Exploring Collective Behavior in Social Computation Through Relational Statistical Models

Elisabeth Baseman and David Jensen

College of Information and Computer Sciences
University of Massachusetts Amherst
Amherst, MA 01003, USA
ebaseman@gmail.com, jensen@cs.umass.edu

Abstract. Humans have a surprising ability to solve hard problems by working collaboratively. Using data from collaborative human problem-solving experiments conducted by Judd et al., we apply statistical models that are both relational and temporal to simulate the strategies that participants in the experiments may have employed to successfully perform hard computations. We compose models of individual behavior under a variety of network structures to model collective social network dynamics. Our approach simulates collective behavior accurately as determined by a metric derived from real human experiments, and provides insight into the ways that individuals chose problem solving strategies. Our findings suggest that in collaborative settings, individuals may make choices based more on their location within a social network than on intrinsic personal characteristics.

Keywords: social computation, collective behavior, graphical models, relational models

1 Introduction

Many important aspects of our society rely on the fact that people efficiently solve complicated problems when working together, but the details of these interactions are not yet well understood [1, 2]. While a wide variety of social science theories explain interactions between individuals, in dyads, triads, etc, there is no consistent quantitative theory that effectively models real-world, problem-solving strategies that lead to collective behavior [3–5]. In particular, researchers do not well understand how an individual's position within an organization affects the problem-solving strategies that the individual will adopt. Some social theories emphasize the importance of an individual's inherent traits in their choice of strategy, while other theories hold that an individual's position within an organization is of primary importance [6, 7].

In this study, we investigate the behavior of individuals in a collaborative setting using data collected at the University of Pennsylvania in which Judd et al. explore human behavior in a collaborative social network context [8]. Thirty-six participants were assigned positions in a variety of network structures. The basic network structure consisted of six cliques — highly connected graph subcomponents — containing six individuals each, as shown in Figure 1a. Sixteen network structures were generated, by varying a 're-wiring' parameter that determined the probability of an intra-clique edge being changed to connect between cliques. We focus on the consensus experiments, in which individuals were given a maximum of three minutes to arrive at a network state where all nodes were the same color. This could be any color of their collective choosing, as long as all nodes agreed. Participants could choose between nine colors, making changes to their own color at any time during the experiment. However, each participant was only able to observe the color of his or her immediate neighbors in the network, the degree of their neighbors, and a measure of how close the entire network was to completion. Performance was measured in terms of how much time was required for participants to come to a consensus. Two experiments were run using the basic 6-clique structure shown in Figure 1a, and in one of these experiments participants did not reach a consensus. The network structure which was shown to be the easiest (required the least amount of time to solve) had a high 're-wiring' parameter and is shown in Figure 1b. This easiest structure is much more connected than the basic 6-clique structure. The hardest structure that participants were consistently able to solve is shown in Figure 1c.

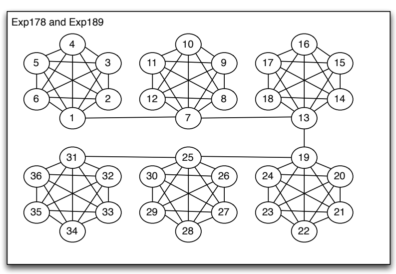

(a) The basic social network structure.

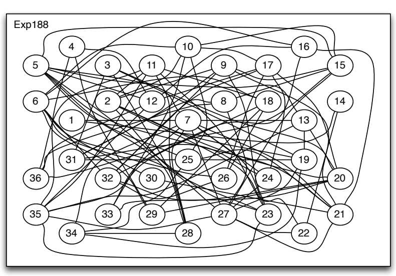

(b) The easiest social network structure.

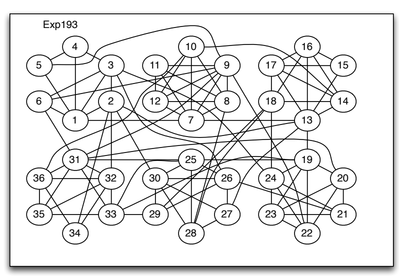

(c) The hardest social network structure humans consistently solved.

Fig. 1: Network structures used in consensus experiments.

