Learning and Predicting Dynamic Network Behavior with Graphical Multiagent Models (Extended Abstract)

Quang Duong*, Michael P. Wellman*, Satinder Singh* and Michael Kearns†
*Computer Science and Engineering, University of Michigan
†Computer and Information Sciences, University of Pennsylvania

Abstract

Factored models of multiagent systems address the complexity of joint behavior by exploiting locality in agent interactions. *History-dependent graphical multiagent models* (hGMMs) further capture dynamics by conditioning behavior on history. The hGMM framework also brings new elements of strategic reasoning and more expressive powers to modeling information diffusion over networks. We propose a greedy algorithm for learning hGMMs from time-series data, inducing both graphical structure and parameters. To evaluate this learning method, we employ human-subject experiment data for a voting consensus scenario, where agents on a network attempt to reach a unanimous vote. We empirically show that the learned hGMMs directly expressing joint behavior outperform alternatives in predicting dynamic voting behavior.

I. Introduction

There has been much interest in modeling and analyzing social behaviors that facilitate the diffusion of information in various online network scenarios. Recommendation network studies for example have identified node-centric and network-related factors dictating the spread of recommendations on videos and books across social networks [15, 14, 12]. Within the realm of electronic commerce, some researchers proposed to model one's product endorsement decisions as a function of his or her friends's choosing the same action [16, 18]. Similarly Bakshy et al. [1] employed information about individuals' adoption rates in capturing the spread of information in the virtual world of Second Life.

These research works mostly focus on online activities, and often fail to incorporate offline interactions. As the expressive power to specify local joint behavior provides advantages over models that assume conditional independence [5], the absence of offline interaction data in online social studies presents a need for directly modeling joint actions. Even if nodes or agents make decisions independently, conditioning actions on each other's prior decisions or on commonly observed history induces interdependencies over time. Moreover, most network research does not represent and capture the nodes' strategic reasoning. We have addressed these concerns in the context of multiagent behavior modeling by introducing the *history-dependent graphical multiagent models* (hGMMs) that express multiagent behavior on a graph, and capture dynamic relations and strategic decisions by conditioning action on history [5].

However, information the modeler may have about the agents' interaction network, such as scientific research partnerships or Facebook's friendship links, is not definitive, as there is no inherent reason that the interaction graph should constitute the ideal structure for a predictive graphical model for agent behavior. Even though actual agent behavior is naturally conditioned on its observable history (as captured by the interaction graph), once we abstract the history representation it may well turn out that non-local historical activity provides more useful predictive information. Moreover, these networks may be too complex for practical computation without imposing strong independence assumptions on behavior. We thus consider learning the graphical structure a necessary part of the modeling effort. As the interaction network structures and the hGMM interaction graphs do not necessarily overlap, learning the hGMM graphical structures can not only help bolster the hGMM's predictive power, but also provide insights on human subjects' behavior on a network. We thus consider learning the graphical structure a necessary part of the modeling effort.

We motivate and empirically evaluate our learning technique with the voting consensus experiments conducted by Kearns et al. [10]. The human subjects in these experiments were arranged on a network, specifying for each subject (also called *player*, or *agent*) the set of other players whose voting decisions he or she can observe. Each agent chooses to vote either blue (0) or red (1), and can change votes at any time. The scenario terminates when: (i) agents converge on action $a \in \{0, 1\}$, in which case agent i receives reward $r_i(a) > 0$, or (ii) they cannot agree by the time limit T, in which case rewards are zero. Agents may have different preferences for the available vote options. As nobody gets any reward without a unanimous vote, agents have to balance effort to promote their own preferred outcomes against the common goal to reach consensus, while taking into account their neighbors' voting patterns. Figure 1 illustrates the dynamic behavior of an example voting experiment network.

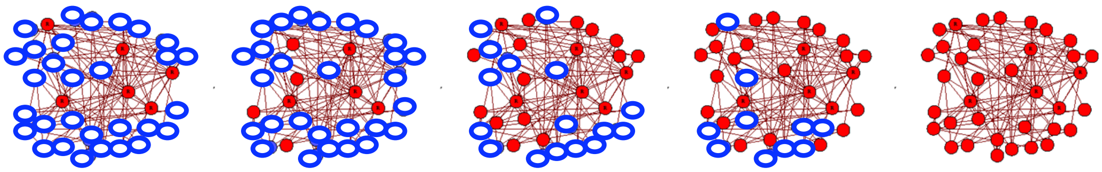

Fig. 1. Time snapshots of an experiment run where the densely connected minority group (red) exerts strong influences on others' votes [10].

