

Wei Pan, Wen Dong, Manuel Cebrian,
Taemie Kim, James H. Fowler, and Alex (Sandy) Pentland

Modeling Dynamical Influence in Human Interaction

[Using data to make better inferences
about influence within social systems]

How can we model influence between individuals in a social system, even when the network of interactions is unknown? In this article, we review the literature on the “influence model,” which utilizes independent time series to estimate how much the state of one actor affects the state of another actor in the system. We extend this model to incorporate dynamical parameters that allow us to infer how influence changes over time, and we provide three examples of how this model can be applied to simulated and real data. The results show that the model can recover known estimates of influence, it generates results that are consistent with other measures of social networks, and it allows us to uncover important shifts in the way states may be transmitted between actors at different points in time.

INTRODUCTION

The concept of influence is extraordinarily important in the natural sciences. The basic idea of influence is that an outcome in one entity can cause an outcome in another entity. Flip over the first domino, and the second domino will fall. If we understand exactly how two dominoes interact—how one domino influences another—and we know the initial state of the dominoes and how they are situated relative to one another, then we can predict the outcome of the whole system.

For decades, social scientists have also been interested in analyzing and understanding who influences whom in social systems [1], [2]. But the analogue with the physical world is not exact. In the social world, influence can be more complicated because internal states are often unobservable, intentional

Signal and Information Processing for Social Learning and Networking



Digital Object Identifier 10.1109/MSP.2011.942737
Date of publication: 17 February 2012

©STOCKPHOTO.COM/ANDREY PROKHOROV

behavior can be strategic, and the situational context of specific interactions can change the effect one actor has on another. And even more challenging, actors can choose with whom they interact, which can confound efforts to infer influence from correlated behaviors between actors [3]. As a consequence, there has been tremendous interest in developing methods for better understanding the effect that networked interactions have on the spread of social behaviors and outcomes.

Social scientists have already carefully studied communication settings like group discussions to better understand the causal mechanisms that underlie influence [4], but recent advances in modern sensing systems such as sociometric badges [5] and cell phones [6] now provide valuable social behavioral signals from each individual at very high resolution in time and space. The challenge for those of us interested in signal processing is how to use this data to make better inferences about influence within social systems.

In this article, we describe the “influence model” first articulated in [7] and the subsequent literature that has refined this approach. Similar definitions on influence in other literature include research on voting models in physics [8], cascade models in epidemiology [1], attitude influence in psychology [9], and information exchange models in economics [10]. The influence model is built on an explicit abstract definition of influence: an entity’s state is influenced by its network neighbors’ states and changes accordingly. Each entity in the network has a specifically defined strength of influence over every other entity in the network, and, equivalently, each relationship can be weighted according to this strength.

We believe that the influence model is a unique tool for social scientists because it can be applied to a wide range of social systems (including those where aggregates like organizations, states, and institutions can themselves be thought of as “actors” in a network). The influence model also enables researchers to infer interactions and dynamics when the network structure is unknown—all that is needed is information about time-series signals from individual observations. And although this method is subject to the same limitations as any observational network study [11], the ordering of behaviors in time and social space makes it less likely that alternative mechanisms like selection effects and contextual heterogeneity can explain the patterns of influence ascertained by the model.

OVERVIEW FOR THE INFLUENCE MODEL

ENTITIES IN A SOCIAL SYSTEM

We describe the influence model here, followed by a review on its history in the section “Previous Applications of the Influence Model.” The model starts with a system of C entities. We assume that each entity c is associated with a finite set of possible states $1, \dots, S$. At different time t , each entity c is in one of the states, denoted by $h_t^{(c)} \in \{1, \dots, S\}$. It is not necessary that each entity is associated with the same set of possible states. Some entities can have more or less states. However, to

simplify our description, we assume that each entity’s latent state space is the same without loss of generality.

The state of each entity is not directly observable. However, as in the hidden Markov model (HMM), each entity emits a signal $O_t^{(c)}$ at time stamp t based on the current latent state $h_t^{(c)}$, following a conditional emission probability $\text{Prob}(O_t^{(c)}|h_t^{(c)})$. The emission probability can either be multinomial or Gaussian for discrete and continuous cases respectively, exactly as in HMM literature [12].

It is important to note here that entities can be anything that has at least one state. For example, they could be people in group discussions who are in a “talking” state or a “silent” state. Or they could be geographical districts with variation in flu incidence that yields some in a “high-incidence” state or a “low-incidence” state. The fundamental question remains in any situation, does the state in one entity influence (cause a change) the state in another entity? It is therefore possible to apply the influence model to a wide range of contexts.

INFLUENCE BETWEEN ENTITIES

The influence model is composed of entities interacting and influencing each other. “Influence” is defined as the conditional dependence between each entity’s current state $h_t^{(c)}$ at time t and the previous states of all entities $h_{t-1}^{(1)}, \dots, h_{t-1}^{(C)}$ at time $t-1$. Therefore, intuitively, $h_t^{(c)}$ is influenced by all other entities.

An important implication of this Markovian assumption is that all effects from states at times earlier than $t-1$ are completely accounted for by incorporating all information from time $t-1$. This does not mean that earlier times had no effect or were unimportant; it just means that their total effect is felt in the immediately previous time period. And even path dependent processes (of which there are many in the social sciences) can operate this way, one time period at a time.


We now discuss the conditional probability

$$\text{Prob}(h_t^{(c)}|h_{t-1}^{(1)}, \dots, h_{t-1}^{(C)}). \quad (1)$$

Once we have $\text{Prob}(h_t^{(c)}|h_{t-1}^{(1)}, \dots, h_{t-1}^{(C)})$, we naturally achieve a generative stochastic process. As in the coupled Markov model [13], we can take a general combinatorial approach (1), and convert this model to an equivalent HMM, in which each different latent state combination of $(h_{t-1}^{(1)}, \dots, h_{t-1}^{(C)})$ is represented by a unique state. Therefore, for a system with C interacting entities, the equivalent HMM will have a latent state space of size S^C , exponential to the number of entities in the system, which generates insurmountable computational challenges in real applications.

