

# The Impact of Locality and Authority on Emergent Conventions: Initial Observations

James E. Kittock\*

Robotics Laboratory  
Computer Science Department  
Stanford University  
Stanford, CA 94305  
jek@cs.stanford.edu

## Abstract

In the design of systems of multiple agents, we must deal with the potential for conflict that is inherent in the interactions among agents; to ensure efficient operation, these interactions must be coordinated. We extend, in two related ways, an existing framework that allows behavioral conventions to emerge in agent societies. We first consider localizing agents, thus limiting their interactions. We then consider giving some agents authority over others by implementing asymmetric interactions. Our primary interest is to explore how locality and authority affect the emergence of conventions. Through computer simulations of agent societies of various configurations, we begin to develop an intuition about what features of a society promote or inhibit the spontaneous generation of coordinating conventions.

## Introduction

Imagine a society of multiple agents going about their business: perhaps it is a team of construction robots assembling a house. Perhaps it is a group of delivery robots responsible for carrying books, copies, or medical supplies throughout a building. Or perhaps it is a society of software agents, working to collect data from diverse sources in “information space.” Whatever the nature and environment of these agents, they will find it necessary to interact with one another. There is an inherent potential for conflict in such interactions; for example, two robots might attempt to move through a doorway at the same time, or two software agents might try to modify the same file. As designers, we have achieved *coordination* when agents’ actions are chosen specifically to prevent such conflicts. *Conventions* are a straightforward means of implementing coordination in a multi-agent system. When several conflicting strategies are available to agents for approaching a particular task, a convention specifies a common choice of action for all agents.

\*This research was supported in part by the Air Force Office of Scientific Research under grant number F49620-92-J-0547-P00001 and by the National Science Foundation under grant number IRI-9220645.

In general, designing all necessary conventions into a system or developing a centralized control mechanism to legislate new conventions is a difficult and perhaps intractable task (Shoham & Tennenholtz 1992b). It has been shown that it is possible for an agent society to reach a convention without any centralized control if agents interact and learn from their experiences (Shoham & Tennenholtz 1992a). Conventions thus achieved have been called “emergent conventions”, and the process for reaching them has been dubbed “co-learning.” (Shoham & Tennenholtz 1993)

In previous work on the emergence of conventions through co-learning, it was assumed that each agent in a society is equally likely to interact with any other agent (Shoham & Tennenholtz 1992a). This seems an unreasonable assumption in the general case, and we consider ways to extend the framework by allowing for non-uniform interaction probabilities. Conceptually, we can divide limitations on interactions into two categories: those due to inherent separation (geographic distance, limited communication, etc.) and those due to organizational separation (division of labor, segregation of classes, etc.). These notions are two sides of one coin; they are different forms of *locality* within a multi-agent society.

Previous work also assumed that agents have equal influence on one another’s behavior. However, in practice this is not generally true; multi-agent systems often have variations in *authority*. We model differences in authority by implementing asymmetrical interactions in which the less influential agent always receives feedback as a result of its actions, while the more influential agent receives feedback with some probability. As the probability that an agent receives “upward feedback” from its subordinates decreases, the agent’s authority increases. This is intended to model a situation in which the agent can act with impunity, choosing strategies based only upon their perceived effects on other agents (this can also model the case in which an agent is simply stubborn and deliberately chooses to ignore feedback). We do not claim that this is an exhaustive treatment of the notion of authority, but this asymmetry of interaction is one aspect of authority

that is strongly related to the topological organization of the agent society.

Our primary aim in this paper is to explore how various forms of locality and authority affect the emergence of conventions in a multi-agent society. We note that our goal is not to model human society; rather, we seek to gain a preliminary understanding of what global properties we might expect in a society of artificial agents that adapt to one another's behavior.

### The Basic Framework

Our explorations of multi-agent societies take place within a formal model that allows us to capture the essential features of agent interactions without becoming lost in unnecessary detail. In particular, we distill the environment to focus on the interactions between agents: each agent's environment consists solely of other agents in the system. This allows us to be precise about the effects of agents on one another.

In this model, agents must choose from a finite repertoire of *strategies* for carrying out an action. When agents interact, they receive feedback based on their current strategies; this simulates the utility various situations would have to an agent (for example, a robot might get negative feedback for colliding with an obstacle and positive feedback for completing a task). An agent may update its strategy at any time, but must do so based only upon the history of its feedback (its "memory"). As designers in pursuit of emergent coordination in an agent society, our ultimate goal is for all agents to adopt the same strategy. Thus, we must find an appropriate strategy update rule that causes a convention to arise from mixed strategies.

