Russell, A generalized mean field algorithm for variational inference in exponential families, in *Proc. UAI'03*, Acapulco, Mexico, 2003, pp. 583-591.

[27] J. Neville and D. Jensen, Iterative classification in relational data, in *Proc. AAAI'00*, Austin, USA, 2000, pp. 13-20.

[28] X. J. Zhu and J. Lafferty, Harmonic mixtures: Combining mixture models and graph-based methods for inductive and scalable semi-supervised learning, in *Proc. ICML'05*, Bonn, Germany, 2005, pp. 1052-1059.

[29] M. S. Ryoo, Human activity prediction: Early recognition of ongoing activities from streaming videos, in *Proc. ICCV'11*, Barcelona, Spain, 2011, pp. 1036-1043.

[30] W. G. Choi, K. Shahid, and S. Savarese, Learning context for collective activity recognition, in *Proc. CVPR'11*, Colorado Springs, USA, 2011, pp. 3273-3280.

[31] D. Riboni and C. Bettini, Cosar: Hybrid reasoning for context-aware activity recognition, *Personal and Ubiquitous Computing,* vol. 15, no. 3, pp. 271-289, Mar. 2011.

[32] X. L. Shi, J. Zhu, R. Cai, and L. Zhang, User grouping behavior in online forums, in *Proc. KDD'09*, Paris, France, 2009, pp. 777-786.

[33] L. Tang and H. Liu, Relational learning via latent social dimensions, in *Proc. KDD'09*, Paris, France, 2009, pp. 817-826.

[34] L. Backstrom, R. Kumar, C. Marlow, J. Novak, and A. Tomkins, Preferential behavior in online groups, in *Proc. WSDM'08*, Stanford, CA, USA, 2008, pp. 117-128.

[35] N. Eagle, A.S. Pentland, and D. Lazer, Inferring social network structure using mobile phone data, *PNAS,* vol. 106, no. 36, pp. 1-24, 2009.

[36] C. H. Tan, J. Tang, J. M. Sun, Q. Lin, and F. J. Wang, Social action tracking via noise tolerant time-varying factor graphs, in *Proc. KDD'10*, Washington, DC USA, 2010, pp. 1049-1058.

[37] J. N. Rosenquist, J. Murabito, J. H. Fowler, and N. A. Christakis, The spread of alcohol consumption behavior in a large social network, *Annals of Internal Medicine,* vol. 152, no. 7, pp. 426-433, 2010.

[38] J. H. Fowler and N. A. Christaki, Cooperative behavior cascades in human social networks, *PNAS,* vol. 107, no. 12, pp. 1-52, Mar. 2010.

[39] G. Kossinets and D. J. Watts, Empirical analysis of an evolving social network, *Science,* vol. 311, no. 5757, pp. 88-90, Jan. 2006.

[40] J. P. Onnela, J. Saramaki, J. Hyvonen, G. Szabo, D. Lazer, K. Kaski, J. Kertesz, and A. L. Barabasi, Structure and tie strengths in mobile communication networks, *PNAS,* vol. 104, no. 18, pp. 7332-7336, 2006.

[41] A. Ahmed and E. P. Xing, Recovering time-varying networks of dependencies in social and biological studies, *PNAS,* vol. 106, no. 29, pp. 11878-11883, 2009.

[42] T. La Fond and J. Neville, Randomization tests for distinguishing social influence and homophily effects, in *Proc. WWW'10*, Raleigh, USA, 2010, pp. 601-610.

[43] M. Gomez-Rodriguez, J. Leskovec, and A. Krause, Inferring networks of diffusion and influence, in *Proc. KDD'10*, Washington, DC, USA, 2010, pp. 1019-1028.

[44] Y. Z. Sun and J. W. Han, Meta-path-based search and mining in heterogeneous information networks, *Tsinghua Science and Technology,* vol. 18, no. 4, pp. 329-338, Aug. 2013.

Jibing Gong received his PhD from Institute of Computer Technology, Chinese Academy of Sciences in 2012. He is an associate professor in School of Information Science and Engineering at Yanshan University. His research interests include data mining, social network and machine learning.

Jie Tang received his PhD from Tsinghua University in 2006. He is an associate professor at the Department of Computer Science and Technology, Tsinghua University. His main research interests include data mining algorithms and social network theories. He has been visiting scholar at Cornell University, Chinese University of Hong Kong, Hong Kong University of Science and Technology, and Leuven University.

A. C. M. Fong received his PhD from University of Auckland in 1998. He is with the University of Glasgow in Singapore. He has published papers widely in the areas of data mining and communications. He is a senior member of the IEEE and a member of the IEEE Computer Society.