The influence model, on the other hand, uses a much simpler mixture approach with far fewer parameters. Entities $1, \dots, C$ influence the state of c' in the following way:

$$\text{Prob}(h_t^{(c')}|h_{t-1}^{(1)}, \dots, h_{t-1}^{(C)}) = \sum_{c \in \{1, \dots, C\}} \underbrace{\mathbf{R}_{c',c}}_{\text{tie strength}} \times \underbrace{\text{Infl}(h_t^{(c')}|h_{t-1}^{(c)})}_{\text{influence } c \rightarrow c'} \quad (2)$$



[FIG1] Different versions of the sociometric badge are shown in (a) and (b). The sociometric badge is a wearable sensing device for collecting individual behavioral data. In (c), a group brainstorming session is shown where all participants are wearing the sociometric badges.

where \mathbf{R} is a $C \times C$ matrix. (\mathbf{R}_{c_1, c_2}) represents the element at the c_1 th row and the c_2 th column of the matrix \mathbf{R} . \mathbf{R} is row stochastic, i.e., each row of this matrix sums up to one. $\text{Infl}(h_t^{(c)}|h_{t-1}^{(c)})$ is modeled using a $S \times S$ row stochastic matrix $\mathbf{M}^{c, c'}$, so that $\text{Infl}(h_t^{(c)}|h_{t-1}^{(c)}) = \mathbf{M}_{h_{t-1}^{(c)}, h_t^{(c)}}^{c, c'}$, where $\mathbf{M}_{h_{t-1}^{(c)}, h_t^{(c)}}^{c, c'}$ represents the element at the $h_{t-1}^{(c)}$ th row and $h_t^{(c)}$ th column of matrix $\mathbf{M}^{c, c'}$. The row stochastic matrix $\mathbf{M}^{c, c'}$ captures the influence from c over c' , and is very similar to the HMM literature [12].

Equation (2) can be viewed as follows: All entities' states at time $t - 1$ will influence the state of entity c' at time t . However, the strength of influence is different for different entities; the strength of c over c' is captured by $\mathbf{R}_{c, c'}$. As a result, the state distribution for entity c' at time t is a combination of influence from all other entities weighted by their strength over c' . Such definition of influence from neighbor nodes is not unique, and it has been well studied in statistical physics and psychology as well [8], [9]. Because \mathbf{R} captures influence strength between any two entities, we refer to \mathbf{R} as the influence matrix.

Generally, for each entity c , there are C different transition matrices in the influence model to account for the influence dynamics between c and c' , $c' = 1, \dots, C$. However, it can be simplified by replacing the C different matrices with only two $S \times S$ matrices \mathbf{E}^c and \mathbf{F}^c : $\mathbf{E}^c = \mathbf{M}^{c, c}$, which captures the self-state transition; Empirically, in many systems an entity c may influence other entities in the same manner. For instance, a strong politician always asks everyone to support his political view no matter who they are. Therefore, sometimes we can simplify the system by assuming $\mathbf{M}^{c, c'} = \mathbf{F}^c, \forall c' \neq c$.

INFERENCE

The influence model is a generative model defined by parameters $\mathbf{R}, \mathbf{E}^{1:C}, \mathbf{F}^{1:C}$ and the emission probabilities $\text{Prob}(O_t^{(c)}|h_t^{(c)})$, $\forall c$. As in most generative machine learning models, these parameters are not set by users, but they are automatically learned from observations $O_{1:T}^1, \dots, O_{1:T}^C$. The inference algorithms for learning these parameters will be discussed in the section "The Dynamical Influence Model."

The influence model has two key advantages over other machine learning approaches. First, the number of parameters

grows quadratically with the latent space size S and linearly to the number of entities C . As a result, the influence model is resistant to overfitting when training data is limited compared with other approaches [14].

Second, the model captures the tie strength between entities using a $C \times C$ matrix \mathbf{R} . \mathbf{R} inferred by our model can be naturally treated as the adjacency matrix for a directed weighted graph between nodes. This key contribution connects the conditional probabilistic dependence to a weighted network topology. In fact, the most common usage for the influence model in the literature is to use \mathbf{R} to infer social structure [15], [16].

PREVIOUS APPLICATIONS OF THE INFLUENCE MODEL

The influence model has been applied to various social systems, particularly those that have been monitored by sociometric badges like those shown in Figure 1. These badges are personal devices that collect individual behavioral data including audio, location, and movement. Early attempts to analyze data from these badges focused on questions revolving around group interaction and interpersonal influence.

The first application of the influence model [7] attempted to infer influence networks from audio recordings of a group discussion session with five individuals. The researchers used audio features as observations $O_t^{(c)}$ and modeled the latent state space to be either "speaking" or "nonspeaking." They then used the model to infer the underlying pattern of interpersonal influence from the noisy signals measured directly from each individual and their interactions on turn taking.

An important question about these inferences relates to their validity: How do we know that the measure of influence is real? To answer this question, researchers applied the influence model to conversation data from sociometric badges on 23 individuals and showed that the influence strength between individuals learned by the model correlated extremely well with individual centrality in their social networks (with $R = 0.92, p < 0.0004$) [16]. This evidence suggests that the influence matrix defined as the weights in the conditional dependence on states of other entities is an important measure for the social position of the individuals in real interaction data.

In other words, even more abstract concepts related to influence like status or social hierarchy might be captured by the inferences of the influence model.

EVEN MORE ABSTRACT CONCEPTS RELATED TO INFLUENCE LIKE STATUS OR SOCIAL HIERARCHY MIGHT BE CAPTURED BY THE INFERENCES OF THE INFLUENCE MODEL.

The model has also been applied to many other human interaction contexts [15]. For instance, researchers have used the influence model to understand the functional role (such as follower, orienteer, giver, and seeker) of each individual in the mission survival group discussion data set [17]. They found that the inferred influence matrix helped them to achieve better classification accuracy compared with other approaches. The model has also been applied to the reality mining [6] cell phone sensor data. Using information from 80 Massachusetts Institute of Technology (MIT) affiliates as observations and constraining the latent space of each individual to be binary “work” and “home,” researchers found that the influence matrix learned from this data matches well with the organizational relationship between individuals [14].

Recently the influence model has been extended to a variety of systems, including traffic patterns [18] and flu outbreaks [19]. But more importantly, there have been methodological advances that allow the model to incorporate dynamic changes in the influence matrix itself [19]. This new approach, the dynamical influence model, is a generalization of the inference model, and is discussed in the following section.

Related approaches have utilized Bayesian networks to understand and process social interaction time series data. Examples include coupled HMM [13], dynamic system trees [7], and interacting Markov chains [20]. The key difference between these approaches and the influence model is that the influence matrix \mathbf{R} connects the real network to state dependence.

The key idea of the influence model is to define influence as the state dependence for an entity on the weighted sum of states from network neighbors. This idea has been extensively explored by statistical physicists [8], and very recently by psychologists in modeling attitude influence [9].

THE DYNAMICAL INFLUENCE MODEL

Above, we introduced the influence model, where the influence strength matrix \mathbf{R} remains the same for all t . However, there is extensive evidence leading us to think that influence is indeed a dynamical process [21]. This can also be seen from many real-world experiences: Friendship is not static; in negotiations, your most active opponent may change due to shifts in topics or strategies over time. Therefore, we believe that the influence between subjects may fluctuate as well in many social systems.

