Trust, Gossip, and Collaboration in Agents

Michael Prietula
University of Florida

Kathleen Carley
Carnegie Mellon University

Abstract

Much of organization science is based on the premise that members of the organization work cooperatively to get things done. In such a situation trust is imperative. Work on power and politics however, suggests that it may not be appropriate to assume that trust exists. Trust is typically facilitated by face-to-face interaction. Thus, as we move to distributed groups, virtual organizations, and web-based teams the level of basic trust in team members is likely to decrease. We ask, in such a distributed setting – how much trust do you need?

Early work on communication in a search task suggested that truth telling affected organizational behavior, but that the rate of forgiveness did not (Prietula & Carley, 1999). This work also suggested that in a volatile environment, it was difficult to distinguish truth-tellers from liars except through reputation (Prietula & Carley, 1998; Carley & Prietula, 1998), A web based environment, where the rate of "movement" is high, and the lifetime of sites low, is such a volatile environment. In this paper we examine the role of trust and reputation in a volatile distributed team environment.

A simple search task is simulated, where computational agents seek to find objects (e.g., items, websites, patterns) in a search space. The concept of "search" is loosely defined and can be interpreted with either a more or less physical analogy (e.g., seeking items over a geographical area, or
seeking relevant web sites on the Internet). For the agents in this task, the objects they seek are simply integers scattered around a block-metric set of locations. They have a list of integers they are to locate, so they go about their task of searching the block-metric space for those integers.

The first set of simulations examines baseline behaviors of agents in the task without communication. These simulations hold the task constant, and vary the number of agents engaged on the task, as well as two key agent properties, where each agent engages in random searches for their objects,
but does not communicate with other agents. In addition, beyond baselines, these simulations address versions of a question proposed by Cohen (1990) who investigated single versus multi-actor agents in a learning task – when and how do multi-actors do better than single actors on a learning task?

The second set of simulations elaborate on this baseline data by adding agent communication in the form of advice on object locations. The third set of simulations broadens the agent architecture and adds simple trust models that vary in their tolerance for bad advice. Bad advice is generated by randomly disrupting the environment, so that advice provided by agents is invalidated. The results

of the interactions of the factors are presented as a set of observations concerning impacts the factors on three dependent variables measuring organizational effectiveness (percent of the task that was completed) and efficiency (time to complete the task, total effort required by the group).

The fourth set of simulations keep the environment stable, but add agents that are purposefully deceptive, and provide false advice on locations. Thus, this examined how the three observations previously defined would hold over a different source of advice uncertainty.

The final set of simulations adds the capability of generating and spreading gossip about particular agents as credible sources of advice, and examined the impact of gossip on the prior observations and coallition formation. As will be seen, gossip involve several specific decisions for agents that must
be addressed by additional agent component models.

References

Michael Prietula & Kathleen M. Carley, 1999, "Exploring the Effects of Agent Trust and Benevolence in a Simulated Organizational Task," *Applied Artificial Intelligence*, 13(3): 321-338.

Michael J. Prietula & Kathleen M. Carley, 1998, "A Computational Model of Trust and Rumor. " In Proceedings of the 1998 AAAI Fall Symposium Series - Emotional and Intelligent: The Tangled Knot of Cognition, Orlando, FL. October 23-25, 1998.

Kathleen M. Carley & Michael J. Prietula, 1998, "Webbots, Trust, and Organizational Science." Pp. 3-22 in Michael J. Prietula, Kathleen M. Carley & Les Gasser (Eds.), *Simulating Organizations: Computational Models of Institutions and Groups*, Menlo Park, CA: AAAI Press / The MIT Press.