Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence

Learning Temporal Dynamics of Behavior Propagation in Social Networks*

Jun Zhang
Department of Computer Science and Technology
Tsinghua University
Beijing 100084, P. R. China

Chaokun Wang and **Jianmin Wang**
School of Software
Tsinghua University
Beijing 100084, P. R. China

Abstract

Social influence has been widely accepted to explain people's cascade behaviors and further utilized in many related applications. However, few of existing work studied the direct, microscopic and temporal impact of social influence on people's behaviors in detail. In this paper we concentrate on the behavior modeling and systematically formulate the family of behavior propagation models (BPMs) including the static models (BP and IBP), and their discrete temporal variants (DBP and DIBP). To address the temporal dynamics of behavior propagation over continuous time, we propose a continuous temporal interest-aware behavior propagation model, called CIBP. As a new member of the BPM family, CIBP exploits the continuous-temporal functions (CTFs) to model the fully-continuous dynamic variance of social influence over time. Experiments on real-world datasets evaluated the family of BPMs and demonstrated the effectiveness of our proposed approach.

The society is such highly-connected that our behaviors will inevitably affect and meanwhile be affected by others around us (Kelman 1958). The social influence works in either an implicit way via influencing one's interests and emotions, or an explicit way via propagating behaviors from one to another. As shown in the literature, individuals usually have incentives to directly adopt the behaviors of their neighbors in the network (Easley and Kleinberg 2010). We refer to such phenomenon under the explicit and direct social influence as *behavior propagation*.

The behavior propagations are fundamental for the famous viral marketing (Subramani and Rajagopalan 2003) which promotes products via word-of-mouth recommendation and expects to cause cascade adoption behaviors. However, most related work (Chen, Wang, and Wang 2010; Goyal, Lu, and Lakshmanan 2011; Jiang et al. 2011; Chen, Lu, and Zhang 2012) focused on the optimal design of marketing strategies by selecting most influential seed users effectively and efficiently based on the given social network, but didn't study the modeling of propagation-driven behaviors. In our previous work (Zhang et al. 2013b; 2013c), we have studied the direct impact of social influence on people's friend-making behaviors in terms of friendship propagation, and proposed cascade (discrete) approaches to modeling the temporal behaviors by discretizing the continuous time. However, the network is always evolving and the social influence is continuously changing. Discrete approaches are unable to capture the continuous dynamics due to its discrete nature. Thus, an interesting and challenging question is: *Can we depict and learn the fully-continuous temporal dynamics of social influence for behavior propagation?*

In this work, we generalize our previous work and formulate the family of behavior propagation models. More importantly, we address the problem of modeling and learning the continuous temporal dynamics of behavior propagations over the evolving social networks, and to this end we design continuous behavior propagation models.

Why Are Continuous Models Needed?

Intuitively and ideally, a sufficiently fine-grained time discretization may lead to nearly-continuous modeling. Nevertheless, this is impractical because shortening the span of each time interval will inevitably reduce the observed behavior data in each interval and consequently lead to immature or over-fitted models. It also results in the increase of both the number of intervals and the complexity of model learning. Unlike previous work which usually regards the network as static or discrete, we propose to directly model the continuous variation of behavior propagations over continuous-evolving social networks. The proposed native continuous models are remarkably distinguished from the discrete ones by the following inherent merits.

First of all, among the major differences between the discrete and continuous models is that the discrete model can only consider the partial orders among user behaviors (orders are wiped out for behaviors within the same interval) while the continuous model considers the total orders. As one's previous behaviors may influence her interests and subsequent behaviors, we believe the order of behaviors should be carefully considered.

Second, the performance of discrete models is heavily dependent on the quality of discretization. Factually, it's often difficult for us to choose the most appropriate discretization granularity, which varies for different users and different scenarios. In contrast, the continuous approach doesn't treat each interval discretely, but looks at the whole picture.

*Corresponding authors: Chaokun Wang and Jianmin Wang

Third, discrete models can only reveal the dynamics in the past, but not for the future. On the contrary, continuous models can see the long-range dependency of the social influence over time and capture the change tendency in the future. In this sense, continuous models can learn user profiles better, and further predict time-aware behaviors in the future more precisely.

Contributions

To the best of our knowledge, this is the first time the fully-continuous temporal dynamics of behavior propagation is modeled carefully and the family of behavior propagation models (BPMs) is formulated systematically. The main contributions of this paper can be summarized as follows.

Firstly, we formulate the static and discrete behavior propagation models including BP/IBP and DBP/DIBP, and propose the fully-continuous temporal dynamic behavior propagation model, **CBP** and **CIBP**, in which the social influence is modeled as a continuous function of the time. Besides, in CIBP the behaviors are modeled as the co-effect of both behavior propagations and personal interests. These models form the *family of behavior propagation models* (BPMs).

Furthermore, we introduce the continuous-temporal functions (CTFs) to depict the dynamics of social influence over continuous time utilizing the flexible mixture of basis functions. We study how to select the appropriate CTF and control its complexity in experimental study.