Here, we demonstrate how the influence model can be extended to the dynamical case, and we call this generalization the dynamical influence model. Instead of having one single influence strength matrix, \mathbf{R} , we consider a finite set of different influence strength matrices, $\{\mathbf{R}(1), \dots, \mathbf{R}(J)\}$, each representing a different pattern between entities. J is a

hyperparameter set by users to define the number of different interaction patterns. Our approach is basically a switching model, and we also introduce the switching latent state $r_t \in \{1, \dots, J\}, t = 1, \dots, T$,

which indicates the current active influence matrix at time t . Therefore, (2) turns into the following:

$$\text{Prob}(h_t^{(c)} | h_{t-1}^{(1)}, \dots, h_{t-1}^{(c)}) = \sum_{c \in \{1, \dots, C\}} \mathbf{R}(r_t)_{c,c} \times \text{Infl}(h_t^{(c)} | h_{t-1}^{(c)}). \quad (3)$$

As r_t switches to different values between 1 to J at different times t , the dynamics are then determined by different influence matrices $\mathbf{R}(r_t)$.

As shown in the section “Toy Example: Two Interacting Agents,” we note that it is very important to constrain the switching of r_t for two reasons: 1) In many social systems, the change of influence patterns changes slowly and gradually and 2) a prior eliminates the probability of overfitting. Therefore we introduce the following prior for r_t :

$$r_{t+1} | r_t \sim \text{multinomial}(V_{r_t,1}, \dots, V_{r_t,J}), \quad (4)$$

where \mathbf{V} is a system parameter matrix constrained by another hyperparameter $p^V, p^V > 0$. The prior is shown in (5)

$$(V_{r_t,1}, \dots, V_{r_t,J}) \sim \text{Dirichlet}(10^0, 10^0, \dots, 10^{p^V}, \dots, 10^0). \\ \uparrow \quad \uparrow \quad \dots \quad \uparrow \quad \dots \quad \uparrow \quad (5)$$

This prior provides a better control of the process r_1, \dots, r_T . When $p^V = 0$, the Dirichlet prior turns to a uniform distribution. However, the higher p^V gets, the more likely r_{t-1} and r_t will be the same.

Given the model description and hyperparameters J and p^V , we can then write the likelihood function

$$\mathcal{L}(O_{1:T}^{1:C}, h_{1:T}^{1:C}, r_{1:T} | \mathbf{E}^{1:C}, \mathbf{F}^{1:C}, \mathbf{R}(1:J), \mathbf{V}) \\ = \prod_{t=2}^T \left\{ \text{Prob}(r_t | r_{t-1}) \times \prod_{c=1}^C [\text{Prob}(O_t^{(c)} | h_t^{(c)}) \times \text{Prob}(h_t^{(c)} | h_{t-1}^{(1..C)}, r_t)] \right\} \\ \times \prod_{c=1}^C \text{Prob}(O_1^{(c)} | h_1^{(c)}) \text{Prob}(h_1^{(c)}) \text{Prob}(r_1). \quad (7)$$

To demonstrate the difference between the static influence model and the dynamical influence model, we illustrate the Bayesian graph for both models in Figure 2.

Researchers have been studying a variety of alternative time-varying network models: from exponential random graph model (ERGM) [22] to TESLA [23]. ERGM computes a set of features from networks and how they change, and models the distribution of network evolution as the distribution of feature evolution. TESLA uses changing network edges to capture

correlations in node observations with l_1 constraints on edges to ensure sparsity. Another recent version also learns network topology from cascades [24]. Compared with these models, the dynamical influence model serves as a unique generative approach for modeling noisy signals from a dynamical network.

INFERENCE

In signal processing applications, we are given the observation time series signals $O_{1:T}^{(1)}, \dots, O_{1:T}^{(C)}$, and based on these observations we need to learn the distributions of underlying latent variables and system parameters for the dynamical influence model. The inference process for our model is discussed here. Since the dynamical influence model is a generalization of the influence model, the following description is applicable to both models.

Previously, researchers started with a standard exact inference algorithm (junction tree) with exponential complexity, and then moved to an approach based on optimization [16]. Other scholars gradually moved to an approximation approach based on the forward-backward algorithm and variational-Expectation-Maximization (EM) [25], [19]. The influence model can also be trained via other approximations like the mean field method [26].

Here we show some key steps for the variational E-M approach, which has been developed and applied successfully in many data sets. We refer readers to Pan et al. [19] for detail. We denote definition by \equiv , and same distribution by \sim , but the right-hand side of all equations should be normalized accordingly.

E-STEP

We adopt a procedure similar to the forward-backward procedure in the HMM literature. First, we define the following forward parameters for $t = 1, \dots, T$:

$$\alpha_{i,c}^{r_t} \equiv \text{Prob}(h_i^{(c)} | r_t, O_{1:t}), \quad \kappa_t \equiv \text{Prob}(r_t | O_{1:t}), \quad (8)$$

where $O_{1:t}$ denotes $\{O_{1:t}^{(c)}\}_{c=1, \dots, C}$. However, complexity for computing $\alpha_{i,c}^{r_t}$ given $\alpha_{i,c}^{r_{t-1}}$ grows exponentially with respect to C , so we adopt the variational approach [27], and E-M is still guaranteed to converge under variational approximation [27]. We proceed to decouple the chains by

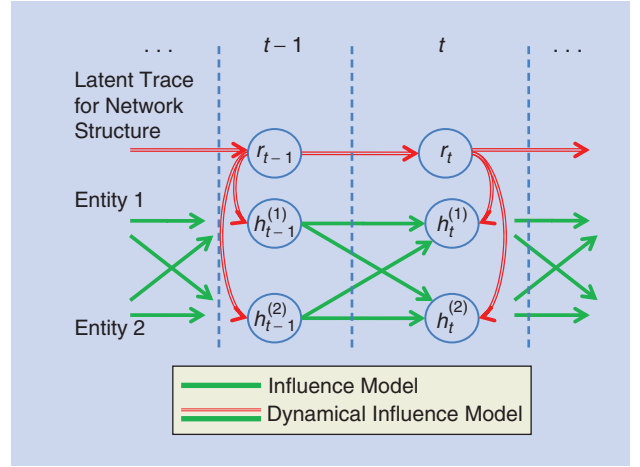
$$\text{Prob}(h_1^{(1)}, \dots, h_t^{(C)} | O_{1:t}, r_t) \approx \prod_c Q(h_i^{(c)} | O_{1:t}, r_t), \quad (9)$$

and naturally

$$\alpha_{i,c}^{r_t} \approx Q(h_i^{(c)} | O_{1:t}, r_t). \quad (10)$$

The approximation adopted here enables us to run inference in polynomial time. Based on this approximation, starting with $\alpha_{i,c}^{r_1}$ and κ_1 , we can compute $\alpha_{i,c}^{r_t}$ and κ_t , $\forall t = 2, \dots, T$ step by step in a forward manner.

