

Social Restricted Boltzmann Machine: Human Behavior Prediction in Health Social Networks

NhatHai Phan	Dejing Dou	Brigitte Piniewski	David Kil
University of Oregon	University of Oregon	PeaceHealth Laboratories	HealthMantic Inc.
Eugene, OR, USA	Eugene, OR, USA	Vancouver, WA, USA	Los Altos, CA, USA
haiphan@cs.uoregon.edu	dou@cs.uoregon.edu	BPiniewski@peacehealthlabs.org	david.kil@healthmantic.com

Abstract—Modeling and predicting human behaviors, such as the activity level and intensity, is the key to prevent the cascades of obesity, and help spread wellness and healthy behavior in a social network. The user diversity, dynamic behaviors, and hidden social influences make the problem more challenging. In this work, we propose a deep learning model named *Social Restricted Boltzmann Machine (SRBM)* for human behavior modeling and prediction in health social networks. In the proposed SRBM model, we naturally incorporate *self-motivation*, *implicit* and *explicit* social influences, and *environmental events* together into three layers which are historical, visible, and hidden layers. The interactions among these behavior determinants are naturally simulated through parameters connecting these layers together. The contrastive divergence and back-propagation algorithms are employed for training the model. A comprehensive experiment on real and synthetic data has shown the great effectiveness of our deep learning model compared with conventional methods.

I. INTRODUCTION

Overweight and obesity are major risk factors for a number of chronic diseases, including diabetes, cardiovascular diseases, and cancers. Once considered a problem only in high-income countries, overweight and obesity are now dramatically on the rise in low- and middle-income countries. Recent studies have shown obesity can spread over the social network [3], bearing similarity to the diffusion of innovation [4] and word-of-mouth effects in marketing [6]. To reduce the risk of obesity related diseases, regular exercise is strongly recommended [14]. However, there have been few scientific and quantitative studies to elucidate how social relationships and personal factors may contribute to macro-level human behaviors.

The Internet is identified as an important source of health information and may thus be an appropriate delivery for health behavior interventions [13]. In addition, mobile devices can track and record the walking/jogging/running distance and intensity of an individual. Utilizing these technologies, our recent study, named YesiWell, was conducted in 2010-2011 as a collaboration between PeaceHealth Laboratories, SK Telecom Americas, and the University of Oregon to record daily physical activities, social activities (i.e., text messages, social games, events, competitions, etc.),

biomarkers, and biometric measures (i.e., cholesterol, BMI, etc.) for a group of 254 individuals. Physical activities were reported via a mobile device carried by each user. All users enrolled an online social network allowing them to make friends and communicate with each other. Users' biomarkers and biometric measures were recorded via monthly medical tests performed at our laboratories. The fundamental problems this study seeks to answer, which are also the key in understanding the determinants of human behaviors, are as follows:

- How could social communities affect individual behaviors?
- How could we leverage individual features and social communities to help predict the individual behavior?

It is nontrivial to determine how much impact social influences could have on someone's behavior. Our starting observation is that human behavior is the outcome of interacting determinants such as *self-motivation*, *social influences*, and *environmental events*. This observation is rooted in sociology and psychology, characterized as *human agency* in social cognitive theory [1]. An individual's level of self-motivation can be captured by learning the correlations between his or her historical and current characteristics. Modeling social influences in health social networks is challenging since they are categorized into *implicit* and *explicit* social influences [2]. In reality, all the users in our program have participated in many other social activities such as offline events, physical activity competitions, social games, etc. Thus, they have unreported relationships which can cause implicit social influences. The social context also is a potential cause of implicit social influences, since human behaviors usually respond to the changing of the social communities [1]. In addition, a user can be influenced by unacquainted users, since they participate in the same program. The hidden social influences are composed of unreported social relationships, unacquainted users, and the changing of social context. This problem is known as the hidden social influence of health social networks [2].

Explicit social influences demand a novel function to better capture the influences on individuals from their friends.

Unlike common developed online social networks (e.g., Facebook, etc.), our health social network is developed from scratch. Users do not have many connections initially. As time goes by, they will have more connections to other users. Thus, conventional social influence functions (e.g., information propagation cascade models [8], [17]) may not fit our context.

Motivated by the above challenging issues, we present a novel *Social Restricted Boltzmann Machine* (SRBM) for human behavior prediction, based on the well-known deep learning model, Restricted Boltzmann Machine (RBM) [19]. In the SRBM model, the human behavior determinants which are self-motivation, explicit and implicit social influences, and environmental events are integrated into three layers which are *historical*, *visible*, and *hidden* layers. Self-motivation can be captured by learning correlations between an individual's historical and current features. The effect of the implicit social influences on an individual is estimated by an aggregation function of the past of the social network. We define a new temporal smoothing and statistical function to capture explicit social influences on individuals from their neighboring users. By combining implicit and explicit social influences into a linear adaptive bias, we are able to model the social influences. The environment events, such as competitions, are integrated into the model as observed variables which will directly affect the user's behaviors. The effectiveness of our model is verified by experiments on both real-world and synthetic health social networks. Our main contributions are as follows:

- We introduce a new method for human behavior prediction based on the well-known deep learning model, Restricted Boltzmann Machine (RBM). The proposed SRBM model incorporates self-motivation, explicit and implicit social influences, and environmental events together.
- Our experimental assessment on both real and synthetic data confirms that our model is very accurate in human behavior prediction.