Using data from the experiments described above, we explore several behavioral assumptions and compare our simulation results to the actual experimental results. We explore a variety of simple behavioral models, based on whether behavior is guided by intrinsic characteristics or the individual's location within the social network. Previous work has investigated the usefulness of these intrinsic and extrinsic characteristics. Campbell et al. investigate these characteristics using a relational approach to classify Twitter users, and find that combining both types of characteristics provides more powerful classification [9].

In this study, we investigate the use of graphical models of temporal and relational data to model individual and system-wide behavior. This means that we do not assume homogeneity among individuals, allowing us to proceed with very few modeling constraints. Then, we learn the parameters of each individual's behavior and compose those individual models into models of system behavior.

We present several ways of modeling individual and collective behavior, along with a predictive task for evaluating our results. We further investigate whether intrinsic personal traits or network location have a greater impact on behavior.

2 Approach

Using the logs from Judd et al.'s consensus experiments, we performed some preliminary analyses. These logs include network structure and placement of individuals within the network, as well as a full time-stamped record of each individual color change. We also obtained records of exit interviews conducted with the human participants by Judd et al.

We initially thought that individuals would tend to choose the color that caused the fewest conflicts with their neighbors. However, an analysis of the experimental data shows that participants chose the color that minimized conflict only about two-thirds of the time. Therefore, people made their decisions in some unanticipated fashion, so we chose to learn a model that predicts observed experimental data, rather than attempting to explicitly parameterize known behavioral strategies. Drawing on well-established machine learning techniques, we leveraged probabilistic graphical models for modeling and simulating the social computation consensus experiments.

Graphical models are a powerful probabilistic technique for representing the joint probability distribution of a system, such as a social network. In our case, this system consists of individuals assigned to nodes in a specific graph structure as well as their color choices over time. Graphical models represent a system's joint probability distribution as a graph, where nodes represent the system's random variables (which can be either observed or latent) and edges represent correlations between variables [11]. Each node stores the conditional probability distribution (CPD) of the associated random variable, which is learned using training data — in our case, data from the actual human experiments. The causal models we employ are required to be acyclic, and follow the causal Markov condition, that each variable is statistically independent of its non-descendants when it is conditioned on the values of its parent variables. This social computation data has a temporal aspect, so we must extend our static graphical models to incorporate a time component. When modeling a system over time with a consistent structure, we simply use the same model at each time step and require that the system also follow the Markov assumption that the values at each time step are dependent only on values in the same time step or one previous time step [12]. In addition, temporal models allow some or all edges to move from a variable at time t to a variable at time $t+1$.

Because our exploratory analyses indicate heterogeneous behavior among participants, we employ statistical relational models, which are designed specifically to model heterogeneous systems of different kinds of interconnected entities [13]. These relational models rely on entity-relationship (ER) diagrams, which include entity types represented as rectangular nodes and relationship types that connect those entities represented as edges with diamonds at their midpoint [10]. Variables of entities are represented as ellipses within entities. Edges may display crow's feet, indicating many-to-one or many-to-many relationships. In our study, we have only one entity type — a node in the social network — with person, context, strategy, and color attributes. We also only have one type of relationship — neighbor — which indicates that two nodes are adjacent in their social network.

We took a relational approach, in which we designed a model of the decision making process for individuals, then built a dynamic Bayesian network — a Bayesian network which allows temporal dependencies between random variables — by combining the models of individuals. We then used this dynamic Bayesian network to learn parameters of individual behavior and simulate collective behavior. Because individuals can only see the color of their immediate neighbors, the individual decision models are fairly simple. In one approach, decisions are made based only on neighbor colors, one's own color, and strategy, as shown in Figure 2. In an even simpler model, individuals may not take into account their own color, removing one of the self-loop edges on the color variable in the template. In the individual models, each node has a person variable – representing which human participant is located at that node, and a context variable. The context variable encodes attributes of the node's location in the network, such as degree. Dependency edges that pass outside of the node indicate dependencies between neighbor nodes, and dependencies labeled with t indicate a temporal dependency.