Moreover, we conduct extensive experiments on real-world datasets to evaluate the performance of our proposed CIBP model and the whole behavior propagation model family. Results show that the CIBP outperforms both the state-of-the-art static and dynamic models, and can improve the performance of behavior prediction significantly.

Related Work

There is a popular belief that friends may influence each other and thus tend to exhibit similar behaviors (Chua, Lauw, and Lim 2011). The social relationships, recently, have been found beneficial for item recommendation and behavior prediction tasks (Konstas, Stathopoulos, and Jose 2009; Ma 2013; Cheng et al. 2012), and approaches have been proposed to incorporate the social relationships into predictive models (Ma, King, and Lyu 2009; Ma, Lyu, and King 2009; Ma et al. 2011; Ma 2013; Zhao et al. 2013). A major line of research focuses on the better and faster viral marketing based on the social influence (Chen, Wang, and Wang 2010; Goyal, Bonchi, and Lakshmanan 2011; Goyal, Lu, and Lakshmanan 2011; Guo et al. 2013; Zhang et al. 2013d; Jiang et al. 2011; Chen, Lu, and Zhang 2012), but most of them focused on the cascade behaviors over the static network topology. Other researchers took the quantitative analysis of social influence recently (Tang et al. 2009; Goyal, Bonchi, and Lakshmanan 2010; Liu et al. 2010; Wang et al. 2012; Bakshy et al. 2011; Cui et al. 2011). Furthermore, Ye et al. (2012) and Chen et al. (2013) respectively developed behavior models incorporating the social influence on peoples' interests, and we studied the direct friendship propagation using behavior models in our previous work (Zhang et al. 2013a; 2013b; 2013c) . In this paper, we propose to study the variance of social influence from the perspective of continuous modeling, which can capture both the microscopic and macroscopic variation of social influence and overcome the deficiencies of static or discrete approaches.

Temporal modeling has attracted attentions in other areas. Among them are the temporal topic models for documents. The DTM (Blei and Lafferty 2006) and DMM (Wei, Sun, and Wang 2007) are two famous dynamic topic models for discrete data based on the Markov assumptions over state transitions in the time domain. Unlike them, the TOT model (Wang and McCallum 2006) parameterizes a continuous distribution over time associated with each topic, but assumes the word distribution in each topic stays invariant. On the contrary, the cDTM (Wang, Blei, and Heckerman 2008) captures the continuous variance of the word distributions of each topic using the Markov chain modeling. The collaborative filtering with temporal dynamics (Koren 2009) has also been studied. However, none of them studied the temporal dynamics of behavior propagations and social influence, which is quite different from the traditional topic analysis.

Behavioral Propagation Models

In our previous work (Zhang et al. 2013a; 2013b), we proposed a probabilistic generative model, called LaFT-LDA, to model the friend-making behaviors based on the transitivity of friendship. Here we extend this approach as the static **BP** model for general behavior propagation modeling. Let $U = \{u_1, u_2, \cdots, u_N\}$ and $V = \{v_1, v_2, \cdots, v_M\}$ be the user set and item set, respectively. The BP model is a kind of behavior model which assumes the behavior is caused by the behavior propagation via friends (or intermediaries). For each user u and one of her adopted item v, as illustrated in Fig. 1(a), the BP model explains this behavior using a two-step process: u selects an intermediary y from $Y(u)$, i.e. the friends set of u, and then adopts the item v which is already adopted and then recommended by y. Based on the assumption that u and v are independent conditioned on y, the probability of the behavior (u, v), i.e. user u adopts item v, can be written as:

$$P(v|u) = \sum_{y \in Y(u)} P(v|y)P(y|u). \tag{1}$$

As the personal interests are shown beneficial for behavior prediction in our later work (Zhang et al. 2013c), here we incorporate the interests into the BP model to build the interest-aware behavior propagation model, i.e. **IBP** model. As illustrated in Fig. 1(b), the IBP model considers the co-effect of personal interests and behavior propagations by combining the classic pLSA (Hofmann 1999) and BP model. Let $Z = \{z_1, z_2, \cdots, z_K\}$ be the set of latent interest areas. Now the generative probability of the aforementioned item-adoption behavior denoted by (u, v) is defined as:

$$\begin{aligned} P(v|u) &= \sum_{y \in Y(u)} \sum_{z \in Z} P(v|z, y)P(y|u)P(z|u) \\ &= \sum_{y \in Y(u)} P(v|y)P(y|u) \cdot \sum_{z \in Z} P(v|z)P(z|u), \end{aligned} \tag{2}$$

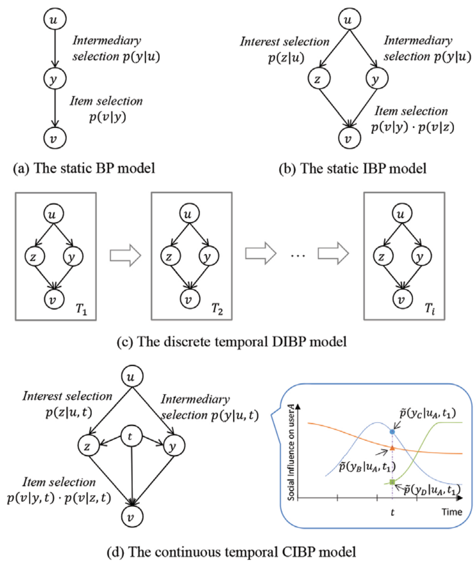

(a) The static BP model (b) The static IBP model (c) The discrete temporal DIBP model (d) The continuous temporal CIBP model

Figure 1: The family of behavior propagation models.[1]

where y and z are assumed independent.