In Sec. 2, we introduce the RBM and related works. We present our SRBM model in Sec. 3. The experimental evaluation is in Sec. 4 and we conclude the paper in Sec. 5.

II. THE RBMS AND RELATED WORKS




Figure 1. The RBM.

The Restricted Boltzmann Machine (RBM) [19] is a deep learning structure that has a layer of visible units fully connected to a layer of hidden units but no connections within a layer (Figure 1). Typically, RBMs use stochastic binary units for both visible and hidden variables. The stochastic binary units of RBMs can be generalized to any distribution that falls in the exponential family

[22]. To model real-valued data, a modified RBM with

binary logistic hidden units and real-valued Gaussian visible units can be used. In Figure 1, v_i and h_j are respectively used to denote the states of visible unit i and hidden unit j . a_i and b_j are used to distinguish biases on the visible units and hidden units. The RBM assigns a probability to any joint setting of the visible units, \mathbf{v} and hidden units, \mathbf{h} :

$$p(\mathbf{v}, \mathbf{h}) = \frac{\exp(-E(\mathbf{v}, \mathbf{h}))}{Z} \quad (1)$$

where $E(\mathbf{v}, \mathbf{h})$ is an energy function,

$$E(\mathbf{v}, \mathbf{h}) = \sum_i \frac{(v_i - a_i)^2}{2\sigma_i^2} - \sum_j b_j h_j - \sum_{ij} \frac{v_i}{\sigma_i} h_j W_{ij} \quad (2)$$

where σ_i is the standard deviation of the Gaussian noise for visible unit i . In practice, fixing σ_i at 1 makes the learning work well. Z is a partition function which is intractable as it involves a sum over the (exponential) number of possible joint configurations: $Z = \sum_{\mathbf{v}, \mathbf{h}} \exp(-E(\mathbf{v}, \mathbf{h}))$. W is a weight matrix which bipartitely connects the hidden and visible layers together, e.g., W_{ij} links v_i and h_j . The conditional distributions (assuming $\sigma_i = 1$) are:

$$p(h_j = 1 | \mathbf{v}) = \sigma(b_j + \sum_i v_i W_{ij}) \quad (3)$$

$$p(v_i | \mathbf{h}) = \mathcal{N}(a_i + \sum_j h_j W_{ij}, 1) \quad (4)$$

where $\sigma(\cdot)$ is a logistic function, $\mathcal{N}(\mu, V)$ is a Gaussian.

Given a training set of state vectors, the weights and biases in a RBM can be learned following the gradient of contrastive divergence [5]. The learning rule are:

$$\Delta W_{ij} = \langle v_i h_j \rangle_d - \langle v_i h_j \rangle_r; \quad \Delta b_j = \langle h_j \rangle_d - \langle h_j \rangle_r \quad (5)$$

where the first expectation $\langle \cdot \rangle_d$ is based on the data distribution and the second expectation $\langle \cdot \rangle_r$ is based on the distribution of “reconstructed” data. The reconstructions are generated by starting a Markov chain at the data distribution. The hidden units can be updated by sampling Eq. 3, then updating the visible units by sampling Eq. 4.

To incorporate temporal dependencies into the RBM, the CRBM [20] adds autoregressive connections from the visible and hidden variables of an individual to his/her historical variables. The CRBM simulates well human motion in the single agent scenario. However, it cannot capture the social influences on individual behaviors in the multiple agent scenario. Li et al. [12] proposed a *ctRBM* model for link prediction in dynamic networks. The ctRBM simulates the social influences by adding the prediction expectations of local neighbors on an individual into a dynamic bias. However, the ctRBM cannot directly incorporate personal features with social influences to predict human behaviors. This is because the visible layer in the ctRBM does not take individual features as an input.

Regarding human behavior prediction, social behavior has been studied recently such as analysis of user interactions

on Facebook [21], activity recommendation [11], and user activity level prediction [25]. In [25], the authors focus on predicting users who have a tendency to decline their activity levels. This problem is known as churn prediction. Churn prediction aims to find users who will leave a network or a service. By finding such users, service providers could analyze the reasons and figure out the strategies to retain such users. Social churn prediction has been studied in different applications, including online social games [7], Q&A forums [23], etc. Our study is similar to social churn prediction because we also aim to predict the future activity level of a user. The users in these applications usually have simple user behaviors. Meanwhile our models enrich the application area by incorporating various personal features. The work most closely related to our study is the Socialized Gaussian Process Model (SGP) [18].

III. THE SRBM MODEL

In this section, we present our SRBM model for human behavior prediction. Given an *online social network* $G = \{U, E, \mathcal{F}\}$ where U is a set of all users and E is a set of edges, every edge $e_{u,m}$ exists in E if u and m are friends with each other in G ; otherwise $e_{u,m}$ does not exist. Each user has a set of individual features $\mathcal{F} = \{f_1, \dots, f_n\}$. The social network G grows from scratch over a set of time points $T = \{t_1, \dots, t_m\}$. To illustrate this we use $E_T = \{E_{t_1}, \dots, E_{t_m}\}$ to denote the topology of the network G over time, where E_t is a set of edges which have been made *until* time t in the network, and $\forall t \in T : E_t \subseteq E_{t+1}$. For each user, the values of individual features in \mathcal{F} also change over time. We denote the values of individual features of a user u at time t as \mathcal{F}_u^t . At each time point t , each user u associates with a binomial behavior $y_u^t \in \{0, 1\}$. y_u^t could be “decrease/increase exercise”, “inactive/active in exercise”. We will describe y_u^t clearly in our experimental result section.

