

Organizations of Self-Interested Agents

Foster McGeary

Computer and Information Sciences
University of Delaware

Existing agent-based systems apply the productive power of Multi-Agent Systems (MAS) and Distributed Artificial Intelligence (DAI) techniques. However, little work has demonstrated agents autonomous enough to form concurrently operating non-trivial, non-predetermined organizations where agents form such organizations specifically to improve their individual welfare through improving their collective productivity.

In my work, organizations are defined as sets of autonomous agents that reach, and generally abide by, agreements to exchange, over periods of time, computational services for value, and where such agreements are voluntary by each agent. The dissertation¹ identifies and explores techniques by which autonomous agents may form organizations that persist over time, in dynamic environments, and interact with agents outside the organization.

Curiosity about the computational basis of persistent organizations is a driving force behind this work. Extensive writings in Economics, Sociology, and Political Science purport to explain why certain human organizations operate the way they do. From a computational view, and in particular from an Artificial Intelligence perspective, there is no well-developed modeling approach that supports building models with the large number of agents necessary to model human social behavior using the terms of these other areas of study when they appeal to methodological individualism on a large scale. If we project recent rates of advancement in computational power, we find we can project the power to support detailed and explicit models of the behavior of each of the individual agents within an organization or society. My dissertation addresses the methods needed to harness that growing power to the creation of large-scale, detailed, models of interrelationships among large numbers of individual decision units, which decision units are “agents” in Distributed Artificial Intelligence work.

My thesis is that the formation and operation of organizations is rational and computable, and can be modeled using Artificial Intelligence tools. I extend Multi-Agent Systems technologies by introducing procedures by which sovereign computational agents may form organizations. A sovereign

agent is a computational agent with its own motives — an agent that cannot be forced to abide by any convention to which it is not a willing party. Organizations of sovereign agents are expected to occur in a variety of forms and, to the extent these organizational forms are viable, the terms of traditional organizational theory will be used to catalog them. Organizations formed with the intention to sell products at a profit (and thus to enrich the agents owning the organization) are taken to be firms, a principal interest of the research.

I develop the agent environment sequentially. First, domain-level (Virtual Food Court) technologies are created to make products when properly instantiated with the skills and resources (raw materials or other products) they require. Second, agents are endowed with skills that, when applied within technologies, have value derived from the products of the technology. Third, agents negotiate compensation for the application of their skills. Fourth, agents are provided with computable preferences for products and the desire to acquire products to maximize their private welfare. Finally, the agents are permitted to act in accordance with these components and the data they collect themselves.

Agent welfare (products consumed) is expected to be enhanced by certain behaviors. Sovereign agents require basic protocols simply to engage in barter transactions. More elaborate protocols facilitate contracting for services through time. Still other methods allow each agent to identify other agents, to determine what information it is helpful to acquire from or disseminate to other agents, and thus permit agents to engage in useful “social” behavior.

Organizations are to be contingent entities, but with their form, size, and activities traceable to (and computable from) particular sets of technologies, agents, and agent preferences for products. Organizations will also be mildly contingent upon non-deterministic elements in their environments, such that a given set of initial conditions can result in slightly different organizations, depending on uncontrolled activities on the computers on which the agents operate.

The computational procedures under which sovereign agents exchange data, negotiate the application of and compensation for their skills, and make reciprocal commitments are the structures supporting the organization of firms among sovereign computational agents. This form of computational organization is original with my dissertation.

Copyright © 2002, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

¹Based in part on work supported by National Science Foundation Grant Nos. IIS-9733004 and IIS-9812764.