

Modelling Participation in Small Group Social Sequences with Markov Rewards Analysis

Gabriel Murray

Dept. of Computer Information Systems
University of the Fraser Valley
gabriel.murray@ufv.ca

Abstract

We explore a novel computational approach for analyzing member participation in small group social sequences. Using a complex state representation combining information about dialogue act types, sentiment expression, and participant roles, we explore which sequence states are associated with high levels of member participation. Using a Markov Rewards framework, we associate particular states with immediate positive and negative rewards, and employ a Value Iteration algorithm to calculate the expected value of all states. In our findings, we focus on discourse states belonging to team leaders and project managers which are either very likely or very unlikely to lead to participation from the rest of the group members.

1 Introduction

Task-oriented small groups are most effective when all group members have the opportunity to participate and be heard (Duhigg, 2016; Sunstein and Hastie, 2015). The members will have a diversity of viewpoints that can enrich the discussion and improve group problem-solving, and individual members might possess critical information that will remain hidden if the environment is not conducive to their participation (Stasser and Titus, 1985; Sunstein and Hastie, 2015; Forsyth, 2013). A group leader or project manager may be able to foster such an environment that leads to high participation levels by all team members.

In this work, we describe a novel application of Markov Rewards models to social sequence data from highly-structured small group interaction. We represent social sequence elements as complex states that include discourse information

such as dialogue act types, sentiment types, and participant roles. We associate positive and negative rewards with states, such that participation by members other than the leader has a positive reward and participation by the leader has a negative reward. We then employ a Value Iteration algorithm to calculate the expected value of each state. We particularly analyze which discourse states associated with the group leader are most likely or least likely to encourage group participation.

In Section 2, we discuss related work on applying Markov Reward models, relevant work on social sequence data and group dynamics, and various work analyzing discourse-related aspects of multi-modal interaction. In Section 3, we present our state representation, the Markov Rewards model, the Value Iteration algorithm, and the corpus of small-group interactions. We present key results and analysis in Section 4 and conclude in Section 5.

2 Related Work

In this section we survey a wide variety of research related to small group interaction, as well as Markov Rewards models.

Group Dynamics There was a great deal of research on group dynamics in the post-WWII era through to the 1970s, particularly in the fields of social psychology and organizational behaviour. For example, Steiner (1972) analyzed the effects of group factors such as group size, composition, and motivation. Forsyth (2013) summarizes much of this classic work, as well as more recent studies of group dynamics and processes. There has been a resurgence of interest on this topic in recent years, of an inter-disciplinary nature, including formal and computational models of group interaction (Pilny and Poole, 2017). Organizations such as Google (Duhigg, 2016) and Microsoft (Watts, 2016) have conducted large studies of what

makes some internal teams succeed and others fail. Similar empirical studies are described in books by Sunstein and Hastie (2015) and Karlgaard and Malone (2015).

Social Sequence Analysis Primarily falling within the field of sociology, *social sequence analysis* seeks to understand, model, and visualize social sequences, particularly temporal sequences, using a variety of tools (Cornwell, 2015; Bakeman and Quera, 2011). One of the most commonly used techniques is *optimal matching*, based on sequence alignment and editing procedures originally developed within bioinformatics. Social sequence analysis also often involves analysis of social network structure within sequences (Friedkin and Johnsen, 2011; Cornwell, 2015). In contrast to our current work, social sequence analysis often involves temporal sequences spanning days, weeks, or months, while we are examining micro-sequences spanning minutes or hours.

Multimodal Interaction In the field of *multimodal interaction*, multiple modalities of human-human interaction are investigated (Renals et al., 2012). It may be the case that the human interaction being studied takes place through multiple modalities, including face-to-face conversation, email, online chat, and notes. Or it may be the case that within a face-to-face conversation, researchers analyze many different aspects of the interaction, including speech patterns, head movements, gestures, social dynamics, and discourse structure. Multimodal interaction has also been referred to as *social signal processing* (Vinciarelli et al., 2009).