Using the same idea, we can compute the following backward parameters for all t in the backward order (i.e. start with $t = T$, then compute $\beta_{i,c}^{r_t}$ and ν_t for $t = T - 1, T - 2, \dots, 1$):



[FIG2] A graphical representation of our model when $C = 2$. The blue lines show the dependence of the influence model described in the section "Overview for the Influence Model." The red lines indicate the layer that brings additional switching capacity to the influence model, and together with the blue lines they fully describe the variable dependence of the dynamical influence model.

$$\beta_{i,c}^{r_t} \equiv \text{Prob}(h_i^{(c)} | r_t, O_{t:T}), \quad \nu_t \equiv \text{Prob}(r_t | O_{t:T}). \quad (11)$$

M-STEP

With κ_t and ν_t , we can estimate

$$\begin{aligned} \xi_{i,j}^t &\equiv \text{Prob}(r_t = i, r_{t+1} = j | O_{1:T}) \\ &= \text{Prob}(r_t = i | O_{1:t}) \text{Prob}(r_{t+1} = j | O_{t+1:T}) \text{Prob}(r_{t+1} | r_t) / \\ &\quad \sum_{i,j} \text{Prob}(r_t = i | O_{1:t}) \text{Prob}(r_{t+1} = j | O_{t+1:T}) \text{Prob}(r_{t+1} | r_t), \end{aligned} \quad (12)$$

$$\lambda_i^t = \text{Prob}(r_t = i | O_{1:T}) = \frac{\sum_j \xi_{i,j}^t}{\sum_i \sum_j \xi_{i,j}^t}, \quad (13)$$

and update V by

$$\mathbf{V}_{i,j} \leftarrow \frac{\sum_t \xi_{i,j}^t + k}{\sum_t \sum_j \xi_{i,j}^t + p^V}, \quad (14)$$

where $k = p^V$ if $i = j$, 0 otherwise.

We then compute the joint distribution $\text{Prob}(h_{t+1}^{(1)}, h_{t+1}^{(2)}, r_{t+1} | O_{1:T})$, and update parameters such as influence matrices $\mathbf{R}(1), \dots, \mathbf{R}(J)$, \mathbf{E}^c and \mathbf{F}^c by marginalizing this joint distribution.

APPLICATIONS

TOY EXAMPLE: TWO INTERACTING AGENTS

In this example, we demonstrate how the dynamical influence model can be applied to find structural changes in network

[TABLE 1] THE LEFT SIDE OF THIS TABLE SHOWS PART OF THE TWO INPUT TOY SEQUENCES FOR A TWO-CHAIN DYNAMICAL INFLUENCE PROCESS. RIGHT SIDE OF THE TABLE SHOWS THE ORIGINAL TWO INFLUENCE MATRICES OF THE TOY MODEL AND THE SAME MATRICES LEARNED BY OUR ALGORITHM WITH $J = 3$ AND $p^V = 10^1$.

SEQ. NO.	DATA (PARTIALLY)	R(1)	R(2)
1	221111121212212... TRUE	$\begin{pmatrix} 0.90 & 0.10 \\ 0.10 & 0.90 \end{pmatrix}$	$\begin{pmatrix} 0.05 & 0.95 \\ 0.95 & 0.05 \end{pmatrix}$
2	112111212121122... LEARNED	$\begin{pmatrix} 0.93 & 0.07 \\ 0.10 & 0.89 \end{pmatrix}$	$\begin{pmatrix} 0.08 & 0.92 \\ 0.94 & 0.06 \end{pmatrix}$

dynamics. As a tutorial, we also explain how readers should adjust two hyperparameters J and p^V in using this model.

From a dynamical influence process composed of two interacting entities, we sample two binary time series of 600 steps. Each chain has two hidden states with a random transition biased to remain in the current state. We sample binary observations from a randomly generated multinomial distribution. To simulate a switch in influence dynamics, we sample with influence matrix $\mathbf{R}(1)$ (shown in Table 1) in the first 200 frames, and later on we sample with influence matrix $\mathbf{R}(2)$. We purposely make the two configuration matrices different from each other. Partial data are shown in the left-hand side of Table 1. We use the algorithm in the section “Inference” to infer the dynamical influence model’s parameters \mathbf{V} , $\mathbf{R}(1:J)$, $\mathbf{E}^{1:c}$, $\mathbf{F}^{1:c}$. All parameters (including the emission distribution) are initialized randomly, and they are estimated automatically during the E-M process.

CHOOSING HYPERPARAMETERS

We now discuss the selection of hyperparameters J and p^V . For the number of active influence matrices J , we illustrate their characteristics by running the same example with $J = 3$. We show the poster distribution of r_t [calculated in (13)] in Figure 3(a) and (b). The dynamical influence model discovers the sudden change of influence weights accurately at $t = 200$. Since the toy process only has two true configuration matrices, the posterior probability of the third configuration being active is almost zero for any t . The system properties are fully captured by the other two configuration matrices during the training. The learned configuration matrices (shown in Table 1) are correctly recovered. Based on Figure 3(a) and experiments with other values for J (which we cannot show here due to the article’s space limitation), we suggest that readers should gradually increase J until the newly added configuration matrices are no longer useful in capturing additional dynamical information from the data, by ensuring there is no constant zero posterior probability as in Figure 3(b).

We also demonstrate convergence of the K-L Divergence between the true distributions of the transition probability and the learned distributions in Figure 3(c) with different values of p^V . As can be seen in Figure 3(c), the algorithm converges quickly within 50 iterations. However, when p^V is small, we may encounter over-fitting where the learned model rapidly

switches between different configurations to best suit the data. Therefore, in Figure 3(c), the divergence for $p^V = 0$ remains higher than other p^V values at convergence. In conclusion, we advise users to increase p^V gradually until the posterior of r_t does not fluctuate.

MODELING DYNAMICAL INFLUENCE IN GROUP DISCUSSIONS

DATA SET DESCRIPTION AND PREPROCESSING

Researchers in [5] recruited 40 groups with four subjects in each group for this experiment. During the experiment, each subject was required to wear a sociometric badge on their neck for audio recording [Figure 1(c)], and each group was required to perform two different group discussion tasks: a brainstorming task (referred to as BS) and a problem solving task (referred to as PS). Each task lasted for about three to ten minutes. We refer readers to the original paper [5] for details on data collection and experiment preparations.