Both the BP and IBP models assume the world is static. However, as one may make new friends, and one's trusts and interests are also varying, temporal models are expected to capture such variation. A popular method is to discretize the continuous time into periods, and infer the models for each period based on the Markov assumption (Zhang et al. 2013b; 2013c; Chen, Hsu, and Lee 2013). The DIBP model, shown in Fig. 1(c), is the discretized version of IBP model, and similarly we can draw the discrete DBP model corresponding to BP. As we discussed before, the discrete models face many difficulties and drawbacks. In the following section, we'll present our idea on modeling the fully-continuous temporal dynamics of behavior propagation.

Temporal Dynamics of Behavior Propagation

In this section, we focus on the continuously temporal dynamic interest-aware behavior propagation model, called CIBP based on the IBP and DIBP. The interest-unaware counterpart, i.e. CBP based on BP and DBP, can be derived by removing the interest factor easily.

The CIBP Model

The static IBP model includes three probabilistic selection steps: intermediary selection, interest selection and item selection from corresponding probabilistic distributions, which are assumed static w.r.t. time. In this paper, we argue that all these distributions are varying with time. we propose to use the *continuous-temporal function* (CTF) to depict the fully-continuous dynamic change of the related factors. Specifically, we use the intermediary preference CTF $f_{u,y}(t;\theta)$ and the interest preference CTF $g_{u,z}(t;\phi)$ to express one's preference for intermediaries and interest areas, and use the item favoritism CTF $r_{y,v}(t;\psi)$ and the item popularity CTF $s_{z,v}(t;\chi)$ to denote the recommendations of each item from intermediaries and its popularity in each interest areas. Each function is a real-valued single-variable continuous function over time t defined for each pair of objects. By introducing the CTFs, we can capture both the microscopic variation and global tendency of the hidden distributions, without concerning the granularity of discretization.

Without loss of generality, we'll discuss the design of the CTFs by taking the example of the intermediary preference CTF $f(t;\theta)$. In this study, we adopt the linear basis function model (Bishop 2006) that involves a linear combination of fixed nonlinear functions of the input variable t, of the form

$$f(t;\theta) = \sum_{j=0}^{J} w_j \vartheta_j(t;\theta) = \mathbf{w}^{\mathbf{T}} \vartheta(\mathbf{t};\theta), \tag{3}$$

where each $\vartheta_j(t;\theta)$ is a basis function. Usually, we set the first term as the static time-invariant factor by defining an additional dummy basis function $\vartheta_0(t;\theta) = 1$. The subsequent items are time-variant factors w.r.t. time. In this study we consider 6 types of popular basis functions, including the linear, polynomial, quadratic, Gaussian, sigmoid and exponential function.

Let $Y_t(u)$ and $Y_{\bar{t}}(u)$ be the set of friends (i.e. intermediaries) of user u exactly before and at time t, respectively. Let Z_t and V_t be the set of *available* interest areas and items exactly before time t, respectively. $V_t(y)$ denotes the items adopted by y before time t. Given the CTFs, we define the corresponding continuously temporal probability density functions (PDFs):

$$P(y|u,t,\Theta) = \begin{cases} \frac{f_{u,y}(t;\theta)}{\sum_{y' \in Y_t(u)} f_{u,y'}(t;\theta)} & \text{if } y \in Y_t(u) \\ 0 & \text{otherwise} \end{cases}, \tag{4}$$

$$P(z|u,t,\Theta) = \begin{cases} \frac{g_{u,z}(t;\phi)}{\sum_{z' \in Z_t} g_{u,z'}(t;\phi)} & \text{if } z \in Z_t \\ 0 & \text{otherwise} \end{cases}, \tag{5}$$

$$P(v|y,t,\Theta) = \begin{cases} \frac{r_{y,v}(t;\psi)}{\sum_{v' \in V_t(y)} r_{y,v'}(t;\psi)} & \text{if } v \in V_t(y) \\ 0 & \text{otherwise} \end{cases}, \tag{6}$$

$$P(v|z,t,\Theta) = \begin{cases} \frac{s_{z,v}(t;\chi)}{\sum_{v' \in V_t} s_{z,v'}(t;\chi)} & \text{if } v \in V_t \\ 0 & \text{otherwise} \end{cases}. \tag{7}$$

In such formulation, we allow all the factors including intermediaries, interest areas and items to change at any time, and one's behavior at any time is considered as the interplay of all current available factors by that time.