People Analytics The relatively new fields of *People Analytics* (Waber, 2013) and *Human Resource Analytics* (Edwards and Edwards, 2016) draw on some of the older fields above, in order to study aspects of human interaction and performance, particularly in the workplace. These fields examine how to improve hiring, promotion, collaboration, and group communication for businesses.

Markov Rewards Models Markov Reward models have been used to analyze many diverse phenomena, from the value of various actions in volleyball (Miskin et al., 2010) and hockey (Routley and Schulte, 2015), to a cost analysis of geriatric care (McClean et al., 1998). To our knowl-

edge, Markov Reward models have not been used for studying social sequences in small group interaction. Markov Reward models are probably best known through Markov Decision Processes (MDPs) (Bellman, 1957), which have many applications in artificial intelligence and natural language processing.

3 Small Group Social Sequence Analysis

We focus on social interactions in small group meetings. In the following two sections, we describe the state representation used for representing these social sequences, followed by the details of the Markov Rewards model and Value Iteration algorithm.

3.1 State Representation

In our representation of social sequences in meetings, each state is a 5-tuple consisting of the following information:

- the participant’s role in the group
- the dialogue act type
- the sentiment being expressed (positive, negative, both, none)
- whether the utterance involves a decision
- whether the utterance involves an action item

We are therefore analyzing sequences of complex states rather than simple one-dimensional sequences; in social sequence analysis, this is referred to as an *alphabet expansion* (Cornwell, 2015).

For our dataset (Section 3.3), the participant roles are precisely defined: Project Manager (PM), Marketing Expert (ME), User Interface Designer (UI), and Industrial Designer (ID). For our purposes here, we only care about the distinction between PM and non-PM roles. The dialogue act types are based on the AMI dialogue act annotation scheme (Renals et al., 2012), and are very briefly described in Table 1.

Example states including the following:

- $\langle PM - bck - pos - nodec - noact \rangle$
(the project manager making a positive back-channel comment, unrelated to a decision or action)
- $\langle PM - el.ass - nosent - nodec - yesact \rangle$
(the project manager eliciting feedback about an action item)
- $\langle UI - sug - nosent - yesdec - noact \rangle$
(the UI expert making a suggestion about a decision item)

ID	description
fra	fragment
bck	backchannel
stl	stall
inf	inform
el.inf	elicit inform
sug	suggest
off	offer
el.sug	elicit offer or suggestion
ass	assessment
und	comment about understanding
el.ass	elicit assessment
el.und	elicit comment about understanding
be.pos	be positive
be.neg	be negative
oth	other

Table 1: Dialogue Act Types

3.2 Markov Rewards and Value Iteration

The Markov aspect of the Markov Rewards model is that the probability of a given state depends only on the preceding state in the sequence. The state transition probabilities are estimated directly from the transition counts in the data. In addition to the complex states described in the preceding section, there are START and STOP states representing the beginning and end of a meeting, and the STOP state is absorbing, i.e. there are no transitions out of the STOP state.

The Rewards aspect of the Markov Rewards model is that certain states are associated with immediate rewards. For this study, all of the states are associated with rewards, but some of them are negative (i.e. punishments). Since our area of interest is participation by group members other than the project manager, we associate all non-PM states with a reward of 1, and PM states with a reward of -1. In other words, it is implicit that participation by people other than the project manager is desirable.

We can then differentiate between the immediate reward of a state and the estimated value of the state. For example, a particular PM state has a negative reward because it represents a discourse utterance of the project manager, but it may have a high estimated value if that state tends to lead to contributions from other members of the group. The goal then is to learn the estimated value of being in each state. We do so using a Value Iteration algorithm.

Algorithm 1 shows the Value Iteration algorithm for our Markov Rewards model. It is very similar to the Value Iteration algorithm used with Markov Decision Processes (Bellman, 1957). The inputs are an initial reward vector r containing the immediate rewards for each state, a transition matrix M , and a discount factor γ . The algorithm outputs a vector v containing the estimated values of each state. The core of the algorithm is an update equation that is applied until convergence. In the following pseudo-code, the term v_t represents a vector of estimated state values at time step t , with the initial vector v_0 consisting of just the immediate rewards.