Our aim is to capture dynamic voting behavior, which facilitates the spread of outcome preference information across the network. We propose a greedy algorithm for learning the graphical structure and parameters of an hGMM that can effectively and compactly capture joint dynamic behavior. We then empirically investigate the learned models' predictions of voting behavior and compare their performance with those of different baseline multiagent models, and demonstrate that models expressing joint behavior outperform the alternatives in predicting voting behavior.

II. HISTORY-DEPENDENT GRAPHICAL MULTIAGENT MODELS

We model behavior of n nodes or agents over a time interval divided into discrete periods, $[0, \ldots, T]$. At time t, agent $i \in \{1, \ldots, n\}$ chooses an action a_i^t from its action domain, A_i, according to its *strategy*, σ_i. Agents observe others' and their own past actions, as captured in *history* H^t, up to time t. Limited memory capacity or other computational constraints restrict an agent to focus on a subset of history H_i^t considered in its probabilistic choice of next action: $a_i^t \sim \sigma_i(H_i^t)$.

A *history-dependent graphical multiagent model* (hGMM) [5], $hG = (V, E, A, \pi)$, is a graphical model with graph elements V, a set of vertices representing the n agents, and E, edges capturing pairwise action dependencies between them. $A = (A_i, \ldots, A_n)$ represents the action domains, and $\pi = (\pi_1, \ldots, \pi_n)$ potential functions for each agent. The graph defines a neighborhood for each agent i: $N_i = \{j \mid (i,j) \in E\} \cup \{i\}$, including i and its neighbors $N_{-i} = N_i \setminus \{i\}$. The hGMM captures agent interactions in dynamic scenarios by conditioning joint agent behavior on an abstracted history of actions H^t. The history available to agent i, $H^t_{N_i}$, is the subset of H^t pertaining to agents in N_i. Each agent i is associated with a potential function $\pi_i(a^t_{N_i} \mid H^t_{N_i})$: $\prod_{j \in N_i} A_j \to R^+$, which specifies a local action configuration's likelihood of being included in the global outcome, conditioning on history. Specifically, the joint distribution of the system's actions taken at time t is the product of neighbor potentials [2, 5]:

$$\Pr(a^t \mid H^t) = \frac{\prod_i \pi_i(a^t_{N_i} \mid H^t_{N_i})}{Z}. \quad (1)$$

We construct and examine four multiagent behavior model forms for capturing voting behavior dynamics in the voting consensus experiments. All are expressible as hGMMs. Only the first, however, exploits the flexibility of hGMMs to express dependence of actions within a neighborhood given history by (1), hence we refer to this as the *joint behavior model* (JBM). JBM incorporates the historical frequency of local configurations a_{N_i} for each i as a summarized version of history $H^t_{N_i}$ of length h. The other three forms model agent behaviors individually: for each agent we specify a probabilistic strategy $\sigma_i(H_i^t) = \Pr(a_i^t \mid H_i^t)$. The agents' actions are probabilistically dependent, but conditionally independent given this common history, yielding the joint distribution

$$\Pr(a^t \mid H^t) = \prod_i \sigma_i(H_i^t). \quad (2)$$

We refer to a dynamic multiagent model expressible by (2) as an *individual behavior hGMM* (IBMM). Conditional independence given history is a compelling assumption for autonomous agents. However, it is often infeasible to specify the entire history for conditioning, and the assumption may not hold with respect to partial history. One of these three IBMMs is designed as an independent behavior version of JBM; thus, we call it the *individual behavior model* (IBM). The remaining two models are based on proposals and observations from the original experimental analysis [10], and are labeled *proportional response model* (PRM) and *sticky proportional response model* (sPRM), respectively.