Problem Formulation: Given the health social network in M timestamps $T_{data} = \{t - M, \dots, t\}$, we would like to predict the behavior of all the users in the next timestamps $t + 1$. In formal, given $\{\mathcal{F}_u^t, y_u^t, E_t | t \in T_{data}, u \in U\}$ we aim at predicting $\{y_u^{t+1} | u \in U\}$.

Figure 2 illustrates the proposed SRBM model. Our model includes three layers which are visible layer \mathbf{v} , hidden layer \mathbf{h} , and historical layer \mathcal{H} . Given a user, each visible variable v_i in the *visible layer* \mathbf{v} corresponds to an individual feature f_i at time t . All the visible variables of all the users in the previous N time intervals $\{t - N, \dots, t - 1\}$ (i.e., $N < M$) are included in a *historical layer*, denoted by \mathcal{H}_{t-} . In addition, all the variables in the historical layer are called *historical variables*. Obviously, we will have $|\mathcal{F}| \times |U| \times N$ historical variables. The hidden layer \mathbf{h} consists of $|\mathbf{h}|$ hidden variables. The issue now becomes connecting the three layers together and modeling the variables. Let us start with the modeling of self-motivation.




Figure 2. The SRBM model.

Self-motivation. The self-motivation is composed of many dimensions including attitudes, intentions, effort, and withdrawal which can all affect the motivation that an individual experiences [16]. In our YesiWell study, individual features are specially designed to capture self-motivation of each user. Some of the key measures are as follows:

- *Personal ability:* BMI, fitness, cholesterol, etc.
- *Attitudes:* the number of off-line events each user participates in, individual sending/receiving messages, the number of goals set and achieved, Wellness-score [9], etc. Wellness-score is a measure to evaluate how well a user lives their life. In essence, be active in social activities, setting and achieving more goals, and getting higher Wellness-score result in a better attitude of a user.
- *Intentions:* the number of competitions each user joins, the number of goals set, etc. They are intent to exercise, so they join the competitions and set more goals.
- *Effort:* the number of exercise days, walking/running steps, the distances, and speed.
- *Withdrawal:* BMI slope, Wellness-score slope, etc. The increase of BMI slope or decrease of Wellness-score [9] indicate not good signs in the self-motivation. The users may give up their progresses.

In order to model self-motivation of a user u , we first fully bipartite connect the hidden and visible layers via a *weight matrix* W (Figure 2). Then each visible variable v_i and hidden variable h_j will be connected to all the historical variables of u , denoted by $\mathcal{H}_{f,u,t-k}$ where $f \in \mathcal{F}$ and $k \in \{1, \dots, N\}$. These connections are presented by the two weight matrices A and B (Figure 2). Each historical variable $\mathcal{H}_{f,u,t-k}$ essentially is the state of feature f of the user u at time point $t - k$. Note that all the historical variables are treated as additional *observed* inputs. The hidden layer can learn the correlations among the features and the effect of the historical variables to capture self-motivation. This effect can be integrated into a dynamic bias:

$$\hat{b}_{j,t} = b_j + \sum_{k \in \{1, \dots, N\}} \sum_{f \in \mathcal{F}} B_{jf,u,t-k} \mathcal{H}_{f,u,t-k} \quad (6)$$

which includes a static bias b_j , and the contribution from

the past of the social network. B_j is a $|\mathbf{h}| \times |U| \times N$ weight matrix which summarizes the autoregressive parameters to the hidden variable h_j . This modifies the factorial distribution over hidden variables: b_j in Eq. 3 is replaced with $\hat{b}_{j,t}$ to obtain

$$p(h_{j,t} = 1 | \mathbf{v}_t, \mathcal{H}_{t_{<}}) = \sigma(\hat{b}_{j,t} + \sum_i v_{i,t} W_{ij}) \quad (7)$$

where $h_{j,t}$ is the state of hidden variable j at time t , the weight W_{ij} connects v_i and h_j . The self-motivation has a similar effect on the visible units. The reconstruction distribution in Eq. 4 becomes

$$p(v_{i,t} | \mathbf{h}_t, \mathcal{H}_{t_{<}}) = \mathcal{N}(\hat{a}_{i,t} + \sum_j h_{j,t} W_{ij}, 1) \quad (8)$$

where $\hat{a}_{i,t}$ is also a dynamic bias:

$$\hat{a}_{i,t} = a_i + \sum_{k \in \{1, \dots, N\}} \sum_{f \in \mathcal{F}} A_{ifu, t-k} \mathcal{H}_{fu, t-k} \quad (9)$$

Implicit Social Influences and Environmental Events.