Algorithm 1: Value Iteration for Markov Rewards Model

Input: reward vector r , transition matrix M , discount factor γ

Output: A vector v containing the estimated values of all states

```

 $v_0 = r'$ 
 $t = 1$ 
repeat
   $v_t = r' + (M * (\gamma * v_{t-1}))$ 
   $t = t + 1$ 
until convergence;
return  $v_{t-1}$ 

```

The update equation $v_t = r' + (M * (\gamma * v_{t-1}))$ essentially says that the states at step t of the algorithm have an estimated value equal to their immediate reward, plus the discounted value of the states that can be transitioned to, as calculated at the previous step $t - 1$. The discount factor γ can be set to a value between 0 and 1, and controls how much weight is given to future rewards, compared with immediate rewards. For our experiments, we set $\gamma = 0.9$. Further work will examine the impact of varying the γ value. Software for running Value Iteration and replicating these results is available at <https://github.com/gmfraser>.

3.3 Corpus

For this study, we use the AMI meeting corpus (Carletta et al., 2005), a corpus of scenario and non-scenario meetings. In the scenario subset of the corpus, each meeting consists of four participants who are role-playing as members of a company tasked with designing a remote control unit. The participants are assigned the roles mentioned previously: project manager (PM), user interface

expert (UI), marketing expert (ME), and industrial designer (ID). While the scenario given to each team is artificial and structured, the participation and interaction of the group members is not scripted. The conversation is natural and spontaneous, and the groups can make whatever decisions they see fit. For these experiments, we rely on the AMI gold-standard annotations for dialogue act type, sentiment type, decision items, and action items (Renals et al., 2012). We report results on a set of 131 scenario meetings.

4 Results

For this paper, we focus on the estimated value of states belonging to the project manager, since we are interested more generally in how team leaders can encourage participation. Table 2 shows the key results, highlighting the top 10 and bottom 10 states according to estimated value, for states belong to the PM. The table also shows the frequency of each state within the set of meetings. The top two states both represent the PM expressing positive sentiment, in the form of a backchannel and an assessment, respectively. Specifically, the second state $\langle PM - ass - pos - yesdec - noact \rangle$ involves the PM making a positive assessment about a decision item. Importantly, five states in the top 10 involve the PM explicitly trying to elicit information from the other participants. This is a less obvious finding than it may seem, for the following reason: a team leader might assume that team members will feel welcome and willing to participate in the discussion of their own volition, when in fact it may take deliberate action by the leader to elicit information from people and involve them in the discussion.

In contrast, most of the low-value states involve the PM either informing or stalling. In fact, the most frequently occurring low-value state $\langle PM - stl - nosent - nodec - noact \rangle$ represents the PM stalling, and the two lowest-value states involve the PM stalling while expressing sentiment.

While we focus here on analyzing the PM states, we briefly note that of the non-PM states, all of the top 10 states in terms of value involve a non-PM group member expressing positive or negative sentiment, and the top 5 all involve stalling. The lowest-value states involve suggestions, assessments, or back-channels. Making a suggestion or assessment regarding a decision is particularly likely to bring the PM back into the

State	Value	Freq.
PM-bck-pos-nodec-noact	2.86	42
PM-ass-pos-yesdec-noact	2.75	11
PM-el.inf-nosent-yesdec-noact	2.66	23
PM-el.ass-nosent-yesdec-noact	2.64	21
PM-bck-nosent-nodec-noact	2.60	3333
PM-el.inf-nosent-nodec-noact	2.52	1527
PM-oth-pos-nodec-noact	2.47	13
PM-und-pos-nodec-noact	2.41	22
PM-el.inf-pos-nodec-noact	2.39	11
PM-el.ass-nosent-nodec-noact	2.30	832
PM-inf-pos-nodec-noact	1.72	269
PM-inf-nosent-yesdec-noact	1.68	304
PM-off-nosent-nodec-noact	1.55	488
PM-inf-nosent-nodec-yesact	1.39	220
PM-fra-neg-nodec-noact	1.38	13
PM-stl-nosent-nodec-noact	1.35	2958
PM-inf-pos-yesdec-noact	1.31	30
PM-fra-pos-nodec-noact	1.23	27
PM-stl-neg-nodec-noact	1.22	19
PM-stl-pos-nodec-noact	0.94	74

Table 2: Top 10 and Bottom 10 States, by Estimated Value (full meeting set)

discussion. At the workshop, we will present further analysis of other interesting high- and low-value states belonging to all participants. In general, we see that all participants tend to express positive or negative sentiment while stalling, as a way of engaging in floor-holding.