Then, we use the actual social network structure employed in the experiments in combination with our model template to generate what is known as a "ground graph". This graph is created by using a copy of the template model for each individual in the social network, which results in a Bayesian network for one time step. Then, the entire ground graph, a large dynamic Bayesian network that will simulate the entire network, is generated by rolling out the graph over multiple time steps, as shown in Figure 3.

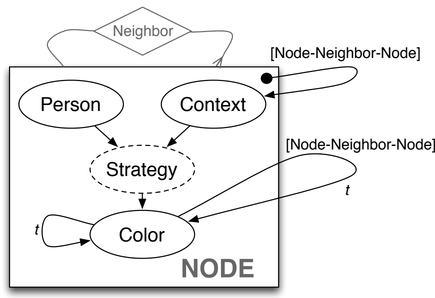

Fig. 2: Example of a model that includes the node's previous color.

3 Probabilistic Models

Based on our exploratory observations, we hypothesized 6 different decision-making strategies to explain participants' behaviors. For each strategy, we considered 2 model variations — one where an individual takes into account his own color, and one where he only considers his neighbors' colors — for a total of 12 probabilistic models. These models are organized into two categories: behavioral and contextual. In behavioral models, we assume that the participants' strategies are based on some intrinsic personal characteristics, rather than on their location within the social network. In contrast, contextual models assume that participants' strategies are based primarily on some attribute of their location within the network. For completeness, we include two homogeneous models where all participants employ the same strategy, which we refer to as "pooled."

3.1 Behavioral Models

We call our two general behavioral models "person" and "toggle." In the "person" model, we assume that each individual uses his own unique strategy based on the color proportions of his neighbors, regardless of location in the network or exactly how many neighbors he has. In the "toggle" model, we observe that certain individuals have a tendency to change color at a higher frequency than the rest of the population. We calculate the average time between color changes for each human participant, and designate those who change color at an average frequency higher than one standard deviation above the average as fast changers, those within one standard deviation of the average as normal changers, and those lower than one standard deviation away from the average as slow changers. We then simulate choices based on whether the human participant assigned to the node was a fast, normal, or slow changer. We base this "toggle" model on exit interviews from participants who indicated that they switched colors rapidly in an attempt to communicate with their neighbors.

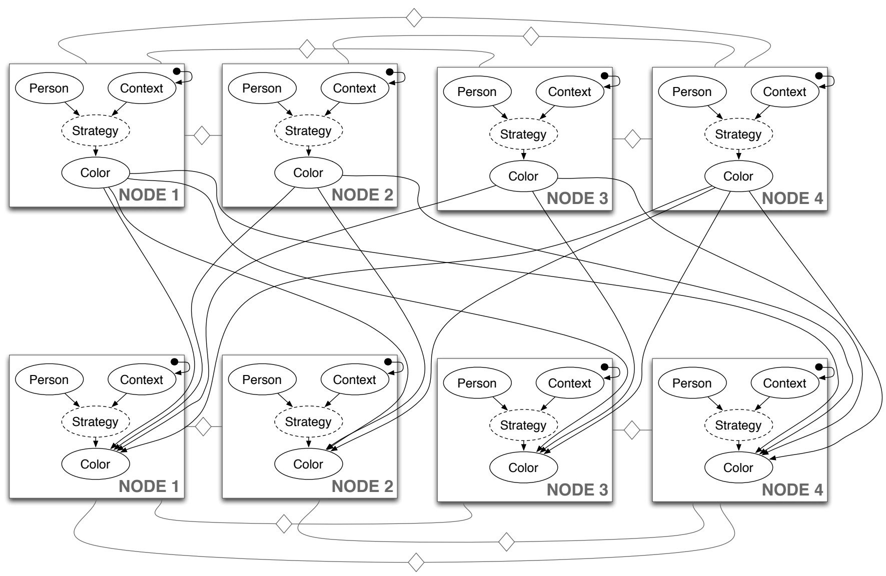

Fig. 3: Example of a ground graph, based on the model template in Figure 2 and a fictitious social network structure with four individuals

.