The hidden social influences are composed of unobserved social relationships, unacquainted users, and the changing of social context [2]. In other words, a user can be influenced by any users via any features in health social networks. It is hard to exactly define the implicit social influences. Fortunately, the dynamic of neural networks offers us a great solution to capture the flexibility of implicit social influences. In fact, given a user u , each visible variable v_i and hidden variable h_j are connected to all historical variables of all other users. It is similar to self-motivation modeling, the influence effects of each user, and the social context on the user u are captured via the weight matrices A and B . These effects can be integrated into the dynamic biases $\hat{a}_{i,t}$ and $\hat{b}_{j,t}$ as well. The dynamic biases in Eq. 6 and Eq. 9 become:

$$\hat{b}_{j,t} = b_j + \sum_{k \in \{1, \dots, N\}} \sum_{f \in \mathcal{F}} \sum_{u \in U} B_{jfu, t-k} \mathcal{H}_{fu, t-k} \quad (10)$$

$$\hat{a}_{i,t} = a_i + \sum_{k \in \{1, \dots, N\}} \sum_{f \in \mathcal{F}} \sum_{u \in U} A_{ifu, t-k} \mathcal{H}_{fu, t-k} \quad (11)$$

The *environmental events* such as the number of competitions and meet-up events are included in individual features. Therefore, the effect of environmental events is well embedded into the model. It will interact with self-motivation and implicit social influences to capture the behaviors of the users. Next, we will incorporate explicit social influence into our model.

Statistical Explicit Social Influences. Individuals tend to be influenced to perform similar behaviors as their friends (*homophily principle*). Let us define user similarity as follows. Given two neighboring users u and m , a simple way to quantify their similarity is to applying a *cosine function* of their individual features (i.e., \mathbf{v}^u and \mathbf{v}^m). In addition, the hidden layer \mathbf{h} detects higher features of \mathbf{v} and thus it




Figure 3. A sample of user similarity distributions.




Figure 4. The cumulative number of friend connections in the YesiWell study.

can reflect the user similarity as well. The user similarity between u and m at time t , denoted $s_t(u, m)$, is defined as:

$$s_t(u, m) = \cos_t(u, m | \mathbf{v}) \times \cos_t(u, m | \mathbf{h}) \quad (12)$$

where $\cos_t(\cdot)$ is a cosine similarity function, i.e.,

$$\cos_t(u, m | \mathbf{v}) = \frac{p(\mathbf{v}_t^u | \mathbf{h}_t^u, \mathcal{H}_{t_{<}}^u) \cdot p(\mathbf{v}_t^m | \mathbf{h}_t^m, \mathcal{H}_{t_{<}}^m)}{\|p(\mathbf{v}_t^u | \mathbf{h}_t^u, \mathcal{H}_{t_{<}}^u)\| \|p(\mathbf{v}_t^m | \mathbf{h}_t^m, \mathcal{H}_{t_{<}}^m)\|} \quad (13)$$

$$\cos_t(u, m | \mathbf{h}) = \frac{p(\mathbf{h}_t^u | \mathbf{v}_t^u, \mathcal{H}_{t_{<}}^u) \cdot p(\mathbf{h}_t^m | \mathbf{v}_t^m, \mathcal{H}_{t_{<}}^m)}{\|p(\mathbf{h}_t^u | \mathbf{v}_t^u, \mathcal{H}_{t_{<}}^u)\| \|p(\mathbf{h}_t^m | \mathbf{v}_t^m, \mathcal{H}_{t_{<}}^m)\|} \quad (14)$$

Figure 3 illustrates a sample of user similarity spectrum (i.e., $s_t(\cdot, \cdot)$) of all the edges in our social network over time. We randomly select 35 similarities of neighboring users for each day in ten months. Apparently the distributions are not uniform and different time intervals present various distributions. To well qualify the similarity between individuals and their friends, it potentially requires a *cumulative distribution function* (CDF). This quality of similarities is considered as the explicit social influences of local neighbors on individuals. In addition, our health social network is developed from scratch. Users do not have many friends initially. As time goes by, they will have more connections to other users (Figure 4). Thus a temporal smoothing is demanded to better capture the explicit social influences. To formulate this idea, we propose a *statistical explicit social influence function* as follows:

Definition 1: The statistical explicit social influence η_t^u of a user u at time t is defined as an exponential similarity

average of the cumulative distribution function (CDF) of the instant similarity over the user similarity spectrum, i.e.,

$$\eta_t^u = \alpha \eta_{t-\tau}^u + (1-\alpha) \frac{1}{|\mathbb{Z}_t^u|} \sum_{m=1}^{|\mathbb{Z}_t^u|} l_t(u, m) \times p(s_t \leq s_t(u, m))$$

where $\mathbb{Z}_t^u = \sum_{m=1}^{|\mathbb{Z}_t^u|} l_t(u, m)$, and the indicator function l_t is 1 if user u is connected to user m until time t (i.e., $e_{u,m} \in E_t$), and 0 otherwise. s_t is the similarity between two arbitrary neighboring users in the social network at time t . $p(s_t \leq s_t(u, m))$ represents the probability that similarity is less than or equal to the instant similarity $s_t(u, m)$. α and τ are two parameters to capture the dynamic of η .

The effect of explicit social influences is integrated to the dynamic biases of hidden and visible variables. The dynamic bias $\hat{b}_{j,t}$ and $\hat{a}_{i,t}$ in Eq. 10 and Eq. 11 now becomes:

$$\begin{aligned} \hat{b}_{j,t} &= b_j + \sum_{k \in \{1, \dots, N\}} \sum_{f \in \mathcal{F}} \sum_{u \in U} B_{jfu, t-k} \mathcal{H}_{fu, t-k} + \beta_j \eta_t^u \\ \hat{a}_{i,t} &= a_i + \sum_{k \in \{1, \dots, N\}} \sum_{f \in \mathcal{F}} \sum_{u \in U} A_{ifu, t-k} \mathcal{H}_{fu, t-k} + \beta_i \eta_t^u \end{aligned} \quad (15)$$

where parameters β_i and β_j present the ability to observe the explicit social influences η_t^u of user u given v_i and h_j .