We have limited our preliminary investigation to pairwise interactions, and we can write the possible outcomes of agent interactions as a matrix in which each entry is the feedback that the agents involved will receive as a result of their choice of strategies.<sup>1</sup> We model coordination by the following matrix:

	A	B
A	+1, +1	-1, -1
B	-1, -1	+1, +1

In this case, agents have two strategies from which to choose. It is not important which particular strategy a given agent uses, but it is best if two interacting agents use the same strategy.<sup>2</sup> A simplified example of such a situation from a mobile robotics domain is deciding who enters a room first: a robot going in or a robot going out. If some robots use one strategy and some

<sup>1</sup>Although the matrix we use to model feedback is analogous to the payoff matrix formalism in game theory, it is important to note that we assume neither that agents are "rational" nor that they can access the contents of the matrix directly.

<sup>2</sup>We limit our discussion here to the two-strategy case, but the results are qualitatively similar when more strategies are available to agents.

robots use the other, there will be much unnecessary maneuvering about (or knocking of heads) when two robots attempt to move through a door simultaneously. If all robots use the same strategy, the system will run much more smoothly. This is reflected in the matrix entries: there is positive feedback for two agents using the same strategy and negative feedback for two agents using different strategies. Since there is no *a priori* preference between the two available strategies, the feedback matrix is symmetric with respect to them.

Agents update their strategies based upon the contents of a finite memory. Each memory element records the time of an interaction, the strategy used by the agent in that interaction, and the feedback the agent received as a result of that interaction. When an agent receives new feedback, it discards its oldest memory, to maintain the memory at a fixed size. Currently, we make the rather weak assumption that interactions are anonymous; although this reduces the amount of information available to each agent, we believe that exploring the behavior of societies of simple agents will yield insight into the behavior we can expect from more complex agents.

For our preliminary investigations, we have chosen to use a learning rule similar to the Highest Cumulative Reward rule used by Shoham and Tennenholtz (Shoham & Tennenholtz 1993). To decide which strategy it will use, an agent first computes the cumulative reward for each strategy by summing the feedback from all interactions in which it used that strategy and then chooses the strategy with the highest cumulative reward (HCR). There are, of course, many other possible learning rules agents could use, including more sophisticated reinforcement learning techniques such as Q-learning (Watkins 1989); however, using the simpler HCR rule fits with our program of starting with a simpler system.

In preliminary experimentation, we found that small memory sizes allow the agent society to achieve a convention rapidly; we consistently used an agent memory of size 2 in the experiments described in this paper. Thus, each agent chooses a new strategy based only on the outcome of the previous two interactions.

### Locality and Authority

In practice, locality can arise in two general ways, either as an inherent part of the domain or as a design decision; we will examine localization models that reflect both of these sources. Whatever its origin, we implement localization with non-uniform interaction probabilities: an agent is more likely to interact with those agents to which it is "closer" in the society.

**Two-Dimensional Grids** Consider two systems of mobile robots. In both societies, there is an inherent locality: each robot has some restricted neighborhood in which it moves (perhaps centered on its battery charger). In one society, each robot is confined to a

small domain and only interacts with other robots that are nearby. In the other system, the robots' domains are less rigid; although they generally move in a small area, they occasionally wander over a greater distance. Now assume that in both systems, each robot randomly chooses between two available strategies. The robots then interact pairwise, receive feedback as specified by the matrix describing coordination, and update their strategies according to the HCR rule.

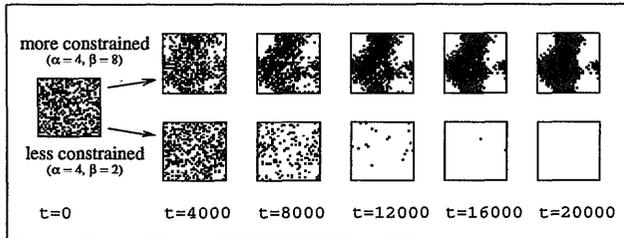


Figure 1: Time evolution example for agents on a grid.

A typical example of the time evolution of two societies fitting this description can be seen in Figure 1. In our initial investigations, the agents (robots) occupy a square grid, with one agent at each grid site; in this case there are 1024 agents on a 32 by 32 grid. The agents are colored white or black according to their choice of strategy. In the society at the top of the figure, the agents are tightly constrained, while those in the society at the bottom have more freedom to wander. Both systems start from the same initial configuration, but their evolution is quite different. In the system of agents with limited interaction, we see the spontaneous development of coordinated sub-societies. These macroscopic structures are self-supporting—agents on the interior of such a structure will have no cause to change strategy. The only changes will come at the edges of the sub-societies, which will wax and wane until eventually all of the agents are using one of the strategies. In the system of agents with more freedom of interaction, sub-societies do not appear to arise. The strategy which (perhaps fortuitously) gains dominance early is able to spread its influence throughout the society, quickly eliminating the other strategy.