III. Learning Parameters and Graphical Structures

We first address the problem of learning the parameters of an hGMM hG given the underlying graphical structure and data in the form of a set of joint actions for m time steps, $X = (a^0, \ldots, a^m)$. For ease of exposition, let θ denote the set of all the parameters that define the hGMM's potential functions. We seek a θ maximizing the log likelihood of X using gradient ascent for updating the parameters.

We developed a structure learning algorithm that produces graphs for hGMMs within specified complexity constraints. Our learning algorithm starts with a completely disconnected graph and keep alternating between greedily adding edges that maximally improve the testing data's likelihood and tuning the model's parameters using gradient ascent optimization.

We evaluate the learned multiagent models by their ability to predict future outcomes, as represented by a test set Y. Given two models M_1 and M_2, we compute the ratio of their corresponding log-likelihood measures for the test data set Y: $R_{M_1/M_2}(Y) = \frac{L_{M_1}(Y)}{L_{M_2}(Y)}$. We are particularly interested in the ratios $R_{\mathsf{JBM/IBMM}}$, for $\mathsf{IBMM} \in \{\mathsf{IBM}, \mathsf{PRM}, \mathsf{sPRM}\}$. Note that since log-likelihood is negative, $R_{\mathsf{JBM/IBMM}} < 1$ indicates that JBM is better than IBMM at predicting Y, and vice versa if the ratio exceeds one.

IV. Empirical Study

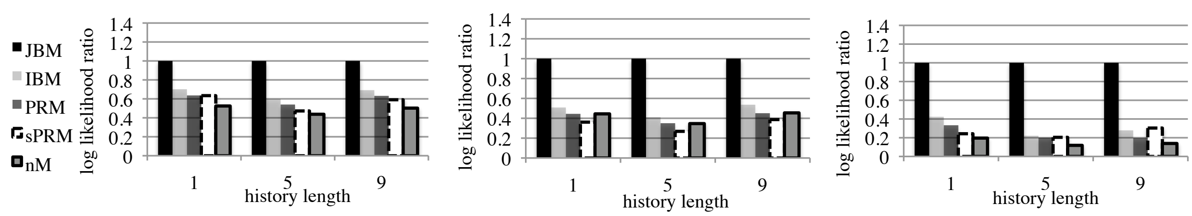

Fig. 2. JBM provides better predictions than IBM, PRM, sPRM, and nM in three experiment sets: coER_0.5 (left), coER_2 (middle), and power22 (right).

We empirically evaluate the predictive power of JBM in comparison with IBM, PRM, and sPRM, using the voting consensus experiment data from Kearns et al. [10]. We also compare JBM against a naive guessing model, nM, which initially assigns each a_i a probability proportional to $r_i(a_i)$ and linearly converges to a uniform distribution of agent actions as the game progresses. The human-subject experiments are divided into nine different sets, each associated with a network structure. We reuse Kearns et al.'s labels, coER_0.5, coER_2 and power22, for the three network sets studied in this analysis.

In our study, we learn predictive models for each network structure, pooling data across subject assignments and incentive schemes. This approach is based on the premise that network structure is the main factor governing the system's collective behavior, in line with the findings of Kearns et al. [10]. In each experiment set, we use four of the nine trials for training the predictive models for each form. The hGMM graphical structures are learned with node degree constraint 10, while the maximum degree of the individual behavior models is restricted to the greatest node degree of the original network. We then evaluate these models based on their predictions over a test set comprising the other five experimental trials. This process is repeated five times, each of which uses a different training trial set randomly chosen from the original trials.

We first examine hGMMs' predictions of players' votes in each time period conditional on available history. Figure 2 shows the log likelihood ratios between JBM and $\mathsf{M} \in \{\mathsf{IBM}, \mathsf{PRM}, \mathsf{sPRM}, \mathsf{nM}\}$, $R_{\mathsf{JBM/M}}(Y)$, computed on the test data set. We observe that JBM performs significantly better than IBM, PRM, sPRM, and nM in predicting dynamic agent behavior in the voting consensus experiments (differences significant at $p < 0.02$). As the history length h decreases, this difference in prediction performance decreases, which is likely a consequence of the fact that shorter history lengths significantly reduce the amount of information summarized in the historical frequency functions. These outcomes in general demonstrate JBM's ability to capture joint dynamic behavior, especially behavior correlations induced by limited historical information, as opposed to different IBMMs and a naive guessing model.