Inference and Learning. Inference in the SRBM model is no more difficult than in the RBM. The states of the hidden variables are determined by both the input they receive from the visible variables and the input they receive from the historical variables. The conditional probability of hidden and visible variables at time interval t can be computed as in Eq. 7, Eq. 8, and Eq. 15. The energy function becomes:

$$E(\mathbf{v}_t, \mathbf{h}_t | \mathcal{H}_{t-}, \theta) = \sum_{i \in \mathbf{v}} \frac{(v_{i,t} - \hat{a}_{i,t})^2}{2\sigma_i^2} - \sum_{j \in \mathbf{h}} \hat{b}_{j,t} h_{j,t} - \sum_{i \in \mathbf{v}, j \in \mathbf{h}} \frac{v_{i,t}}{\sigma_i} h_{j,t} W_{ij}.$$

Contrastive divergence [5] is used to train the SRBM. The updates for the symmetric weights, W , as well as the static biases, a and b , have the same form as Eq. 5. However they have a different effect because the states of the hidden and visible variables are now influenced by the implicit and explicit social influences. The updates for the directed weights, A and B , are also based on simple pairwise products. The gradients are summed over all the training time intervals $t \in T_{train} = T_{data} \setminus \{t - M + 1, \dots, t - M + N\}$:

$$\begin{aligned} \Delta W_{ij} &= \sum_t (\langle v_{i,t} h_{j,t} \rangle_d - \langle v_{i,t} h_{j,t} \rangle_r); \Delta a_i = \sum_t (\langle v_{i,t} \rangle_d - \langle v_{i,t} \rangle_r); \Delta b_j = \sum_t (\langle h_{j,t} \rangle_d - \langle h_{j,t} \rangle_r); \\ \Delta A_{ifu, t-k} &= \sum_t (\langle v_{i,t} \mathcal{H}_{fu, t-k} \rangle_d - \langle v_{i,t} \mathcal{H}_{fu, t-k} \rangle_r); \\ \Delta B_{jfu, t-k} &= \sum_t (\langle h_{j,t} \mathcal{H}_{fu, t-k} \rangle_d - \langle h_{j,t} \mathcal{H}_{fu, t-k} \rangle_r); \\ \Delta \beta_i &= \sum_t (\langle v_{i,t} \rangle_d - \langle v_{i,t} \rangle_r) \eta_t^u; \Delta \beta_j = \sum_t (\langle h_{j,t} \rangle_d - \langle h_{j,t} \rangle_r) \eta_t^u. \end{aligned}$$

We train the SRBM for each user independently. At any time we update the parameters, we will update the explicit social influences for all the users.

Human Behavior Prediction. On top of the SRBM model, we put an output layer which is commonly called softmax layer for a user behavior prediction task. Our goal

is to predict whether a user *increases* or *decreases* physical exercises. Thus the softmax layer contains a single output variable \hat{y} and binary target values: 1 for increases, and 0 for decreases. The output variable \hat{y} is fully linked to the hidden variables by weighted connections S which includes $|\mathbf{h}|$ parameters s_j . We use the logistic function as an activation function to saturate the two target values, i.e.,

$$\hat{y} = \sigma(c + \sum_{j \in \mathbf{h}} h_j s_j) \quad (16)$$

where c is a static bias. Given a user $u \in U$, a set of training vectors $X = \{\mathcal{F}_u^t, E_t | t \in T_{train}\}$ and an output vector $Y = \{y_t | t \in T_{train}\}$, the probability of a binary output $y_t \in \{0, 1\}$ given the input x_t as follows:

$$P(Y|X, \theta) = \prod_{t \in T_{train}} \hat{y}_t^{y_t} (1 - \hat{y}_t)^{1-y_t} \quad (17)$$

where $\hat{y}_t = P(y_t = 1 | x_t, \theta)$.

A loss function to appropriately deal with the binomial problem is *cross-entropy error*. It is given by

$$C(\theta) = - \sum_{t \in T_{train}} (y_t \log \hat{y}_t + (1 - y_t) \log(1 - \hat{y}_t)) \quad (18)$$

As this last training, *Back-propagation* is used to fine-tune all the parameters together. The derivatives of the objective $C(\theta)$ with respect to all the parameters over all the training cases $t \in T_{train}$ can be computed as:

$$\begin{aligned} \frac{\partial C(\theta)}{\partial s_j} &= - \sum_t (y_t - \hat{y}_t) h_j; \frac{\partial C(\theta)}{\partial c} = - \sum_t (y_t - \hat{y}_t) \\ \frac{\partial C(\theta)}{\partial W_{ij}} &= - \sum_t (y_t - \hat{y}_t) s_j h_j (1 - h_j) v_i \\ \frac{\partial C(\theta)}{\partial a_i} &= - \sum_t (y_t - \hat{y}_t) s_j h_j (1 - h_j) W_{ij} \\ \frac{\partial C(\theta)}{\partial b_j} &= - \sum_t (y_t - \hat{y}_t) s_j h_j (1 - h_j) \\ \frac{\partial C(\theta)}{\partial A_{ifu, t-k}} &= - \sum_t (y_t - \hat{y}_t) s_j h_j (1 - h_j) W_{ij} \mathcal{H}_{fu, t-k} \\ \frac{\partial C(\theta)}{\partial B_{jfu, t-k}} &= - \sum_t (y_t - \hat{y}_t) s_j h_j (1 - h_j) \mathcal{H}_{fu, t-k} \\ \frac{\partial C(\theta)}{\partial \beta_i} &= - \sum_t (y_t - \hat{y}_t) s_j h_j (1 - h_j) W_{ij} \eta_t^u \end{aligned}$$