Now we can incorporate the continuous time fully into the basic static IBP model to build the CIBP model using the temporal PDFs defined in terms of CTFs. As illustrated in Fig. 1(d), CIBP assumes the following generative process for each item-adoption behavior (u, v) at time t:

1. Sample an intermediary $y \in Y_t(u)$ according to the temporal dynamic intermediary selection probability $P(y|u,t)$ at time t;
2. Sample an interest area $z \in Z_t$ according to the temporal dynamic interest selection probability $P(z|u,t)$ at time t;

[1]The DBP and CBP models are omitted here due to space limitations, but they can be derived from the corresponding DIBP and CIBP models easily by removing the interest factor z.

3. Sample an item $v \in V_t$ according to the joint temporal dynamic item selection probability $P(v|y,t) \cdot P(v|z,t)$ at time t.

Thus the probability of the behavior that u adopts an item v at time t can be written as:

$$P(v|u,t,\Theta) = \sum_{y \in Y_t(u)} P(v|y,t,\Theta) P(y|u,t,\Theta) \cdot \sum_{z \in Z_t} P(v|z,t,\Theta) P(z|u,t,\Theta), \tag{8}$$

where Θ is the collection of parameters of the model.

Model Inference

The parameters $\Theta = \{\theta, \phi, \psi\}$ of CIBP can be obtained by maximizing the likelihood. Let $\mathbb{V}_u = \{\langle v_1, t_1\rangle, \langle v_2, t_2\rangle, \cdots, \langle v_{m(u)}, t_{m(u)}\rangle\}$ be the behavior log of user u, and $\mathbb{V} = \bigcup_{u \in U} \mathbb{V}_u$. Let $\mathbb{Y}$ and $\mathbb{Z}$ record the latent intermediary and interest area for each behavior in $\mathbb{V}$. Here, instead of looking at the likelihood for the observed data (U and $\mathbb{V}$), we consider the log-likelihood function of the complete data ($U, \mathbb{Y}, \mathbb{Z}$ and $\mathbb{V}$):

$$\begin{aligned} &\log L(\Theta; U, \mathbb{Y}, \mathbb{Z}, \mathbb{V}) = \log P(\mathbb{Y}, \mathbb{Z}, \mathbb{V}|U, \Theta) \\ &= \sum_{u \in U} \sum_{\langle v,t\rangle \in \mathbb{V}_u} \Big(\log P(y|u,t,\Theta) + \log P(z|u,t,\Theta) \\ &\quad + \log P(v|y,t,\Theta) + \log P(v|z,t,\Theta) \Big). \end{aligned} \tag{9}$$

It's intractable to learn the parameters directly for the existence of hidden variables y and z. In this study we employ the expectation-maximization (EM) algorithm for model inference. Instead of maximizing the log-likelihood defined in Eq. 9, the EM algorithm here maximizes the expectation of the log-likelihood iteratively.

In the *E-step*, we calculate the expectation of the log-likelihood given current estimation Θ^g. For each observable behavior (u, v) at time t, we compute the posterior distribution of hidden variables y and z, given the data and the current values of parameters:

$$\begin{aligned} &P(y,z|u,v,t,\Theta^g) \\ &= \frac{P(y,z|u,t,\Theta^g) \cdot P(v|y,z,t,\Theta^g)}{P(v|u,t,\Theta^g)} \\ &= \frac{P(y|u,t,\Theta^g) \cdot P(z|u,t,\Theta^g) \cdot P(v|y,t,\Theta^g) \cdot P(v|z,t,\Theta^g)}{\sum_{y' \in Y_t(u)} P(v|y',t,\Theta) P(y'|u,t,\Theta) \cdot \sum_{z' \in Z_t} P(v|z',t,\Theta) P(z'|u,t,\Theta)}. \end{aligned} \tag{10}$$

Then we compute the expectation of the log-likelihood given Θ^g:

$$\begin{aligned} &Q(\Theta,\Theta^g) = \mathbb{E}\left(\log L(\Theta; U, \mathbb{Y}, \mathbb{Z}, \mathbb{V})|\Theta^g\right) \\ &= \sum_{u \in U} \sum_{\langle v,t\rangle \in \mathbb{V}_u} \sum_{y \in Y_t(u)} \sum_{z \in Z_t} \Bigg(\log \frac{f_{u,y}(t;\theta)}{\sum_{y' \in Y_t(u)} f_{u,y'}(t;\theta)} \\ &+ \log \frac{r_{y,v}(t;\psi)}{\sum_{v' \in V_t(y)} r_{y,v'}(t;\psi)} + \log \frac{g_{u,z}(t;\phi)}{\sum_{z' \in Z_t} g_{u,z'}(t;\phi)} \\ &+ \log \frac{s_{z,v}(t;\chi)}{\sum_{v' \in V_t} s_{z,v'}(t;\chi)} \Bigg) \cdot P(y,z|u,v,t,\Theta^g). \end{aligned} \tag{11}$$

