A cognitive computational model of collective search with social information

Sabina J. Sloman, Robert L. Goldstone and Cleotilde Gonzalez

For many if not most of the decisions we make in day-to-day life, such as picking a restaurant for dinner when traveling in a new city or choosing a university to attend or a job offer to accept, our knowledge about the valuations of possible outcomes is incomplete. However, many current decision theories specify how to make decisions when the possible outcomes and their valuations are fully specified.

When information is incomplete, the quality of that information and how we use it to inform our decisions is especially important. *Social information* is a critical source of information that we use to assess the costs and benefits of our actions: Usually, we extrapolate from not only our own experiences, but also from the experiences of others, e.g., anecdotes told by friends or reviews read online. The structure of the social environments in which we receive information from others—our social *networks*—can vary in important ways. While a friend's anecdote may only be told to a few in a tight-knit social network, online reviews are available to everyone who cares to look. It turns out the structure of our social networks can have dramatic effects on our ability to identify and make the best decision possible.

In a 2008 experiment, Mason, Jones, and Goldstone explored the effect of participants' social network structure on their abilities to identify the best decision given incomplete information. Mason, Jones, and Goldstone (2008) designed a spatial search task where participants selected integers on a number line and received points as a function of the number they guessed. Participants had access to not only their own outcome information, but also to the outcome information of their network neighbors. Mason et al. (2008) found that when the payoff function was smooth with an obvious global maximum—when extrapolating from a relatively small amount of information was usually sufficient to identify the best decision—members of more interconnected networks more frequently guessed within the global maximum. By contrast, for jagged problems with more than one local maximum—complex problems for which identifying the best decision usually required exploration—members of more dispersed networks more frequently guessed within the global maximum.[1]

Our work builds on this result. We synthesize work from various areas of cognitive science into a computational cognitive model of search in a social context: the Social Interpolation Model (SIM). We then explore the implications of our model by running simulations of interacting agents whose behavior is determined by the SIM. By embedding these agents in the same task structure as the one designed by Mason et al. (2008), we explore how these dynamics are affected by the structure of the agents' social networks.

All of our simulations were run with groups of 15 agents, arranged in the same four network structures as the participants in Mason et al. (2008)'s studies. On each of 15 consecutive rounds, these agents "guessed" integers between 0 and 100. On a given round, an agent's guess was informed by the outcome information generated by their own previous guesses and the guesses of their network neighbors.[2]

The SIM posits that agents rely on similarity-based generalization (Shepard, 1987) to integrate the outcome information they've already seen in order to infer the number of points they are likely to receive from decisions whose payoffs are as-yet unknown. An agent's probability of selecting a particular number is proportional to the number of points they think that decision is most likely to confer. The SIM has three free parameters, or avenues for individual difference: 1) the breadth of the agent's generalization gradient, 2) the quality of the agent's uninformed prior about unseen options, and 3) the degree to which the agent weights their own experiences more heavily than the experiences of others.

Like Mason et al. (2008), we find that network structure matters: Interconnected networks perform well when the best decision is easy to identify, but comparatively worse when the payoff function is more complex. We find that the most effective parameter settings also depend on the complexity of the problem: Agents who generalize broadly or attach a high value to unobserved options do well on more complex landscapes—and

[1] The most interconnected network configuration was a fully-connected network, in which every agent had access to the outcome information for every other agent. The most dispersed network configuration was a regular lattice, in which group members were arranged in a circle and had access to the outcome information of only their two nearest neighbors (with some members also connected to others two steps away). Two other network structures were characterized by average path lengths that fall between the extremes defined by the fully-connected and regular lattice networks. See Mason et al. (2008) for complete details.

[2] While the full version of the SIM allows for retention and aggregation of outcome information from multiple rounds, agents in the simulations reported here retain and aggregate information from only the most recent round.