Algorithm 1 Node State Calculation

procedure GETNODESTATE(*neighbors*)
    distribution $\leftarrow$ []
    for n in *neighbors* **do**
        increment *distribution*[*n.color*]
    for $i = 0$, $i < distribution.length$, $++i$ **do**
        $distribution[i] \leftarrow \frac{distribution[i]}{neighbors.length}$
    return *distribution.sort*

3.2 Contextual Models

We refer to our three general contextual models as "degree," "bottleneck," and "clique." In the "degree" design, participants with the same degree in the network use the same strategy. The "degree" model assumes that humans who are trying to synthesize similar amounts of information behave similarly; participants with 7 neighbors are all trying to digest more information than participants with only 2 neighbors. In the "bottleneck" design, we assume that individuals are highly influenced by whether or not they are located at a node likely to be a "bottleneck" — a location that is likely to be a point of conflict in the network (nodes 1, 7, 13, 19, 25, and 31, which interface between cliques in Figure 1a). We assume that participants located at bottleneck nodes use the same strategy, while participants not located at bottlenecks employ a different strategy. Finally, in the "clique" model, we assume that participants who have been assigned to

Algorithm 2 Simulating an Experiment

```
procedure SimulateExperiment(network)
    ColorChangeCount ← 0
    while ColorChangeCount < 600 do
        if network.isSolved then return
        for node in network do
            node.color ← predictColorChoice(node)
```

Algorithm 3 Predicting a Node's Color Choice

```
procedure PredictColorChoice(node)
    if node.neighbors.agree then return node.neighbors.color
    node.state ← getNodeState(node.neighbors)
return color ∼ Categorical(color | node.equivalenceClass, node.state)
```

the same 6-node clique use the same strategy (for a total of 6 different strategies in the simulation), hypothesizing that behavior may be influenced by who is in your neighborhood of the network, rather than your degree.

3.3 Homogeneous Model

We refer to our homogeneous model as "pooled." In this design, we investigate the hypothesis that all individuals use the same decision-making strategy, based on the current color distribution of their neighbors, regardless of their network position. There is only one strategy employed, shared across all participants.

4 Experimental Setup

To simulate the consensus experiments, we first need to define the idea of node state. Because we do not care about the specific color chosen, and only that the network reaches some uniform color, we need the notion of state to be color-agnostic. We achieve this by calculating the normalized frequency of a node's neighbor's colors and defining the state to be this probability distribution, independent of the specific elements of the distribution, as shown in Algorithm 1. That is, a node with three neighbors with colors blue, blue and green, would have a state represented as 66%:33%. A node with six neighbors, 4 red and 2 blue, would also have state 66%:33%. In this way, we can encode the state of a node in a color-agnostic way so that we can match a node's current state with times it saw that same state in our training examples.

We simulate the experiments as shown in Algorithm 2, where a network is said to be solved when all nodes agree on color choice. Our models define equivalence classes of nodes — for the "degree" model, a node's equivalence class includes all other nodes with the same degree, etc. We simulate each node's color choice by computing the color transition probabilities using the maximum likelihood estimate based on how nodes in the same equivalence class behaved when in the same state as the current node, as shown in Algorithm 3.

For each model, we examine 10 simulations of each network structure using the same person-node assignment as in the human experiments, and record the number of color changes executed. When simulating collective behavior given a particular network structure, we use data from human behavior in all other network structures for training. To simulate users making decisions simultaneously in real time, we take a randomized turn-based approach. We observe that in the human experiments, no two users ever make color changes at exactly the same moment, so we believe we are justified in simulating the activity using a turn-based approach.

4.1 Evaluation Criteria

We are interested in accurately modeling the collective behavior of a social network. We chose "number of color changes performed before reaching a solution" as a good quantitative measure of collective behavior and how difficult the network is for humans, or a particular simulation model, to solve. For each of our models, we ran 10 simulations on each network structure. For each model-network structure combination, we then calculated the average number of color changes taken to reach a consensus in simulation. Using these averages, we ranked the network structures by how difficult they were for each model to solve. Finally, we compared the network structure difficulty rankings, using Spearman rank-order correlation, for each model to the actual difficulty as measured from the human subject experiments. If the collective behavior of our simulations closely models the collective behavior observed in the real experiments, we expect the resulting network structure ranking based on difficulty to be similar. The model with simulations that provided the network difficulty ranking most similar to the network difficulty ranking observed from human experiments will obtain the highest Spearman rank-order correlation coefficient.