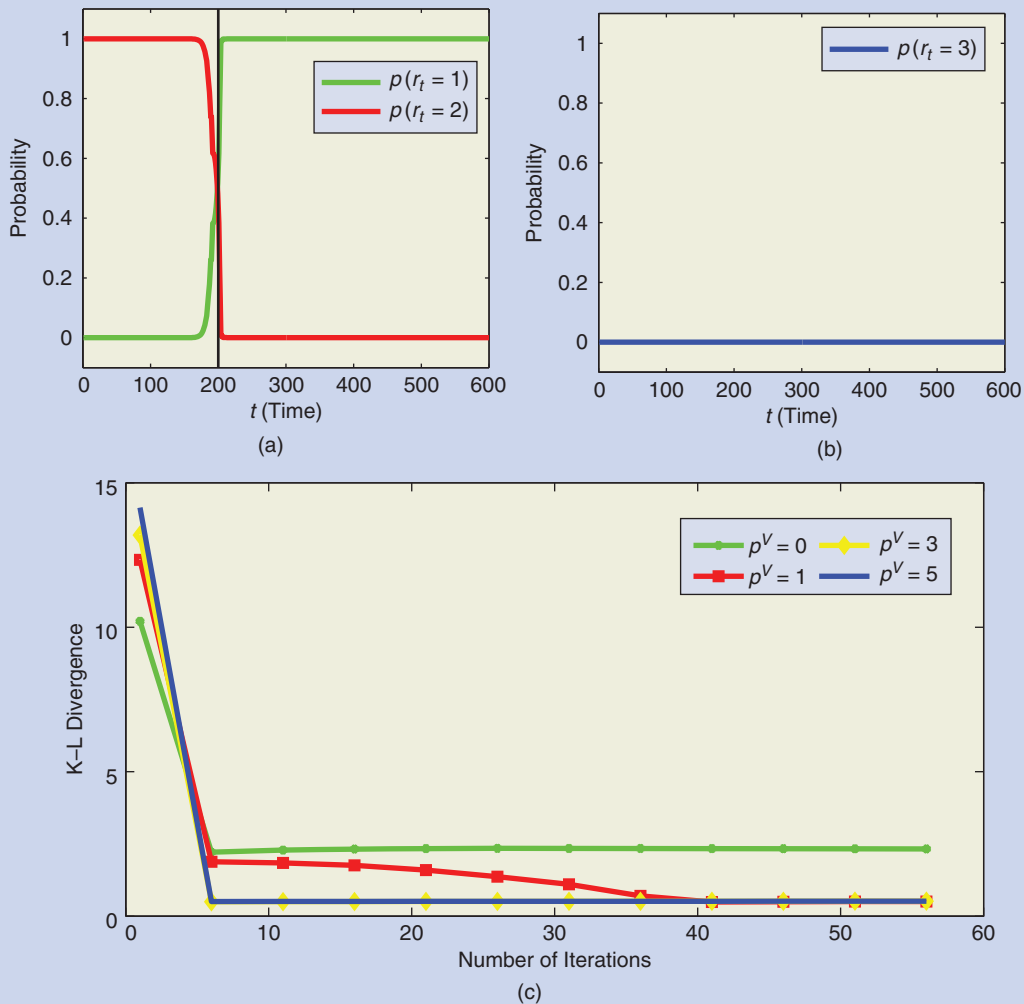
The groups were asked to perform these tasks in two different settings: one in which people were colocated in the same room around a table (referred to as CO), and one in which two pairs of people were placed in two separate rooms with only audio communication available between them (referred to as DS). The badges are deployed in both cases for audio collecting. We separated all samples according to their tasks (BS/PS) and their settings (CO/DS), and we ended up with four categories: DS+BS, DS+PS, CO+BS, and CO+PS. Since discussions were held in four-person groups, each sample for a discussion session is composed of four sequences collected by the four badges on participants’ chests. The audio sequence picked up by each badge was split into one-second blocks. Variances of speech energy were calculated for each block. We then applied a hard threshold to convert them into binary sequences. In all experiments, we only used binary sequences as data input.

PREDICTING TURN TAKING IN DISCUSSION

Here we explain an application of the dynamical influence model to predict turn taking, and we show that it is possible to achieve good accuracy in prediction given only the audio volume variance observations, with no information from the audio content.

Ten occurrences of turn-taking behavior from each sample are selected for prediction purposes. “Turn taking” here is defined as incidences in which the current speaker ceases speaking and another speaker starts to speak.

For the dynamical influence model, we model each person as an entity c , and the observed audio variances at time t as $O_t^{(c)}$. Each person also has two hidden states, representing speaking or not speaking. The hidden layer eliminates error due to noise and nonvoicing speaking in audio signals [16]. Therefore, influence here is set to capture how each person’s speaking/nonspeaking hidden states dynamically change other people’s speaking/nonspeaking states (i.e., how people influence each others’ turn taking). All parameters are initialized randomly and learned by the E-M inference algorithm in this



[FIG3] The posterior of r_t is shown with $J = 3$ after convergence. The middle black vertical line in (a) indicates the true switch in r_t . The probability of R(1) being active and R(2) being active are shown in (a); R(3) is shown in (b), which remains inactive. (c) The K-L divergence between learned parameters and the true distributions with respect to number of iterations.

example. We train the dynamical influence model using data up to $t - 1$, sample observations at time t from it, and mark the chain that changes the most toward the high-variance observations as the turn taker at t . The emission probability $\text{Prob}(O_t^{(c)} | h_t^{(c)})$ is modeled using a multinomial distribution and is estimated automatically during the E-M process.

For comparison, we also show results using TESLA and nearest neighbors methods. For TESLA, we use the official implementation [23] to obtain the changing weights between pairs of nodes, and we pick the node that has the strongest correlation weight to other nodes at $t - 1$ as the turn taker at t . To predict the turn taking at time t using the nearest neighbor method, we look over all previous instances of turn-taking behaviors that have the same speaker as the one in $t - 1$, and predict by using the most frequent outcomes.

The accuracy for each algorithm is listed in Table 2. We also show the prediction accuracy for the half of all samples that have more complex interactions, i.e., higher entropy. For

the dynamical influence approach, we list error rates for $J = 1$ (which is simply the influence model), $J = 2$, and $J = 3$. Except for DS+BS, we notice that the dynamical influence model outperforms others in all categories with different J . This performance is quite good considering that we are using only volume and that a human can only predict at around 50% accuracy for similar tasks [28].