Causal Generation. In the prediction task, we need to predict the y_{t+1}^u without observing the \mathcal{F}_u^{t+1} . In other words, the visible and hidden variables are not observed at the future time point $t + 1$. Thus we need a causal generation step to initiate these variables. Causal generation from a learned SRBM can be done just like the learning procedure. We always keep the historical variables fixed and perform alternating Gibbs sampling to obtain a joint sample of the visible and hidden variables from the SRBM. To start alternating Gibbs sampling, a good choice is to set $\mathbf{v}_t = \mathbf{v}_{t-1}$ since \mathbf{v}_{t-1} can be considered as a strong prior of \mathbf{v}_t . This picks new hidden and visible variables that are compatible with each other and with the recent historical variables. Afterward, we aggregate the hidden variables to evaluate the output \hat{y} .

IV. EXPERIMENTS

We have carried out a series of experiments using both real-world and synthetic health social networks to validate our proposed SRBM model. We first elaborate the

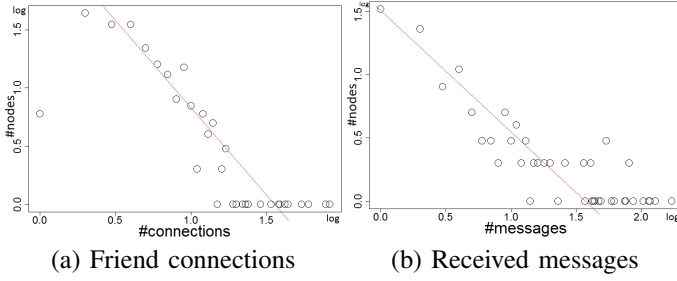


Figure 5. Some distributions in our dataset.

experiment configurations, evaluation metrics, and baseline approaches. Then, we introduce the experimental results.

The YesiWell data and Experiment Configurations.

The YesiWell social network data were collected from Oct 2010 to Aug 2011 as a collaboration between PeaceHealth Laboratories, SK Telecoms Americas, and University of Oregon to record daily physical activities, social activities (i.e., text messages, competitions, etc.), biomarkers, and biometric measures (i.e., cholesterol, BMI, etc.) for a group of 254 individuals. Physical activity information including the number of walking and running steps were reported via a mobile device carried by each user. All users enrolled an online social network allowing them to make friends and communicate each other. Users' biomarkers and biometric measures are recorded via daily/weekly/monthly medical tests performed at home (i.e., individually) or at PeaceHealth's laboratories.

In total, we have 30 features taken into account (Table I). All the features are weekly summarized. Figure 5 illustrates the distributions of friend connections, and the number of received messages in the health social network. They clearly follow the Power law distribution. Note that, given a week if a user does exercise more than the previous week, he/she is considered increasing exercise in that week. Otherwise, the user will be considered decreasing exercise. The number of hidden units and the number of previous time intervals N respectively are set to 200 and 3. The weights are randomly initialized from a zero-mean Gaussian with a standard deviation of 0.01. All the learning rates are set to 10^{-3} . A contrastive divergence CD_{20} [5] is used to maximize likelihood learning. We train the model for each user independently.

Evaluation Metrics. In the experiment, we leverage the previous 10 week records to predict the behaviors of all the users (i.e., increase or decrease exercises) in the next week. The prediction quality metric, i.e., *accuracy*, is as follows:

$$accuracy = \frac{\sum_{i=1..|U|} I(y_i = \hat{y}_i)}{|U|} \quad (19)$$

where y_i is the true user activity of the user u_i , and \hat{y}_i denotes the predicted value, I is the indication function.

Competitive Prediction Models. We compare the SRBM model with the conventional methods reported in [18].

Table I
PERSONAL CHARACTERISTICS.

Behaviors	#joining competitions	#exercising days
	#goals set	#goals achieved
	\sum (distances)	avg(speeds)
Social Communications (the number of inbox messages)	Encouragement	Fitness
	Followup	Games
	Competition	Personal
	Study protocol	Technique
	Progress report	Meetups
	Social network	Goal
	Wellness meter	Feedback
	Heckling	Explanation
	Invitation	Notice
	Technical fitness	Physical
Biomarkers	Wellness Score	BMI
	Wellness Score slope	BMI slope

The competitive methods are divided into two categories: personalized behavior prediction methods and socialized behavior prediction methods. Personalized methods only leverage individual's past behavior records for future behavior predictions. Socialized methods use both one person's past behavior records and his or her friends' past behaviors for predictions. Specifically, five models reported in [18] are Socialized Gaussian Process (**SGP**) model, Socialized Logistical Autoregression (**SLAR**) model, Personalized Gaussian Process (**PGP**) model, Logistical Autoregression (**LAR**) model, and Behavior Pattern Search (**BPS**) model.