5 Results

We find significant differences in the predictive performance of our contextual and behavioral models, with only two models out-performing the homogeneous baseline ("pooled") model, which we consider as a naïve baseline. Table 1 shows each model ordered by the obtained Spearman rank-order correlation coefficients. The model that best predicted network difficulty was the "degree" model where individuals do consider their own color in the decision making process. We note that only the "degree" and "bottleneck" models that include the individual's own color perform better than the "pooled" model. Both of these models are context-based, implying that we only do better than the baseline when we assume that people behave more based on their degree within their social network rather than on intrinsic behavioral characteristics. Other contexual models, such as "bottleneck," do not out-perform the baseline either, suggesting that individuals

Model Grouping	Self	Spearman Rank-Order Correlation Coefficient
degree	with	0.753
bottleneck	with	0.718
pooled	without	0.691
pooled	with	0.688
toggle	without	0.683
toggle	with	0.682
bottleneck	without	0.641
clique	with	0.479
degree	without	0.324
clique	without	0.267
person	without	0.013
person	with	-0.057

Table 1: Spearman rank-order results for network difficulty.

are sensitive more to their specific degree, rather than how much conflict they encounter among their neighbors.

The simulated and real behavior for the best-performing model in terms of ranking network difficulty on the easiest and hardest network structures are shown in Figures 4 and 5, respectively. In the figures, the blue line shows the average behavior in the simulations while the gray series show the behavior in each individual simulation, and the green series shows observations from the human behavior experiment. Note that the average line seems to be a good representation of the general behavior, with no obvious outliers. In addition, the number of color changes required to reach a solution appears to be in the correct order of magnitude for the simulated experiments.

The worst-performing model for network difficulty was the "person" model where individuals did include their own color, but each behaved with their own unique strategy. The simulated behavior for the worst-performing model is compared to human behavior for the easiest network in Figure 6. Even in the easiest network, none of the simulated experiments reach a consensus.

We find that in terms of network difficulty, the model where individuals base their decisions on their degree in the network and include their own color in their reasoning performs best. This implies that participants placed at nodes in the network with the same degree tended to behave similarly, regardless of who they were. It has long been well-known that humans will change their behavior when they are arbitrarily placed somewhere in a hierarchical structure, but these consensus experiments did not include an explicit hierarchy among individuals [14, 15]. Instead, individuals seem to adjust their behavior simply based on the number of social connections available to them, even when faced with a straightfoward computational task.

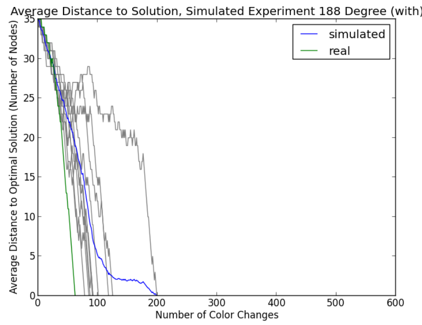

Fig. 4: Participants behave based on degree, including self in color distribution (easiest network).

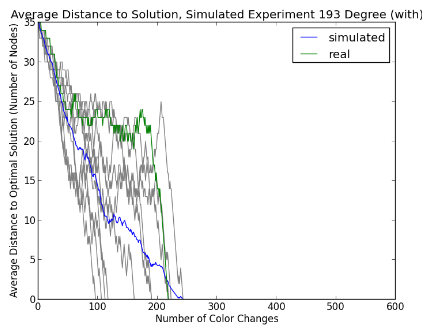

Fig. 5: Participants behave based on degree, including self in color distribution (hardest network).

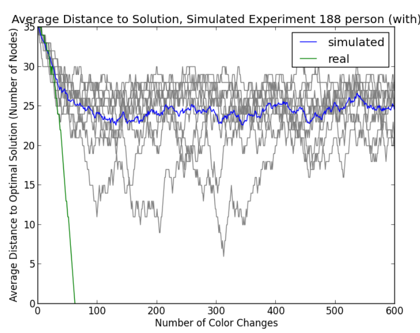

Fig. 6: Each participant behaves differently, including self in color distribution (easiest network).