We also evaluated the models' capacity to predict the end state of a voting consensus experiment. The original aim of modeling in these domains was to predict this final outcome. Indeed, the convergence of PRM to consensus strongly correlates with observed experimental results. In contrast, we find that simulated runs drawn from JBM* rarely converge to a consensus, and thus this model is not directly useful for predicting end state specifically. That the model's success in capturing transient dynamics fails to translate to outcome prediction is an interesting anomaly. It will be worth further investigating whether incorporating time dependence or other additional factors can remedy the discrepant effectiveness.

REFERENCES

[1] Bakshy, E., Karrer, B., and Adamic, L. (2009). Social influence and the diffusion of user-created content. In *Tenth ACM Conference on Electronic Commerce*, pages 325–334, Stanford, California.

[2] Daskalakis, C. and Papadimitriou, C. H. (2006). Computing pure Nash equilibria in graphical games via Markov random fields. In *Seventh ACM conference on Electronic Commerce*, pages 91–99, Ann Arbor, MI.

[3] Duong, Q., Vorobeychik, Y., Singh, S., and Wellman, M. P. (2009). Learning graphical game models. In *Twenty-First International Joint Conference on Artificial Intelligence*, pages 116–121, Pasadena, CA.

[4] Duong, Q., Wellman, M. P., and Singh, S. (2008). Knowledge combination in graphical multiagent models. In *Twenty-Fourth Conference on Uncertainty in Artificial Intelligence*, pages 153–160, Helsinki.

[5] Duong, Q., Wellman, M. P., Singh, S., and Vorobeychik, Y. (2010). History-dependent graphical multiagent models. In *Ninth International Conference on Autonomous Agents and Multiagent Systems*, Toronto.

[6] Ficici, S. G., Parkes, D. C., and Pfeffer, A. (2008). Learning and solving many-player games through a cluster-based representation. In *Twenth-Fourth Conference on Uncertainty in Artificial Intelligence*, pages 187–195.

[7] Gal, Y. and Pfeffer, A. (2008). Networks of influence diagrams: A formalism for representing agents' beliefs and decision-making processes. *Journal of Artificial Intelligence Research*, 33:109–147.

[8] Heckerman, D., Geiger, D., and Chickering, D. M. (1995). Learning Bayesian networks: The combination of knowledge and statistical data. *Machine Learning*, 20:197–243.

[9] Jiang, A. X., Leyton-Brown, K., and Bhat, N. A. R. (2008). Action-graph games. Technical Report UBC CS TR-2008-13, University of British Columbia.

[10] Kearns, M., Judd, S., Tan, J., and Wortman, J. (2009). Behavioral experiments on biased voting in networks. *Proceedings of the National Academy of Sciences*, 106(5):1347–52.

[11] Kearns, M., Littman, M. L., and Singh, S. (2001). Graphical models for game theory. In *Seventeenth Conference on Uncertainty in Artificial Intelligence*, pages 253–260, Seattle.

[12] Kiss, C. and Bichler, M. (2008). Identification of influencers: Measuring influence in customer networks. *Decision Support Systems*, 46(1):233–253.

[13] Koller, D. and Milch, B. (2003). Multi-agent influence diagrams for representing and solving games. *Games and Economic Behavior*, 45:181–221.

[14] Leskovec, J., Adamic, L., and Huberman, B. (2007). The dynamics of viral marketing. *ACM Transactions on the Web*, 1(1):5–43.

[15] Leskovec, J., Singh, A., and Kleinberg, J. (2006). Patterns of influence in a recommendation network. *Advances in Knowledge Discovery and Data Mining*, pages 380–389.

[16] Salganik, M. J., Dodds, P. S., and Watts, D. J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, 311(5762):854–856.

[17] Vorobeychik, Y., Wellman, M. P., and Singh, S. (2007). Learning payoff functions in infinite games. *Machine Learning*, 67(2):145–168.

[18] Wu, F. and Huberman, B. (2007). Novelty and collective attention. *Proceedings of the National Academy of Sciences*, 104(45):17599–17601.