Modeling Dynamical Influence in Human Interaction Patterns

Wei Pan
MIT Media Lab
Cambridge, MA 02139
panwei@media.mit.edu

Manuel Cebrian
MIT Media Lab
Cambridge, MA 02139
cebrian@mit.edu

Wen Dong
MIT Media Lab
Cambridge, MA 02139
wdong@media.mit.edu

Taemie Kim
MIT Media Lab
Cambridge, MA 02139
taemie@media.mit.edu

Alex (Sandy) Pentland
MIT Media Lab
Cambridge, MA 02139
sandy@media.mit.edu

Abstract

We present a new way for modeling 1) *social influence* and 2) the well-observed property of social influence – the influence strength between individuals changes over time (e.g., friendships break and reform). We show that our unsupervised generative switching Bayesian approach can simultaneously captures the system dynamics as the outcome of both (i) the influence between individuals (each modeled as an HMM), and (ii) the changes of influence itself using only individual observations. We describe here a variational Expectation-Maximization (EM) algorithm for inference. In our experiments, we illustrate applications with synthetic and real data of detecting structural change, predicting turn taking by analyzing a real group discussion behavior dataset and understanding flu influence patterns between US states. Results demonstrate that our approach is a strong alternative for modeling complex interacting social systems. [1]

1 Introduction

Our proposed model tackles the problem of analyzing and understanding *who influences whom* in a social system, such as a group discussion process, which has been an interesting question for social scientists for the last six decades [1]. Influence is also interesting in the context of leadership where the influence between one another has been recognized as a significant factor of group performance [2]. However, it remains a difficult question to define and model the concept of *influence* in a formal mathematical way.

In this paper, we handle this problem by modeling each agent in a social system as a Hidden Markov Chain with a finite set of states, and all chains interact with each other according to a family of influence configurations, each of which describes a different interaction pattern among nodes. Influence between two agents is modeled as how the current state of one agent can effect the future states of the other agents.

In the prevailing studies on social computing, quantitative efforts have focused on the *static* picture of the influence [3] [4], namely who is influencing whom in a social system when longitudinal data on human interactions is aggregated in a snapshot. However, there is extensive evidence leading us to think that influence is indeed a dynamical process[5][6]. This can also be seen from many real-world experiences: Friendship is not static, and the person who currently possesses the most

[1] A full version of this paper can be found at `http://arxiv.org/abs/1009.0240`.

influence over you may be different after some time; In a tedious negotiation with many parties involved, your most active opponent may change due to topic shift and strategy shift over time... Therefore, we believe that, in a social system such as a group discussion session, the influence between subjects fluctuates as well, and a better model should take the changes of influence itself into consideration.

Our approach is in essence a switching version of influence model [7], a special type of Bayesian network. We are interested in the challenge of inferring influence and learning parameters in a social system based solely on individual observations over time, i.e., without actually knowing the individual interaction patterns. We define influence between two nodes as the conditional probability between the internal states of these two nodes in consecutive time frames. We approach the problem of dynamical influence by introducing a family of different influence configurations, each of which captures a different interaction pattern among nodes. A latent trace r_t is included also to represent the index of the current active influence configuration at time t. r_t gradually switches between different influence configurations, and itself is treated as a stochastic process as well. Therefore, our model not only captures the dynamics of individual behaviors, but also the underlying latent variables tracing changes in influence. It should be noted that in our approach system dynamics and changing influence are learned simultaneously in a unified framework, and the learning algorithm is unsupervised.

We are fully aware of the class of time-varying models: from EGRM [16] to TESLA [17], to name two. We model the dynamics of every node and edge in a network instead of feature functions as in EGRM. Our work is also significantly different from TESLA: We consider changing influence, which are continuous real values, as the topological dynamics. As shown in our turn-taking experiment, our generative model is more suitable in capturing the interaction of nodes and the dynamics of the interaction strength simultaneously.