More importantly, the dynamical influence model seems to perform much better than the competing methods for more complex interactions. For simple interactions, it seems that $J = 1$ or even the nearest neighbor (NN) method perform the best due to the fact that there is little shift in the influence structure during the discussion. However, when handling complex interaction processes, the introduction of a switching influence dynamic dramatically improves the performance as shown in Table 2. This result suggests that the dynamical influence assumption is reasonable and necessary in modeling complex group dynamics, and it can improve

[TABLE 2] ACCURACY FOR DIFFERENT TURN-TAKING PREDICTION METHODS ON BOTH THE FULL DATA SET AND THE HALF OF THE DATA SET WITH MORE COMPLEX INTERACTIONS. THE RANDOM GUESS ACCURACY IS 33%. HUMAN ACCURACY IS TYPICALLY AROUND 50% FOR SIMILAR TASKS [28].

METHODS	ACCURACY ALL SAMPLES				ACCURACY COMPLEX INTERACTION SAMPLES			
	DS+BS	DS+PS	CO+BS	CO+PS	DS+BS	DS+PS	CO+BS	CO+PS
TESLA	0.41	0.42	0.32	0.25	0.44	0.37	0.37	0.17
NN	0.58	0.60	0.48	0.50	0.47	0.47	0.38	0.26
OURS($J = 1$)	0.45	0.67	0.75	0.63	0.45	0.56	0.77	0.62
OURS($J = 2$)	0.46	0.58	0.65	0.34	0.47	0.58	0.67	0.46
OURS($J = 3$)	0.50	0.60	0.55	0.48	0.47	0.73	0.65	0.65

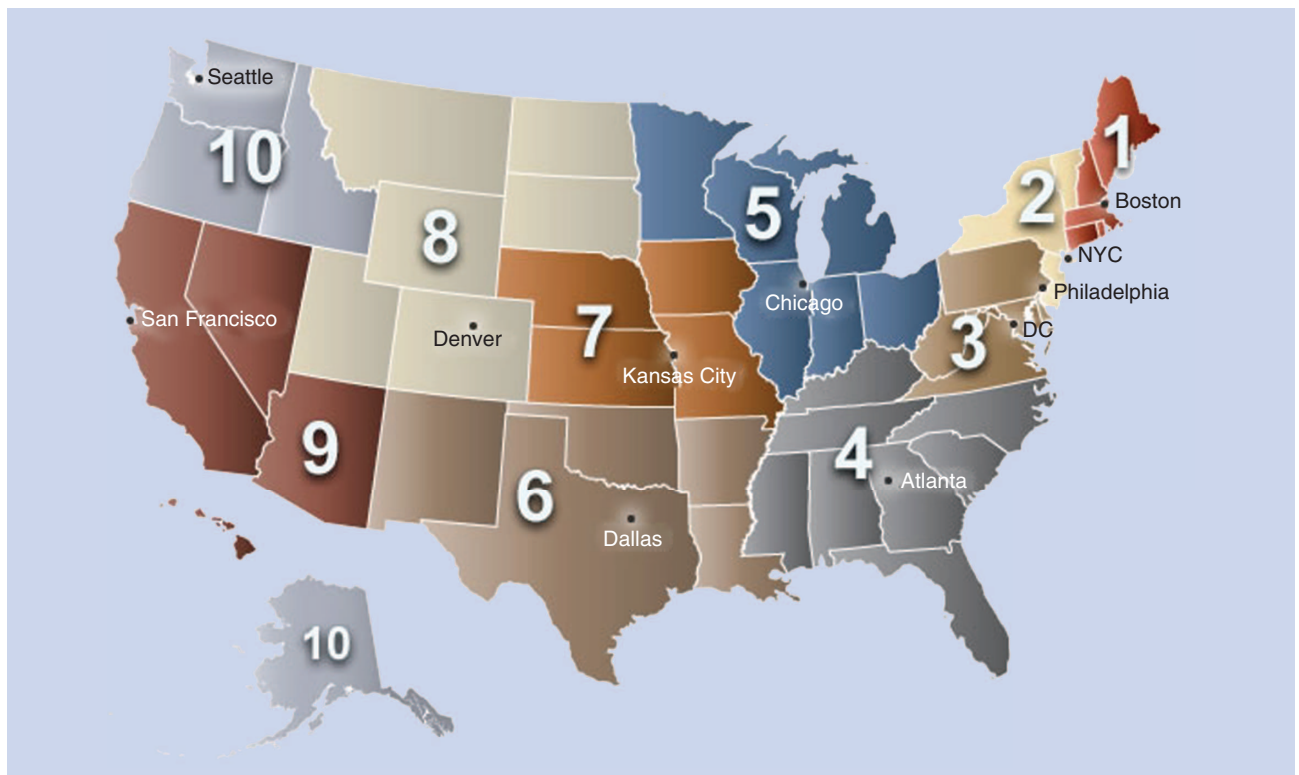
prediction accuracy significantly. However, in simple cases, the model achieves the highest performance when $J = 1$ (i.e., the influence matrix is static), and a higher J will only lead to overfitting.

The fact that turn taking is predictable using our dynamical influence assumption is indeed surprising, because group turn taking dynamics are complicated and related to content as well [4]. We think that the dynamical influence model tracks the two main mechanisms of group discussions noted in Gibson [4], even though the model does not incorporate the content of speaker statements. First, the time variant assumption in the dynamical influence model captures the latency factor in group dynamics. And second, our abstract conditional probability definition of influence is essentially a generalization of the conversational obligation mechanism that Gibson described.