In addition, we also consider the RBM related extensions, i.e., the **CRBM** and **ctRBM**, as competitive models. The CRBM can be directly apply to our problem by ignoring the implicit and explicit influences in our SRBM. Since the ctRBM cannot directly incorporate individual features with social influences to model human behaviors, we only can apply its social influence function into our model. In fact, we replace our statistical explicit social influence function by the ctRBM's social influence function. We call this version of ctRBM as a *Socialized ctRBM* (**SctRBM**).

Validation of the SRBM Model. Our task of validation concerns on three key issues: 1) which configurations of the parameters α and τ produces the best-fit social influence distribution, 2) which of potential social influence functions and our statistical explicit social function produce a better-fit social influence distribution, and 3) whether the SRBM model is better than the competitive models in terms of prediction accuracy.

We carry out the validation through three approaches. One is by conducting the human behavior prediction with various settings of α and τ . By this we look for an optimal configuration for the statistical explicit social influence function. The second approach is to compare the optimal setting of our statistical explicit social influence function with its different forms and existing algorithms. The third approach is to compare our SRBM model with the competitive models in terms of prediction accuracy.

• Figure 6a illustrates the surface of the behavior prediction accuracy of the SRBM model with variations of the




Figure 6. Validation of the SRBM model on the YesiWell health social network.

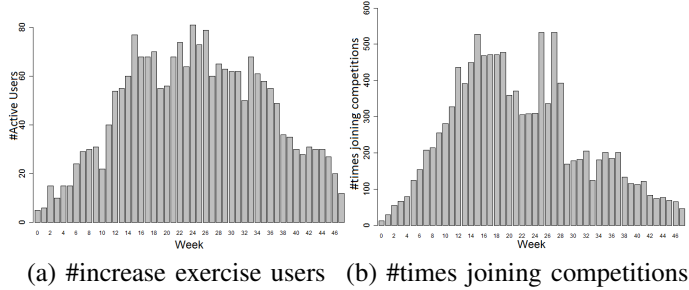


Figure 7. The distributions of users' activities.

two parameters α and τ on our health social network. We observed that the smaller values of τ tend to have higher prediction accuracies. It is quite reasonable since the more recent behaviors have stronger influences. The temporal smoothing parameter τ has similar effects to a time decay function [25]. Meanwhile, the middle range values of α offer better prediction accuracies. Obviously, the optimal setting values of α and τ are 0.5 and 1 respectively.

- To test the correctness, we compare our optimal setting (i.e., $\tau=1$, $\alpha=0.5$) of the explicit social influence function with its different forms such as: 1) without temporal smoothing component (i.e., $\eta_t^u = \frac{1}{|\mathbb{P}_t^u|} \sum_{m \in \mathbb{P}_t^u} p(s_t \leq s_t(u, m))$), and 2) replacing our function by the social influence function in the ctRBM [12], this becomes the aforementioned SctRBM. Figure 6b shows that either the SRBM without temporal smoothing or the SctRBM have significant lower prediction accuracies compared with the SRBM with optimal setting. So, our function is effective, and the optimal setting improves the prediction accuracy by 12% in our health social network. It also offers a better performance for the SRBM compared with other social influence functions. This is because our health social network is developed from scratch. Users do not have many connections initially. They will have more connections over time (Figure 4). In other words, our function produces a better-fit social influence distribution in our health social network. From now, we use the optimal setting of the SRBM in other experiments.

- To examine the prediction accuracy, we compare the proposed SRBM model with the competitive models in terms of human behavior prediction. Figure 6c shows the accuracy comparison in 37 weeks in our health social network. It




Figure 8. Accuracies on the synthetic data.

is clear that the SRBM outperforms the other models. The accuracies of the competitive models tend to be dropped in the middle period of the study. In essence, all the behavior determinants and their interactions potentially become stronger since all the users improve their activities such as walking and running, participating more competitions, and etc. (Figure 7) in the middle weeks. Either insufficiently or lacking of modeling one of the determinants and their interactions results in a low and unstable prediction performance. This is the case of the competitive models. They do not well capture the social influences and environmental events.

Meanwhile, the SRBM comprehensively models all the determinants. The correlation between the personal characteristics and the implicit social influences can be detected by the hidden variables. Thus, much information has been leveraged to predict individual behaviors. In addition, our prediction accuracy stably increases over time. That means our model well captures the growing of our health social network (Figure 4). In obvious, our model achieves *higher prediction accuracy* and a more *stable* performance. The SRBM achieves the best accuracy in average as 0.887.

Synthetic Health Social Network. To illustrate that our model can be generally applied on different datasets, we perform further experiments on a synthetic health social network. To generate the synthetic data, we use software Pajek¹ to generate graphs under the Scale-Free/Power Law Model². However the vertices in the current synthetic graph do not have individual features similar to the real-word data.

¹<http://vlado.fmf.uni-lj.si/pub/networks/pajek/>
²**Scale-Free/Power Law Model (SF)** is a network model whose node degrees follow the power law distribution or at least asymptotically.

An appropriate solution to this problem is to apply a graph matching algorithm to map pairwise vertices between the synthetic and real social networks. In order to do so, we first generate a graph with 254 nodes and the average node degree is 5.4 (i.e., similar to the real YesiWell data). Then, we apply PATH [24] which is a very well-known and efficient graph matching algorithm to find a correspondence between vertices of the synthetic network and vertices of the YesiWell network. The source code of the PATH algorithm is available in the graph matching package GraphM (<http://cbio.ensmp.fr/graphm/>). Then, we can assign all the individual features and behaviors of any real user to corresponding vertices in the synthetic network. Consequently, we have a synthetic health social network which imitates our real-world dataset. Figure 8 shows the accuracies of the conventional models and the SRBM model on the synthetic data. We can see that the our model still outperforms the conventional models in terms of the prediction accuracy.