In the *M-step*, we try to maximize the above expectation to get the new estimation for Θ. However, the existence of the summations in the denominators makes the optimization complex. To avoid the evaluation of summation each time, we introduce the probabilistic regularization to the CTFs. In this way, the Q-function can be rewritten as:

$$\begin{aligned} Q(\Theta,\Theta^g) = &\sum_{u \in U} \sum_{\langle v,t\rangle \in \mathbb{V}_u} \sum_{y \in Y_t(u)} \sum_{z \in Z_t} \Big(\log f_{u,y}(t;\theta) + \log r_{y,v}(t;\phi) \\ &+ \log g_{u,z}(t;\psi) + \log g_{z,v}(t;\chi) \Big) \cdot P(y,z|u,v,t,\Theta^g) \\ &+ \tau_1 \sum_{u \in U} \sum_t \left(1 - \sum_{y' \in Y_t(u)} f_{u,y'}(t;\theta)\right) \\ &+ \tau_2 \sum_{y \in Y} \sum_t \left(1 - \sum_{v' \in V_t(y)} r_{y,v'}(t;\psi)\right) \\ &+ \tau_3 \sum_{u \in U} \sum_t \left(1 - \sum_{z' \in Z_t} g_{u,z'}(t;\phi)\right) \\ &+ \tau_4 \sum_{z \in Z} \sum_t \left(1 - \sum_{v' \in V_t} s_{z,v'}(t;\chi)\right), \end{aligned} \tag{12}$$

where τ_1, τ_2, τ_3 and τ_4 are the Lagrange multipliers.

As each CTF is also a combination of basis functions, it's rather complicated to obtain the parameters for these functions by maximizing the expectation directly. Here, we take an alternative approach by learning the optimal function value at each time t first and then estimating the best parameters for the functions.

Taking the intermediary preference CTF as example, for each user u and one of her friend y, we take the derivative of the Q-function at any time t w.r.t. $f_{u,y}(t;\theta)$:

$$\frac{\partial Q(\Theta,\Theta^g)}{\partial f_{u,y}(t;\theta)} = \sum_{\langle v,t\rangle \in \mathbb{V}_u} \sum_{z \in Z_t} \frac{P(y,z|u,v,t,\Theta^g)}{f_{u,y}(t;\theta)} - \tau_1 = 0, \tag{13}$$

and we get

$$\tau_1 f_{u,y}(t;\theta) = \sum_{\langle v,t\rangle \in \mathbb{V}_u} \sum_{z \in Z_t} P(y,z|u,v,t,\Theta^g). \tag{14}$$

Summing both sides over y, we get:

$$\tau_1 = \sum_{\langle v,t\rangle \in \mathbb{V}_u} \sum_{z \in Z_t} P(z|u,v,t,\Theta^g). \tag{15}$$

Thus the optimal value of $f_{u,y}(t;\theta)$ is given as:

$$\tilde{f}_{u,y}[t] = \frac{\sum_{v \in \mathbb{V}_{u,t}} \sum_{z \in Z_t} P(y,z|u,v,t,\Theta^g)}{\sum_{v \in \mathbb{V}_{u,t}} \sum_{z \in Z_t} P(z|u,v,t,\Theta^g)}. \tag{16}$$

Similarly, we get the optimal values for other CTFs. With the learned values w.r.t. the time series, the parameters θ, ϕ, ψ and χ for the CTFs can be estimated easily using gradient descent or Newton methods. We omit the details due to the limit of space.

By repeating the E-step and M-step, the EM algorithm improves the estimations of model parameters iteratively until they converge to a local log-likelihood maximum.

Experimental Evaluation

In this section, we evaluate the effectiveness of our model through comprehensive experiments on real-world datasets.

We construct 5 real-world datasets from real-world academic collaborative social networks, including **sci-comp** (Science Computation), **comp-edu** (Computer Education), **simu** (Simulation) and **sec-priv** (Security & Privacy) from Microsoft Academic Search[2], and **comp-ling** (Computational Linguistics) from the ACL Anthology Network[3]. In the above datasets, we consider 3 types of behaviors: *uw* (use a specific word in the paper title), *ct* (cite a specific paper) and *ca* (co-author with another researcher). For the *uw* behavior we removed the stopwords from the paper titles. The statistics of the datasets are shown in Tbl. 1.

Table 1: The statistics of our datasets.

dataset	# users	# pub.	# words	# uw	# ct	# ca
sci-comp	181,035	178,363	69,803	2,165,143	3,602,373	665,512
comp-edu	62,071	47,308	24,451	640,094	561,754	214,594
simu	38,832	29,200	15,941	398,147	587,567	132,718
sec-priv	65,884	63,122	29,493	827,426	1,600,772	279,488
comp-ling	15,835	19,424	10,709	309,637	243,102	126,508

Evaluation Methodology

We evaluate our models by examining their performance in temporal behavior prediction. For each dataset, we collected the data in 1981–2000 for model training, and the next 5 years were taken for testing. We only considered the individuals who had at least 10 friends (co-authors) by 2000.

In this study, we evaluate the performance of all behavior propagation modes, including the static **BP** and **IBP** models, discrete-temporal **DBP** and **DIBP** models and the proposed continuous-temporal **CBP** and **CIBP** models. The hierarchy among them is shown in Fig. 2.