2 Our Model

Our approach, the Dynamical Influence Process, is a switching extension to the existing influence model [9]. It is composed of C interacting chains. In this model, similar to an HMM, each chain $c \in \{1, \ldots, C\}$ takes one of a finite number of latent states at any discrete time t: $h_t^{(c)} \in \{1, \ldots, S\}$. Corresponding to each latent state $h_t^{(c)}$, we observe $O_t^{(c)}$ which follows a conditional probability distribution Prob$(O_t^{(c)}|h_t^{(c)})$, usually known as the emission probability in HMM literature. In practice, it can either be multinomial for discrete observations or Gaussian mixture for continuous observations.

We proceed to describe the cross-chain interactions of this system. $\mathbf{M}_{i,j}$ denotes the ith row and jth column of matrix $\mathbf{M}$ in the following discussion. In this model, we have J different interaction configurations described by matrices $\mathbf{R}^1, ..., \mathbf{R}^J$ and only one configuration $r_t \in \{1, \ldots, J\}$ is active at time t. J is a hyper-parameter set by the user, and we cover the detail of selecting the values of hyper-parameters in the full version[19]. The model changes its interaction configurations $\{r_t\}_{t=1,2,...}$ slowly with respect to the sampling period according to the following Markov prices:

$$r_{t+1}|r_t \sim \text{multi}(\mathbf{V}_{r_t,1}, \cdots, \mathbf{V}_{r_t,J}), \tag{1}$$

where $\mathbf{V}$ is constrained by another hyper-parameter $p^V, p^V > 0$. A large p^V will ensure that our model switch slowly to other influence configurations and tend to remain in the current configuration:

$$(\mathbf{V}_{r_t,1}, \ldots, \mathbf{V}_{r_t,J}) \sim \text{Dirichlet}(\underset{\substack{\uparrow \\ 1,}}{10^0}, \underset{\substack{\uparrow \\ 2,}}{10^0}, \ldots, \underset{\substack{\uparrow \\ r_t,}}{10^{p^V}}, \ldots, \underset{\substack{\uparrow \\ J}}{10^0}). \tag{2}$$

Given that the interaction configuration r_t is in effect, the latent state of chain c in this system at time $t+1$ is determined by another random chain $q_t^{(c)} \in \{1, \ldots, C\}$ according to a multinomial distribution described by the configuration matrix $\mathbf{R}^{r_t}$. This is similar to mixture models such as the Gaussian mixture model [18]:

$$q_t^{(c)}|r_t \sim \text{multi}(\mathbf{R}_{c,1}^{r_t}, ..., \mathbf{R}_{c,C}^{r_t}). \tag{3}$$

The state of chain c at $t+1$ thus is determined by:

$$h_t^{(c)}|h_{t-1}^{(1)},\dots,h_{t-1}^{(C)},q_t^{(c)}=c' \sim \begin{cases} \text{multi}(\mathbf{E}^{(c)}_{h_t^{(c)},1},\dots,\mathbf{E}^{(c)}_{h_t^{(c)},S}) & c'=c \\ \text{multi}(\mathbf{F}^{(c')}_{h_t^{(c')},1},\dots,\mathbf{F}^{(c')}_{h_t^{(c')},S}) & c'\neq c \end{cases}. \tag{4}$$

We present some intuitions to the model description here: Eq. 3 and Eq. 4 define the concept of influence in our model. The intuition is that at each time t each chain c will sample $q_t^{(c)}$ from Eq. 3 to decide which chain will influence it at $t+1$. Notice that since the model is dynamical and r_t is changing, the distribution of $q_t^{(c)}$ is different at different t as we are sampling from different configuration matrices at different t. This is how changing influence is captured by switching between influence configurations in our model. If $q_t^{(c)}$ happens to be the same chain c, we will use transition matrice $\mathbf{E}^{(c)}$ and chain c's current state to determine its state at $t+1$; if $q_t^{(c)}=c'\neq c$, we will use transition matrice $\mathbf{F}^{(c')}$ and chain c''s current state to determine chain c's state at $t+1$. $\mathbf{E}^{(c)}$ and $\mathbf{F}^{(c)}$ are both $S\times S$ matrices, and they are similar to the transition matrix in HMM literature.