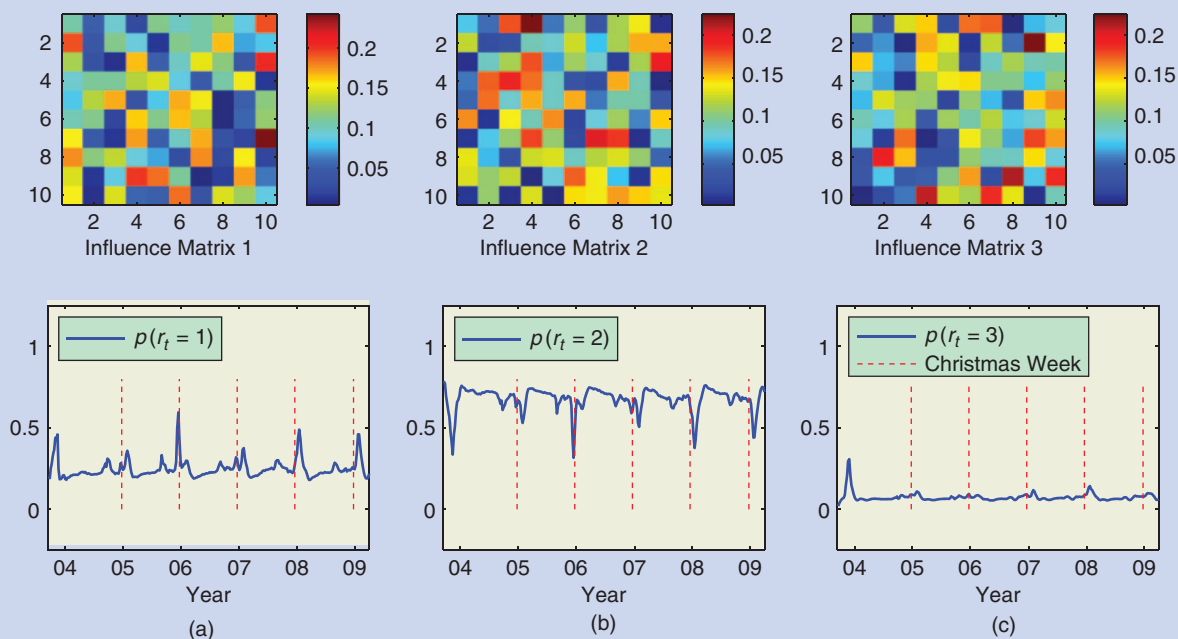
MODELING FLU EPIDEMICS AS INFLUENCE DYNAMICS

The last example is an application of the dynamical influence model to weekly U.S. flu activity data from Google Flu Trends [29]. All 50 states in the United States are divided into ten regions by their geolocation, as shown in Figure 4, and we model each region as an entity in the dynamical influence model.

As the data are continuous, six hidden states are used for each chain, and $p(O_i^{(c)}|h_i^{(c)})$ is modeled with six different Gaussian distributions with different means and the same variance for each hidden state. We set by hand the six mean values so that they represent the six different severe levels for the flu epidemics, from the least severe to the most severe. We train the model using the first 290 weeks (from 2003 to early 2009), and we show the posterior for r_b , the switching parameter, in Figure 5 together with the three learned influence matrices.



[FIG4] Ten regions of the United States defined by the U.S. Department of Health and Human Services (figure used with permission).



[FIG5] The inferred posterior for r_t , given all observations after convergence is shown here. While there are many small peaks indicating changes in influence, the largest peaks occur during the Christmas holiday season, which implies holiday traffic patterns can have a big effect on flu transmissibility between regions. We find that the three configuration matrices (a)–(c) are good enough to capture the flu dynamics.

While there are many small peaks suggesting changes in influence, the probability changes dramatically around Christmas, which suggests that the influence patterns among these ten regions are very different during the holiday season. Note especially that we did not tell the model to search for different patterns on those days—instead, it reveals the fact that transmission dynamics operate differently during a time when many people are engaging in once-a-year travel to visit family and friends. While it is possible that other mechanisms are at work, alternative explanations like the weather are not plausible because they do not change discontinuously during the holiday season.

Influence matrix 1 captures the dynamics during holiday seasons, while influence matrix 2 captures the dynamics during normal seasons. Row i corresponds to the region i in Figure 4. Let's take an example of Row 1, the New England region. Normally (as shown in the first row of influence matrix 2), New England is more likely to be influenced by close regions such as 3 and 4; during the holiday season, New England is more likely to be influenced by all regions, especially distant regions such as region 9. The same phenomena exist for other regions as well.

DISCUSSION

In this article, we have described the influence model and its generalization, the dynamical influence model, and we showed how these can be applied to a variety of social signals to infer how entities in networks affect one other. In particular, we can use the resulting influence matrix \mathbf{R} to connect the underlying

network and the stochastic process of state transition. The switching matrices $\mathbf{R}(1), \dots, \mathbf{R}(J)$ are even able to bridge the state transition to time-varying networks.

The influence model shares the same issues with other machine learning models: inference requires sufficient training data, and tuning is necessary for best results. Future work includes combining known network data into the model to boost performance.

The most important limitation is that we are attempting to infer causal processes from observational data in which many mechanisms are likely at play. If we find that behavior between two individuals is correlated, it could be due to influence, but it could also be due to selection (I choose to interact with people like me) or to contextual factors (you and I are both influenced by an event or a third party not in the data). It has been recently shown that these mechanisms are generically confounded [11], but it is important to remember that this does not make observational data worthless. It just means that we should carefully consider alternative mechanisms that may underlie correlated behavior. The fact that we have time data to test causal ordering and we have asymmetries in network relationships to test direction of effects means that we can have greater (but not complete) confidence than we would if we only had cross-sectional data from symmetric relationships.

ACKNOWLEDGMENTS

This research was sponsored by AFOSR under award number FA9550-10-1-0122 and ARL under Cooperative Agreement

number W911NF-09-2-0053. The views and conclusions contained in this article are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of AFOSR, ARL, or the U.S. Government.

AUTHORS

Wei Pan (panwei@media.mit.edu) is a Ph.D. student in the Human Dynamics Group at the MIT Media Lab. He obtained his B.Eng. degree in computer science from Tsinghua University in Beijing, China. His current research interest is computational social science. He has published over 20 peer-reviewed papers and has had some of his work featured in mainstream media. He and his team recently won the DARPA Network Challenge.

Wen Dong (wdong@media.mit.edu) focuses on modeling human intelligence with network dynamics theory. He has published more than ten papers on the network influence model, group dynamics, and human problem solving, and he has consulted with several business worldwide in constructing successful data products. He received his Ph.D. degree from MIT Media Laboratory. His thesis is "Modeling the Structure of Collective Intelligence."

Manuel Cebrian (mcebrian@ucsd.edu) is a research scientist at UCSD and affiliate with the MIT Media Lab. Previously, he worked at Telefonica Research (2008–2009) and Brown University (2007–2008). He received a Ph.D. degree in computer science from Universidad Autonoma de Madrid (2007).

Taemie Kim (taemie@media.mit.edu) received her Ph.D. degree from the MIT Media Lab. Her research focuses on understanding and improving distributed collaboration using sociometric sensors. Her research is centered around computer supported collaborative work, utilizing disciplines of interaction design, machine learning, organizational behavior, and social psychology. She earned a B.S. degree in electrical engineering and computer science from KAIST, Korea, and an M.S. degree in electrical engineering from Stanford University.

James H. Fowler (jhfowler@ucsd.edu) is a social scientist who received his Ph.D. and B.A. degrees from Harvard and has M.A. degrees from both Harvard and Yale. He is currently a professor of medical genetics and political science at UCSD, and he was recently named a Fellow of John Simon Guggenheim Foundation. He has published more than 50 peer-reviewed articles in journals that span the natural and social sciences, and he is coauthor of *Connected: The Surprising Power of Our Social Networks and How They Shape Our Lives*, which has been published in 20 languages.