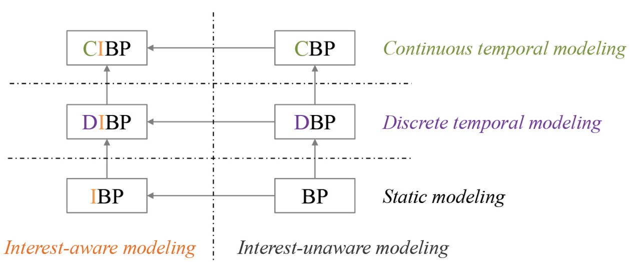

Figure 2: The hierarchy of BPMs family.

We found our data quite sparse, and thus split the data by year for discrete models. To be fair, one year is taken as a basic time unit for continuous models. For each user at each year, only the items had ever adopted by her friends are considered for training and test because our focus is the direct behavior propagation. The items adopted by the user are positive instances and others are negative. Here we only predict the occurrence of each behavior and don't consider the number of the occurrences. We evaluate their prediction performance using MAP (Mean Average Precision) and AUC (Area Under the ROC Curve).

[2] http://academic.research.microsoft.com

[3] http://clair.eecs.umich.edu/aan

Table 2: The overall prediction performance for the *ca* behavior on all datasets.

Dataset	Metric	Static Models		Discrete Models		Continuous Models	
		BP	IBP	DBP	DIBP	CBP	CIBP
sec-priv	MAP	0.4334	0.4435	0.4543	0.4660	0.4657	**0.4723**
	AUC	0.5676	0.6018	0.6020	0.6262	0.6167	**0.6391**
comp-ling	MAP	0.4495	0.4582	0.4725	0.4787	0.4742	**0.4996**
	AUC	0.6018	0.6466	0.6218	0.6673	0.6346	**0.6901**
simu	MAP	0.4064	0.5126	0.4653	**0.5313**	0.4794	0.5189
	AUC	0.5967	0.6790	0.6206	0.6910	0.6446	**0.6915**
sci-comp	MAP	0.4026	0.4139	0.4359	0.4493	0.4519	**0.4717**
	AUC	0.6533	0.6345	0.6589	0.6493	**0.6724**	0.6632
comp-edu	MAP	0.3840	0.4142	0.4093	0.4374	0.4300	**0.4572**
	AUC	0.5412	0.5269	0.5643	0.5303	**0.6150**	0.5330

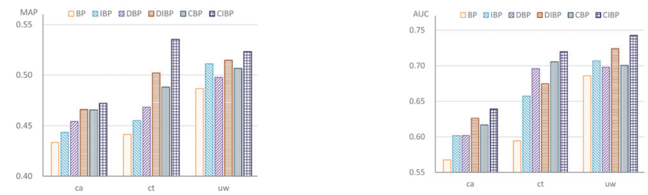

(a) Performance in MAP (b) Performance in AUC

Figure 3: The performance comparison of models for all behaviors prediction on **sec-priv**.

Overall Performance Comparison

Firstly, we present the average performance of the models for each behavior prediction in the 5 testing periods. For the limit of space, we illustrate the results for *ca* behavior on each datasets in Tbl. 2, and results for all of the three behaviors on **sec-priv** in Fig. 3.

First of all, we see the interest-aware models, including IBP, DIBP and CIBP, outperform their corresponding interest-unaware models BP, DBP and CBP, respectively. This reveals that the behavior propagation is usually influenced by personal interests and incorporating the interest into the behavior propagation models can really improve the prediction performance.

Furthermore, the temporal models usually perform better than the static BP and IBP models. By capturing the temporal change, they can get more precise modeling on user profiles and behavior propagations, and achieve better prediction performance.

Moreover, as expected, we observe that the performances of continuous models exceed those of discrete ones. Unlike DBP and DIBP which split the time into discrete periods and assume the evolving models along the periods form a Markov chain, the continuous CBP and CIBP depict the continuous variance of social influence and personal interests over time. Besides, CIBP performs best in most of our experiments by modeling the continuously temporal dynamics of both social influence and personal interests.

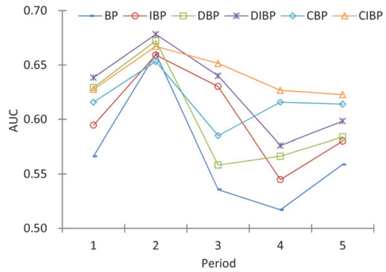

(a) Performance comparison on **sec-priv**

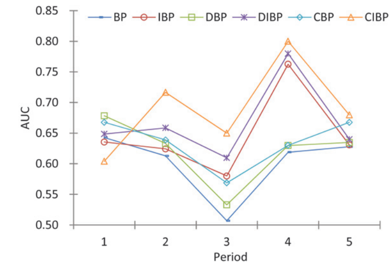

(b) Performance comparison on **comp-ling**

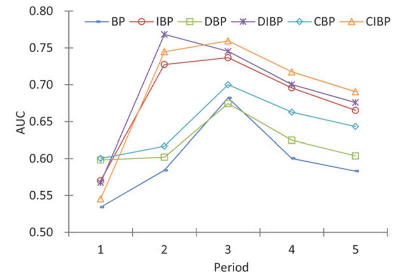

(c) Performance comparison on **simu**

Figure 4: The temporal prediction performance of all models for the *ca* behavior prediction.