Given the model description, the likelihood function is

$$\mathcal{L}(O,h,q,r|\mathbf{E},\mathbf{F},\mathbf{R},\mathbf{V}) \tag{5}$$

$$= \prod_{t=2}^{T}\Big\{\text{Prob}(r_t|r_{t-1})\prod_{c=1}^{C}\Big[\text{Prob}(O_t^{(c)}|h_t^{(c)})\text{Prob}(h_t^{(c)}|h_{t-1}^{(1,\dots,C)},q_{t-1}^{(c)})\text{Prob}(q_t^{(c)}|r_t)\Big]\Big\}$$

$$\times\prod_{c=1}^{C}\text{Prob}(O_1^{(c)}|h_1^{(c)})\text{Prob}(h_1^{(c)})\text{Prob}(r_1). \tag{6}$$

The model is intractable. Therefore we use the variational method for inference. We kindly invite the reader to the full version of this paper for detail in model learning[19].

3 Experiments on Human Interaction Data

We here demonstrate an application of using our dynamical influence process to predict turn taking– who will speak next in the interaction process, and we show that it is possible to achieve good accuracy in prediction given only the binary audio volume variance observations, with no information from the audio content. The detail of the experiment set-up can be found in our full paper[19]. We also implement two competing methods: a)Nearest Neighbor(NN) and b) TESLA [21]. We here present the results in Table 1.

The dataset used in this experiment comes from a group discussion experiment in [20]. Researchers in [20] recruited 40 groups with four subjects in each group for this experiment. During the experiment, each subject was required to wear the sociometric badge on their necks for audio recording, and each group was required to perform two different group discussion tasks: a brainstorming task (BS) and a problem solving (PS) task in two different settings: (a) being *co-located*(CO) in the same room around a table and (b) being *distributed*(DS) in two rooms with only audio communication being available between the pairs. (The badge is deployed in both cases for audio collecting.)

The accuracy for each algorithm is listed in Table 1. We also show the prediction accuracy for the half of all samples that have more complex interactions, i.e., higher entropy. For our dynamical influence based approach, we list error rates for $J=1,2$ and 3. Except DS+BS, We notice that our algorithm 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[22].

More importantly, our model seems to perform much better than the competing methods for more complex interactions. For simple interactions, it seems that $J=1$ or even NN perform the best due to the fact that there is little shift in influence structure during the discussion. However, when handling complex interaction processes, the introduction of a switching influence dynamics improves the performance as shown in Table 1. Our results suggest that the dynamical influence assumption in our model is reasonable and necessary in modeling complex group dynamics, and in one case it

Table 1: Accuracy for different turn taking prediction methods on both the full dataset and the half of the dataset with more complex interactions. The random guess accuracy is 33%. Human accuracy is typically around 50% for similar tasks[22].

	ACCURACY ALL SAMPLES				**ACCURACY COMPLEX INTERACTION SAMPLES**			
METHODS	**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**

can improve prediction accuracy to above 60% for PS tasks. However, in simple cases, the model achieves the highest performance only when $J = 1$, i.e. the influence is static, and a higher J will only lead to overfitting.

We also demonstrate how to use the dynamical influence model for understanding the seasonal difference in flu spreading and flu prediction. The detail can be found in our full paper[19].

4 Conclusions

We have developed an unsupervised generative model that captures the system dynamics as the outcome of both the influence between individuals and the dynamics of influence. Our model directly tackles the important sociological question of analyzing who influence whom in social systems. In our model, C HMM chains interact with each other according to a family of influence matrices describing different interaction patterns, and switch between them over time. A fast variational inference scheme is also developed to handle large system such as the flu epidemic in US. We have demonstrated the performance of our model in applications of detecting structural change, predicting turn taking and understanding epidemic dynamics. Our model provides a new perspective for modeling dynamics in network structures and influence structures.

References

[1] E. Katz and P.F. Lazarsfeld. *Personal influence*. Free Pr., 1955.

[2] G.F.. Farris and F. Lim. Effect of performance on leadership, cohesiveness, influence, satisfaction and subsequent performance. *Journal of Applied Psychology*, 53(6):490–497, 1969.

[3] M. Brand, N. Oliver, and A. Pentland. Coupled hidden Markov models for complex action recognition. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 994–999, 1997.

[4] W. Dong and A. Pentland. Modeling influence between experts. *Lecture Notes in Computer Science*, 4451:170, 2007.