V. CONCLUSIONS AND FUTURE WORKS

This paper introduces the SRBM, a novel deep learning model for human behavior prediction in health social networks. By incorporating all the human behavior determinants which are self-motivation, implicit and explicit social influences, and environmental events, our SRBM model predicts the future activity levels of users more accurately and more stably than conventional methods.

Our work can be extended in several directions. First, we can leverage the domain knowledge to generate explanations for predicted behaviors. Second, although the approach explored in this paper is rooted on the RBM [19], other alternatives are possible, which can be based on CNNs [10] or Sum-Product Networks [15]. We plan to explore and compare these different strategies in the future work.

ACKNOWLEDGMENT

This work is supported by the NIH grant R01GM103309 to the SMASH project. We thank Xiao Xiao, Rebeca Sacks, and Ellen Klowden for their contributions.

REFERENCES

- [1] A. Bandura. Human agency in social cognitive theory. *The American Psychologist*, 1989.
- [2] N.A. Christakis. The hidden influence of social networks. In *TED2010*.
- [3] N.A. Christakis and J.H. Fowler. The spread of obesity in a large social network over 32 years. *The New England Journal of Medicine*, 357:370–9, 2007.
- [4] R.G. Fichman and C.F. Kemerer. The illusory diffusion of innovation: An examination of assimilation gaps. *Information Systems Research*, 10(3):255–275, 1999.
- [5] G.E. Hinton. Training products of experts by minimizing contrastive divergence. *Neural Computation*, 14(8):1771–1800, 2002.
- [6] J. Huang, X.-Q. Cheng, H.-W. Shen, T. Zhou, and X. Jin. Exploring social influence via posterior effect of word-of-mouth recommendations. In *WSDM'12*, pages 573–582.
- [7] J. Kawale, A. Pal, and J. Srivastava. Churn prediction in mmorpgs: A social influence based approach. In *CSE'09*, pages 423–428.
- [8] D. Kempe, J.M. Kleinberg, and E. Tardos. Maximizing the spread of influence through a social network. In *KDD'03*, pages 137–146.
- [9] D. Kil, F. Shin, B. Piniewski, J. Hahn, and K. Chan. Impacts of social health data on predicting weight loss and engagement. In *O'Reilly StrataRx Conference, San Francisco, CA*, October 2012.
- [10] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [11] K. Lerman, S. Intagorn, J.-K. Kang, and R. Ghosh. Using proximity to predict activity in social networks. In *WWW Companion'12*, pages 555–556.
- [12] X. Li, N. Du, H. Li, K. Li, J. Gao, and A. Zhang. A deep learning approach to link prediction in dynamic networks. In *SIAM'14*, pages 289–297.
- [13] A. Marshall, E.G. Eakin, E.R. Leslie, and N. Owen. Exploring the feasibility and acceptability of using internet technology to promote physical activity within a defined community. *Health Promotion Journal of Australia*, 2005(16):82–4, 2005.
- [14] R.R. Pate, M. Pratt, S.N. Blair, and et al. Physical activity and public health. a recommendation from the centers for disease control and prevention and the american college of sports medicine. *Journal of the American Medical Association*, 273(5):402–7, 1995.
- [15] H. Poon and P. Domingos. Sum-product networks: A new deep architecture. In *IEEE ICCV Workshops*, pages 689–690, 2011.
- [16] R. M. Ryan and E. L. Deci. Intrinsic and extrinsic motivations: Classic definitions and new directions. *Contemporary Educational Psychology*, 25(1):54 – 67, 2000.
- [17] K. Saito, R. Nakano, and M. Kimura. Prediction of information diffusion probabilities for independent cascade model. In *KES 2008*, pages 67–75.
- [18] Y. Shen, R. Jin, D. Dou, N. Chowdhury, J. Sun, B. Piniewski, and D. Kil. Socialized gaussian process model for human behavior prediction in a health social network. In *ICDM'12*, pages 1110–1115.
- [19] P. Smolensky. Information processing in dynamical systems: Foundations of harmony theory. chapter Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol. 1, pages 194–281. 1986.
- [20] G.W. Taylor, G.E. Hinton, and S.T. Roweis. Modeling human motion using binary latent variables. In *NIPS'06*, pages 1345–1352.
- [21] B. Viswanath, A. Mislove, M. Cha, and K.P. Gummadi. On the evolution of user interaction in facebook. In *WOSN'09*, pages 37–42.
- [22] M. Welling, M. Rosen-Zvi, and G.E. Hinton. Exponential family harmoniums with an application to information retrieval. In *NIPS'04*, pages 1481–1488.
- [23] J. Yang, X. Wei, M.S. Ackerman, and L.A. Adamic. Activity lifespan: An analysis of user survival patterns in online knowledge sharing communities. In *ICWSM'10*, pages 186–193.
- [24] M. Zaslavskiy, F. Bach, and J.-P. Vert. A path following algorithm for graph matching. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 5099, pages 329–337, 2008.
- [25] Y. Zhu, E. Zhong, S.J. Pan, X. Wang, M. Zhou, and Q. Yang. Predicting user activity level in social networks. In *CIKM'13*, pages 159–168.