Temporal Performance Comparison

Now we evaluate the predictive power of our model for temporal behaviors by observing the performance variance of the models in each testing period. Due to the limit of space, we present the results for the *ca* behavior on three of our datasets in Fig. 4. Comparing with discrete and static methods, our proposed CBP and CIBP exhibit not only better performance but also better stability. Because the static and discrete models can only result in fixed models by the last training period, their prediction capability will inevitably decrease as time goes by. On the contrary, the CBP and CIBP can capture the change tendency with the continuous CTFs and adjust their predictions according to time, and consequently achieve better performance.

CTF Selection

Among the key components of the CBP and CIBP models are the CTFs, which depict the temporal dynamics of each factor. We're interested in that what function can explain the temporal dynamics of social influence better and thus perform better in the prediction tasks.

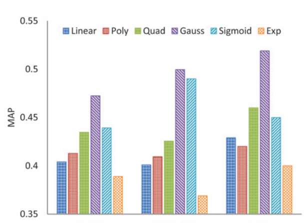

(a) Performance in MAP

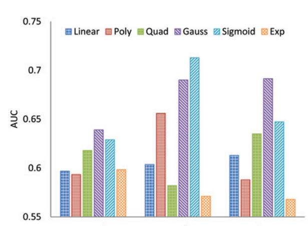

(b) Performance in AUC

Figure 5: The performance comparison of CIBP models with different CTFs for the *ca* behavior prediction on **sec-priv**.

We report the average performance of the CIBP model using different CTFs for the *ca* behavior on **sec-priv** in Fig. 5. Results reveal that the *Gaussian function* and *sigmoid function* perform better than others on most occasions. This is reasonable because the social influence and personal interests usually change slowly and gradually. Compared to others, the two winner functions are smoother and can model

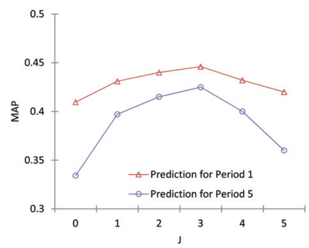

(a) Performance in MAP

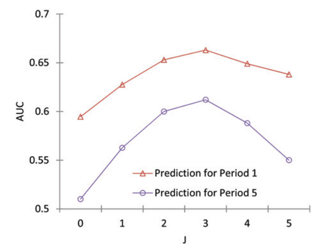

(b) Performance in AUC

Figure 6: The performance variance of CIBP with different number (J) of basis functions in the CTFs for the *uw* behavior prediction on **sec-priv**.

the gentle variance with less risk of over-fitting. The variance depicted by them usually becomes more and more slow with time goes by, and nearly stay at some value in far future. This also conforms to the actual influence change. Besides, *Gaussian function* describes more complete lifecycle of social influence including ups and downs.

Performance vs. Complexity

Now we study the relation between the performance of CIBP and the complexity of the inside CTFs by observing how the prediction performance varies when adjusting the complexity parameter J.

Experimental results of the *ca* behavior prediction on **sec-priv** are shown in Fig. 6. When $J = 0$, our CIBP degenerates to the static IBP and performs worst. Increasing J improves the performance at first and achieves the peak at $J = 3$. After that the continuous increase of J leads to decrease of performance due to the overfitting problem. The prediction for the 5th test period suffers much more than that for the 1st one from that. We note that increasing J also requires longer time for model inference. In our experiments we find the most appropriate value for J is 2 or 3.

Conclusion

As a popular phenomenon in social networks, the behavior propagation is fundamental for many upper-layer applications. Existing research work usually addresses the behavior propagation based on the assumption that the social influence is static all the time or at least in a period. In this paper, we argue the social network is always evolving and the social influence is continuously changing. We firstly formulate the family of behavior propagation models (BPMs) systematically, and then present a continuous temporal generative model, i.e. CIBP, to address the temporal dynamics of behavior propagations. It's a new member of BPMs and designed to capture the continuous varying dynamics of interest-aware behavior propagations using elaborate CTFs. We conduct extensive experiments to evaluate the performance of BPMs for behavior prediction. Evaluation results show that the proposed CIBP outperforms other static or discrete behavior propagation models significantly.

Acknowledgments

We thank Heran Lin and Xiang Ying for collecting the experimental datasets. This work is supported in part by the National Natural Science Foundation of China (No. 61373023, No. 61170064, No. 61133002) and the National High Technology Research and Development Program of China (No. 2012AA011002).

References

Bakshy, E.; Hofman, J. M.; Mason, W. A.; and Watts, D. J. 2011. Everyone's an influencer: Quantifying influence on twitter. In *WSDM '11*, 65–74.

Bishop, C. M. 2006. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag New York, Inc. chapter 3, 138–147.

Blei, D. M., and Lafferty, J. D. 2006. Dynamic topic models. In *ICML '06*, 113–120.

Chen, W.; Hsu, W.; and Lee, M. L. 2013. Modeling user's receptiveness over time for recommendation. In *SIGIR '13*, 373–382.

Chen, W.; Lu, W.; and Zhang, N. 2012. Time-critical influence maximization in social networks with time-delayed diffusion process. In *AAAI '12*.