[5] S.A. Ansari, V.S. Springthorpe, and S.A. Sattar. Survival and vehicular spread of human rotaviruses: possible relation to seasonality of outbreaks. *Reviews of infectious diseases*, 13(3):448–461, 1991.

[6] J.P. Onnela, J. Saramaki, J. Hyvonen, G. Szabó, D. Lazer, K. Kaski, J. Kertész, and A.L. Barabási. Structure and tie strengths in mobile communication networks. *Proceedings of the National Academy of Sciences*, 104(18):7332, 2007.

[7] 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 *Proceedings of the 9th international conference on Multimodal interfaces*, pages 271–278. ACM, 2007.

[8] T. Choudhury and S. Basu. Modeling conversational dynamics as a mixed memory markov process. In *Proc. of Intl. Conference on Neural Information and Processing Systems (NIPS)*. Citeseer, 2004.

[9] S. Basu, T. Choudhury, B. Clarkson, A. Pentland, et al. Learning human interactions with the influence model. *MIT Media Laboratory Technical Note*, 2001.

[10] D. Zhang, D. Gatica-Perez, S. Bengio, and D. Roy. Learning Influence Among Interacting Markov Chains. 2005.

[11] E. B. Fox, E. B. Sudderth, M. I. Jordan, and A. S. Willsky. Nonparametric Bayesian learning of switching linear dynamical systems. In *Neural Information Processing Systems 21*. MIT Press, 2009.

[12] W. Pan and L. Torresani. Unsupervised hierarchical modeling of locomotion styles. In *Proceedings of the 26th Annual International Conference on Machine Learning*. ACM New York, NY, USA, 2009.

[13] M.I. Jordan. *Learning in graphical models*. Kluwer Academic Publishers, 1998.

[14] R.J. Weiss and D.P.W. Ellis. A variational EM algorithm for learning eigenvoice parameters in mixed signals. In *Proceedings of the 2009 IEEE International Conference on Acoustics, Speech and Signal Processing-Volume 00*, pages 113–116. IEEE Computer Society, 2009.

[15] Z. Ghahramani and M.I. Jordan. Factorial hidden Markov models. *Machine learning*, 29(2):245–273, 1997.

[16] F. Guo, S. Hanneke, W. Fu, and E.P. Xing. Recovering temporally rewiring networks: A model-based approach. In *Proceedings of the 24th international conference on Machine learning*, page 328. ACM, 2007.

[17] A. Ahmed and E.P. Xing. Recovering time-varying networks of dependencies in social and biological studies. *Proceedings of the National Academy of Sciences*, 106(29):11878, 2009.

[18] C.M. Bishop et al. *Pattern recognition and machine learning*. Springer New York:, 2006.

[19] W. Pan, M. Cebrian, W. Dong, T. Kim, and A. Pentland. Modeling Dynamical Influence in Human Interaction Patterns. *Arxiv preprint arXiv:1009.0240*, 2010.

[20] T. Kim, A. Chang, L. Holland, and A.S. Pentland. Meeting mediator: enhancing group collaborationusing sociometric feedback. In *Proceedings of the ACM 2008 conference on Computer supported cooperative work*, pages 457–466. ACM, 2008.

[21] A. Ahmed and E.P. Xing. Recovering time-varying networks of dependencies in social and biological studies. *Proceedings of the National Academy of Sciences*, 106(29):11878, 2009.

[22] D. Schlangen. From reaction to prediction: Experiments with computational models of turn-taking. In *Ninth International Conference on Spoken Language Processing*. Citeseer, 2006.

[23] 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*, 457(7232):1012–1014, 2008.

[24] S.M. Siddiqi, B. Boots, and G.J. Gordon. A constraint generation approach to learning stable linear dynamical systems. *Advances in Neural Information Processing Systems*, 2007.

[25] M. Shaw. *Group Dynamics: The Psychology of Small Group Behavior*. McGraw-Hill, 1976.

[26] J.M. DiMicco, KJ Hollenbach, A. Pandolfo, and W. Bender. The Impact of Increased Awareness while Face-to-Face. *Human-Computer Interaction*, 22(1), 2007.

[27] M.C. González, C.A. Hidalgo, and A.L. Barabási. Understanding individual human mobility patterns. *Nature*, 453(7196):779–782, 2008.