6 Conclusion

We applied statistical models that were both relational and temporal to analyze collective behavior in color consensus experiments. We developed 12 probabilistic models to describe different hypotheses regarding the decision-making processes of individual participants and by learning parameters of these models were thereby able to evaluate the extent to which individuals employed a variety of strategies in coming to a consensus. Our results show that individuals' choices are more influenced by their position than by their intrinsic characteristics. While the most predictive approach modeled behavior based on node degree, it was somewhat surprising to find that exclusively considering an individual's own past behavior was the least effective approach.

Across all of our results, we considered the average case of our predictions, however it's clear from our work that there are many probable outcomes for a particular network structure. Because Judd et al. only conducted one experiment for each network structure, we were not able to investigate whether this variance is an artifact of our modeling methodology or a description of variations that would indeed occur if the experiment was repeated many times. Variance is a crucial aspect of any predictive task, and we advocate for further social computation experiments to validate whether our models are predictive of the range of possible outcomes.

Eventually, the findings from our models could contribute to social science, advancing theoretical understanding of collaborative problem solving. Once we have advanced our modeling capabilities to support predictive models of variance, these models could contribute to the design of future social computation experiments, effectively pruning the hypothesis space and guiding resources toward the least understood network structures. Since the monetary cost of complex social science experiments is even more prohibitive than social computation research, investing in effective modeling tools will increase the efficacy and efficiency of real-world social network experiments involving human subjects.

Acknowledgements

The authors wish to thank Hanna Wallach for substantial and extremely useful comments on an earlier version of this paper. The authors also thank Michael Kearns and Stephen Judd, at the Department of Computer and Information Science at the University of Pennsylvania, for assistance in the design of this work and obtaining data.

Funding was provided by MIT Lincoln Laboratory and Defense Advanced Research Projects Agency (DARPA) under Contract Number 7000262845. The content of the information in this document does not necessarily reflect the position or the policy of MIT, MIT Lincoln Laboratory, or the Government, and no official endorsement should be inferred. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

Bibliography

[1] Bazerman, M: Judgement in Managerial Decision Making. John Wiley Sons, Hoboken, NJ (1998)

[2] Newman, M. E.: The Structure of Scientific Collaboration Networks. Proceedings of the National Academy of Sciences 98:2, 404-409 (2001)

[3] Kenny, D. A., La Voie, L.: The Social Relations Model. Advances in Experimental Social Psychology 18, 142-182 (1984)

[4] Sadoka, J.M.: The Checkerboard Model of Social Interaction. The Journal of Mathematical Sociology 1:1, 119-132 (1971)

[5] Wasserman, S.: Social Network Analysis: Methods and Applications. Cambridge University Press, Cambridge, United Kingdom (1994).

[6] Janis, I.L.: Victims of Groupthink: A Psychological Study of Foreign-Policy Decisions and Fiascoes. (1972)

[7] Bazerman, M. H., Schoorman, F. D.: A Limited Rationality Model of Interlocking Directorates. Academy of Management Review 8:2, 206-217 (1983).

[8] Judd, S., Kearns, M., Vorobeychik, Y.: Behavioral Dynamics and Influenced in Networked Coloring and Consensus. Proceedings of the National Academy of Sciences 107:34, 14978-14982 (2010)

[9] Campbell, W., Baseman, E., Greenfield, K.: Content + Context Networks for User Classification in Twitter. Frontiers of Network Analysis: Methods, Models and Applications Workshop at Neural Information Processing Systems (2013)

[10] Neville, J., Jensen, D.: Relational Dependency Networks. Journal of Machine Learning Research 8, 653-692 (2007)

[11] Koller, D., Friedman, N.: Probabilistic Graphic Models: Principles and Techniques. MIT Press, Cambridge, MA (2009)

[12] Murphy, K.: Dynamic Bayesian Networks. (2002)

[13] Getoor, L, Taskar, B.: Introduction to Statistical Relational Learning. MIT Press, Cambridge, MA: (2007)

[14] Haney, C., Banks, C., Zimbardo, P.: A Study of Prisoners and Guards in a Simulated Prison. Naval Research Reviews 30, 4-17 (1973).

[15] Peters, W.: A Class Divided, Then and Now. Yale University Press, New Haven, CT (1987).