5 Conclusion

We have described a novel application of Markov Rewards models to understanding small group social sequence data. By associating positive and negative rewards with particular states, and then running a Value Iteration algorithm, we can determine which states are associated with a particular outcome of interest. In this paper, our outcome of interest was participation by members of the group other than the team leader. We focused on analyzing high- and low-value states belonging to the team leader, and we briefly mentioned interesting states belonging to the other group members.

There are many other possible outcomes of interest in group interaction, and Markov Rewards models should be a useful tool for analyzing social sequences in general. To encourage such research, we are making the Value Iteration software freely available.

References

- R. Bakeman and V. Quera. 2011. *Sequential analysis and observational methods for the behavioral sciences*. Cambridge University Press.
- R. Bellman. 1957. A markovian decision process. Technical report, DTIC Document.
- J. Carletta, S. Ashby, S. Bourban, M. Flynn, M. Guillemot, T. Hain, J. Kadlec, V. Karaiskos, W. Kraaij, and M. Kronenthal et al. 2005. The ami meeting corpus: A pre-announcement. In *International Workshop on Machine Learning for Multimodal Interaction*. Springer, pages 28–39.
- B. Cornwell. 2015. *Social sequence analysis: Methods and applications*, volume 37. Cambridge University Press.
- C. Duhigg. 2016. What google learned from its quest to build the perfect team. *The New York Times Magazine*.
- M. Edwards and K. Edwards. 2016. *Predictive HR Analytics: Mastering the HR Metric*. Kogan Page.
- D. Forsyth. 2013. *Group dynamics*. Wadsworth Publishing.
- N. Friedkin and E. Johnsen. 2011. *Social influence network theory: A sociological examination of small group dynamics*, volume 33. Cambridge University Press.
- R. Karlgaard and M.S. Malone. 2015. Team genius.
- S.I. McClean, B. McAlea, and P.H. Millard. 1998. Using a markov reward model to estimate spend-down costs for a geriatric department. *Journal of the Operational Research Society* 49(10):1021–1025.
- M. Miskin, G. Fellingham, and L. Florence. 2010. Skill importance in women’s volleyball. *Journal of Quantitative Analysis in Sports* 6(2):5.
- A. Pilny and M. Poole. 2017. *Group Processes: Data-Driven Computational Approaches*. Springer.
- S. Renals, H. Bourlard, J. Carletta, and A. Popescu-Bellis. 2012. *Multimodal Signal Processing: Human Interactions in Meetings*. Cambridge University Press.
- K. Routley and O. Schulte. 2015. A markov game model for valuing player actions in ice hockey. In *Proc. of UAI*. AUAI Press, pages 782–791.
- G. Stasser and W. Titus. 1985. Pooling of unshared information in group decision making: Biased information sampling during discussion. *Journal of personality and social psychology* 48(6):1467.
- I.D. Steiner. 1972. *Group Process and Productivity*. Academic Press Inc.
- C. Sunstein and R. Hastie. 2015. *Wiser: Getting beyond groupthink to make groups smarter*. Harvard Business Press.
- A. Vinciarelli, M. Pantic, and H. Bourlard. 2009. Social signal processing: Survey of an emerging domain. *Image and vision computing* 27(12):1743–1759.
- B. Waber. 2013. *People analytics: How social sensing technology will transform business and what it tells us about the future of work*. FT Press.
- D. Watts. 2016. The organizational spectroscopy. <https://medium.com/@duncanjwatts/the-organizational-spectroscope-7f9f239a897c>. (Accessed on 05/06/2017).