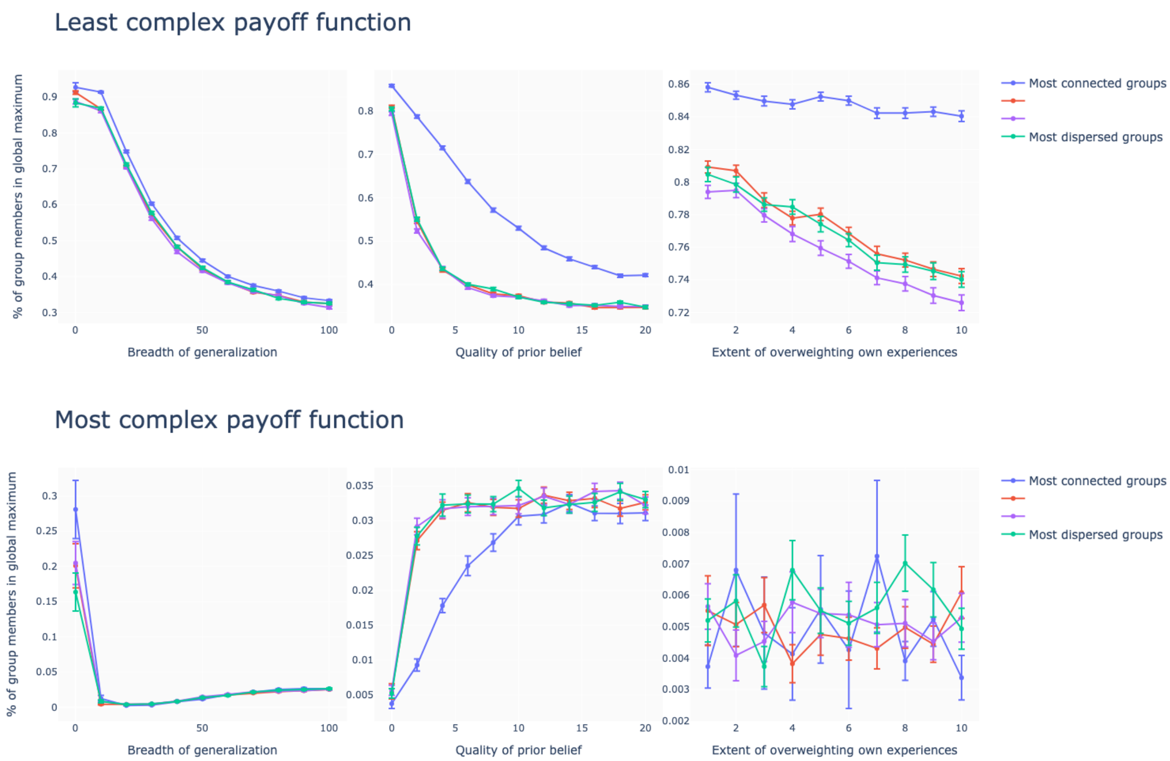

Figure 1: Percentage of simulated agents' guesses that were within one unit of variance of the payoff functions' global maximum.

worse on simpler landscapes (see Figure 1).[3] Taken together, our results show that the effect of reducing the flow of social information by putting agents in a more dispersed social network can have similar effects to inducing the same agents to generalize more broadly or to be more optimistic about unseen options.

In summary, our main contributions are 1) the development of a computational cognitive model of search in a social context, 2) an exploration of the effects of the values of the SIM's free parameters, 3) the deployment of our theory in an agent-based model, and 4) an exploration of the effect of different social and reward environments on the SIM's dynamics. Our work has important practical and theoretical implications. Practically, our agent-based framework can allow exploration of the effects of different interventions in different contexts. Theoretically, we synthesize various areas of cognitive science into a single model that can make predictions about individual- and group-level behavior in decision-making environments characterized by incomplete information and the availability of social information.

References

Mason, W. A., Jones, A., & Goldstone, R. L. (2008). Propagation of innovations in networked groups. *Journal of Experimental Psychology: General*, *137*(3), 422–433. doi: 10.1037/a0012798

Shepard, R. (1987, September). Toward a universal law of generalization for psychological science. *Science*, *237*(4820), 1317–1323. doi: 10.1126/science.3629243

[3]Figure 1 shows that the effect of generalization on the complex payoff function is actually U-shaped: Agents in the most connected groups who generalize extremely narrowly perform the best. We speculate that this exception to the general trend reveals a tension between two components of success in this environment: One the one hand, exploration helps agents escape local maxima; on the other, narrow generalization more accurately reflects the ruggedness of the payoff function.