Alex (Sandy) Pentland (pentland@media.mit.edu) is a pioneer in organizational engineering, mobile information systems, and computational social science. He directs the MIT Human Dynamics Lab and the Media Lab Entrepreneurship Program, advises the World Economic Forum, and is cofounder of the Institute for Data Driven Design (<http://idcubed.org>). He is among the most-cited computer scientists in the world. His most recent book, *Honest Signals: How They Shape Our World*, was published in 2008 by the MIT Press.

REFERENCES

- [1] D. J. Watts and P. S. Dodds, "Influentials, networks, and public opinion formation," *J. Consumer Res.*, vol. 34, no. 4, pp. 441–458, 2007.
- [2] A. W. Woolley, C. F. Chabris, A. Pentland, N. Hashmi, and T. W. Malone, "Evidence for a collective intelligence factor in the performance of human groups," *Science*, vol. 330, no. 6004, pp. 686–688, 2010.
- [3] C. F. Manski, "Identification of endogenous social effects: The reflection problem," *Rev. Econ. Stud.*, vol. 60, no. 3, pp. 531–542, 1993.
- [4] D. R. Gibson, "Making the turn," *Social Psychol. Quart.*, vol. 73, no. 2, pp. 132–151, 2010.
- [5] T. Kim, A. Chang, L. Holland, and A. S. Pentland, "Meeting mediator: Enhancing group collaboration using sociometric feedback," in *Proc. ACM 2008 Conf. Computer Supported Cooperative Work*. New York: ACM, 2008, pp. 457–466.
- [6] N. Eagle and A. Pentland, "Reality mining: Sensing complex social systems," *Pers. Ubiquit. Comput.*, vol. 10, no. 4, pp. 255–268, 2006.
- [7] S. Basu, T. Choudhury, B. Clarkson, and A. Pentland, "Learning human interactions with the influence model," MIT Media Lab., Cambridge, MA, Tech. Rep. 539, 2001.
- [8] C. Castellano and V. Loreto, "Statistical physics of social dynamics," *Rev. Modern Phys.*, vol. 81, no. 2, pp. 591–646, 2009.
- [9] N. E. Friedkin, "The attitude-behavior linkage in behavioral cascades," *Social Psychol. Quart.*, vol. 73, no. 2, pp. 196–213, 2010.
- [10] D. Acemoglu, A. Ozdaglar, and A. Parandeh-Gheibi, "Spread of (MIS) information in social networks," *Games Econ. Behav.*, pp. 194–227, 2010.
- [11] C. R. Shalizi and A. C. Thomas. (2010). Homophily and contagion are generically confounded in observational social network studies. [Online]. Available: Arxiv preprint arXiv:1004.4704
- [12] C. M. Bishop, *Pattern Recognition and Machine Learning*. New York: Springer-Verlag, 2006.
- [13] M. Brand, N. Oliver, and A. Pentland, "Coupled hidden Markov models for complex action recognition," in *Proc. IEEE Computer Society Conf. Computer Vision and Pattern Recognition*, 1997, pp. 994–999.
- [14] W. Dong and A. Pentland, "AI for human computing," in *Modeling Influence Between Experts (Lecture Notes in Computer Science 4451)*. Berlin: Springer-Verlag, 2007, pp. 170–189.
- [15] A. Madan and A. Pentland, "Modeling social diffusion phenomena using reality mining," in *Proc. AAAI Spring Symp. Human Behavior Modeling*, Palo Alto, CA, 2009, pp. 43–48.
- [16] T. Choudhury and S. Basu, "Modeling conversational dynamics as a mixed memory Markov process," in *Proc. Intl. Conf. Neural Information and Processing Systems (NIPS)*, 2004, pp. 281–288.
- [17] W. Dong, B. Lepri, A. Cappelletti, A. S. Pentland, F. Pianesi, and M. Zancanaro, "Using the influence model to recognize functional roles in meetings," in *Proc. 9th Intl. Conf. Multimodal Interfaces*. New York: ACM, 2007, pp. 271–278.
- [18] W. Dong and A. Pentland, "A network analysis of road traffic with vehicle tracking data," in *Proc. AAAI Spring Symp. Human Behavior Modeling*, Pal Alto, CA, 2009, pp. 7–12, 2009.
- [19] W. Pan, M. Cebrian, W. Dong, T. Kim, and A. Pentland. (2010). Modeling dynamical influence in human interaction patterns. [Online]. Available: Arxiv preprint arXiv:1009.0240
- [20] D. Zhang, D. Gatica-Perez, S. Bengio, and D. Roy, "Learning influence among interacting Markov chains," in *Proc. Advances in Neural Information Processing Systems 18 (NIPS 2005)*, pp. 1577–1584.
- [21] S. A. Ansari, V. S. Springthorpe, and S. A. Sattar, "Survival and vehicular spread of human rotaviruses: Possible relation to seasonality of outbreaks," *Rev. Infect. Dis.*, vol. 13, no. 3, pp. 448–461, 1991.
- [22] F. Guo, S. Hanneke, W. Fu, and E. P. Xing, "Recovering temporally rewiring networks: A model-based approach," in *Proc. 24th Intl. Conf. Machine Learning*. New York: ACM, 2007, pp. 321–328.
- [23] A. Ahmed and E. P. Xing, "Recovering time-varying networks of dependencies in social and biological studies," *Proc. Nat. Acad. Sci.*, vol. 106, no. 29, pp. 11878–11883, 2009.
- [24] M. G. Rodriguez, J. Leskovec, and A. Krause, "Inferring networks of diffusion and influence," in *Proc. 16th ACM SIGKDD Intl. Conf. Knowledge Discovery and Data Mining*. New York: ACM, 2010, pp. 1019–1028.
- [25] W. Dong and A. Pentland, "Multi-sensor data fusion using the influence model," in *Proc. Intl. Workshop Wearable and Implantable Body Sensor Networks*, 2006, pp. 75–78.
- [26] W. Dong, "Modeling the structure of collective intelligence," Ph.D. dissertation, Media Lab, Massachusetts Inst. Technol., Cambridge, MA, 2010.
- [27] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul, "An introduction to variational methods for graphical models," *Mach. Learn.*, vol. 37, no. 2, pp. 183–233, 1999.
- [28] D. Schlangen, "From reaction to prediction: Experiments with computational models of turn-taking," in *Proc. 9th Intl. Conf. Spoken Language Processing*, 2006.
- [29] J. Ginsberg, M. H. Mohebbi, R. S. Patel, L. Brammer, M. S. Smolinski, and L. Brilliant, "Detecting influenza epidemics using search engine query data," *Nature*, vol. 457, no. 7232, pp. 1012–1014, 2008.