Chen, W.; Wang, C.; and Wang, Y. 2010. Scalable influence maximization for prevalent viral marketing in large-scale social networks. In *KDD '10*, 1029–1038.

Cheng, C.; Yang, H.; King, I.; and Lyu, M. 2012. Fused matrix factorization with geographical and social influence in location-based social networks. In *AAAI '12*.

Chua, F. C. T.; Lauw, H. W.; and Lim, E.-P. 2011. Predicting item adoption using social correlation. In *SDM '11*.

Cui, P.; Wang, F.; Yang, S.; and Sun, L. 2011. Item-level social influence prediction with probabilistic hybrid factor matrix factorization. In *AAAI '11*.

Easley, D., and Kleinberg, J. 2010. *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*. Cambridge University Press. chapter 19, 497–536.

Goyal, A.; Bonchi, F.; and Lakshmanan, L. V. S. 2010. Learning influence probabilities in social networks. In *WSDM '10*, 241–250.

Goyal, A.; Bonchi, F.; and Lakshmanan, L. V. S. 2011. A data-based approach to social influence maximization. *Proc. VLDB Endow.* 5(1):73–84.

Goyal, A.; Lu, W.; and Lakshmanan, L. V. S. 2011. Celf++: Optimizing the greedy algorithm for influence maximization in social networks. In *WWW '11*, 47–48.

Guo, J.; Zhang, P.; Zhou, C.; Cao, Y.; and Guo, L. 2013. Personalized influence maximization on social networks. In *CIKM '13*, 199–208.

Hofmann, T. 1999. Probabilistic latent semantic analysis. In *UAI '99*, 289–296.

Jiang, Q.; Song, G.; Gao, C.; Wang, Y.; Si, W.; and Xie, K. 2011. Simulated annealing based influence maximization in social networks. In *AAAI '11*.

Kelman, H. 1958. Compliance, identification, and internalization: Three processes of attitude change. *Journal of Conflict Resolution* 2:51–60.

Konstas, I.; Stathopoulos, V.; and Jose, J. M. 2009. On social networks and collaborative recommendation. In *SIGIR '09*, 195–202.

Koren, Y. 2009. Collaborative filtering with temporal dynamics. In *KDD '09*, 447–456.

Liu, L.; Tang, J.; Han, J.; Jiang, M.; and Yang, S. 2010. Mining topic-level influence in heterogeneous networks. In *CIKM '10*, 199–208.

Ma, H.; Zhou, D.; Liu, C.; Lyu, M. R.; and King, I. 2011. Recommender systems with social regularization. In *WSDM '11*, 287–296.

Ma, H.; King, I.; and Lyu, M. R. 2009. Learning to recommend with social trust ensemble. In *SIGIR '09*, 203–210.

Ma, H.; Lyu, M. R.; and King, I. 2009. Learning to recommend with trust and distrust relationships. In *RecSys '09*, 189–196.

Ma, H. 2013. An experimental study on implicit social recommendation. In *SIGIR '13*, 73–82.

Subramani, M. R., and Rajagopalan, B. 2003. Knowledge-sharing and influence in online social networks via viral marketing. *Commun. ACM* 46(12):300–307.

Tang, J.; Sun, J.; Wang, C.; and Yang, Z. 2009. Social influence analysis in large-scale networks. In *KDD '09*, 807–816.

Wang, X., and McCallum, A. 2006. Topics over time: A non-markov continuous-time model of topical trends. In *KDD '06*, 424–433.

Wang, T.; Srivatsa, M.; Agrawal, D.; and Liu, L. 2012. Microscopic social influence. In *SDM '12*, 129–140.

Wang, C.; Blei, D.; and Heckerman, D. 2008. Continuous Time Dynamic Topic Models. In *UAI '08*.

Wei, X.; Sun, J.; and Wang, X. 2007. Dynamic mixture models for multiple time series. In *IJCAI'07*, 2909–2914.

Ye, M.; Liu, X.; and Lee, W.-C. 2012. Exploring social influence for recommendation: A generative model approach. In *SIGIR '12*, 671–680.

Zhang, J.; Wang, C.; Ning, Y.; Liu, Y.; Wang, J.; and Yu, P. S. 2013a. LaFT-Explorer: Inferring, Visualizing and Predicting

How Your Social Network Expands. In *KDD '13*, 1510–1513.

Zhang, J.; Wang, C.; Wang, J.; and Yu, P. S. 2013b. LaFT-Tree: Perceiving the Expansion Trace of One's Circle of Friends in Online Social Networks. In *WSDM '13*, 597–606.

Zhang, J.; Wang, C.; Yu, P. S.; and Wang, J. 2013c. Learning Latent Friendship Propagation Networks with Interest Awareness for Link Prediction. In *SIGIR '13*, 63–72.

Zhang, M.; Dai, C.; Ding, C.; and Chen, E. 2013d. Probabilistic solutions of influence propagation on social networks. In *CIKM '13*, 429–438.

Zhao, T.; Li, C.; Li, M.; Ding, Q.; and Li, L. 2013. Social recommendation incorporating topic mining and social trust analysis. In *CIKM '13*, 1643–1648.