In our model, we describe a robot's motion as a statistical profile that gives the likelihood of the robot wandering a particular distance from its "home" point. The probability of two robots interacting is thus a function of their individual movement profiles; we have modeled this by a simple function,  $p(r) \propto [1 + (\alpha r)^\beta]^{-1}$ , where  $r$  is the distance between the two robots, measured as a fraction of the size of the grid. This function was chosen because the parameters allow us to independently control the overall size of an agent's domain of interaction ( $\alpha$ ) and the "rigidness" of the domain boundary ( $\beta$ ). Figure 2 shows the function for a variety of parameter settings. With this function, we can

model robots confined to a large domain, robots confined to a small domain, robots that usually move in a small domain but occasionally wander over a larger area, etc. Note that the parameter  $\alpha$  controls where the probability is halved, regardless of the value of  $\beta$ ; in the limit  $\beta \rightarrow \infty$ , if  $r > 1/\alpha$  then  $p(r) = 0$ . We can think of  $1/\alpha$  as the "effective radius" of the probability distribution.

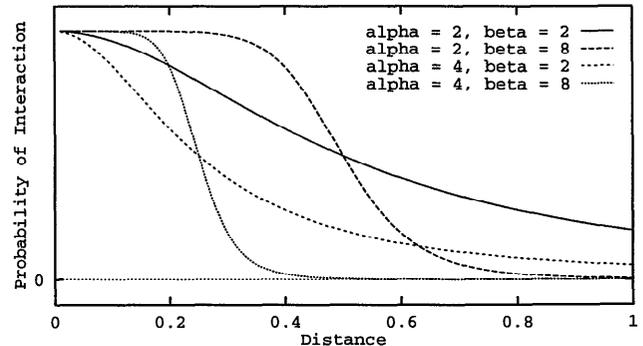


Figure 2:  $p(r) \propto [1 + (\alpha r)^\beta]^{-1}$ .

**Trees and Hierarchies** In many human communities, the societal organization is built up of layers. Hierarchies in large companies and "telephone trees" for distributing messages are examples of such structures. For our purposes, *trees* are defined in the standard fashion: each node has one parent and some number of children; one node, the root node, has no parent.<sup>3</sup> Agents organized as a tree can interact only with their parents and children. If we allow agents to interact with their peers—other agents on the same level—we call the resulting structure a *hierarchy*. We believe that these localizing structures may be useful in societies of artificial agents for some of the same reasons they have served humans well: delegation of responsibilities, rapid distribution of information, etc. Furthermore, trees and hierarchies provide a natural setting for investigating the effects of authority.

**Implementing Authority** In the present experiments, an agent in a tree or hierarchy is equally likely to interact with any other agent to which it is connected, be it parent, child, or peer. However, by giving agents the ability to selectively ignore feedback they receive from interactions with agents at a lower level, we can implement a simple form of authority. We refer to the feedback an agent receives when interacting with its child in the organization as "upward feedback," and

<sup>3</sup>We limit our discussion here to trees with a branching factor of two; additional experiments have shown that increasing the branching factor does not change the relative qualitative behavior of the tree and hierarchical organization schemes.

varying the probability that this upward feedback is incorporated into an agent's memory can be thought of as modeling a range of management styles, from egalitarian bosses who listen to and learn from their subordinates to autocratic bosses who expect their subordinates to unquestioningly follow the rule "Do as I do." In this preliminary model agents always heed feedback from their parents and peers.

### Experimental Results

For each experiment, the number of agents, social structure, and other parameters are fixed. Each agent's probability distribution for interacting with the other agents is computed, and the system is run multiple times. At the beginning of each run, the agents' initial strategies are chosen randomly and uniformly from the two available strategies. In each iteration of the run, a first agent is chosen randomly and uniformly from the society; a second agent is then chosen randomly according to the first agent's probability distribution. The agents interact and possibly update their strategies as described above. The system is run until all of the agents are using the same strategy, i.e. until we have 100% convergence. Each experiment was run 1000 times (with different random seeds), and the convergence time was computed by averaging the number of iterations required in each run.

**Two-Dimensional Grids** We begin our survey of experimental results by considering agents on a square grid, with interaction probabilities defined by the function  $p(r)$ . In Figure 3, we see how the time for all of the agents to agree upon a particular strategy scales with the size of the system, for various parameter values.

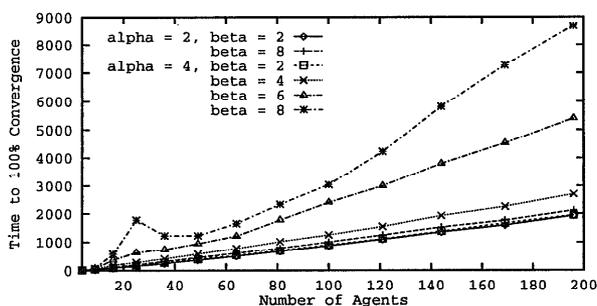


Figure 3: Convergence time vs. number of agents for two-dimensional grid.

In general, the convergence time appears to be polynomial in the number of agents in the system. Fitting curves to the data yields a range of empirical estimates from  $O(n^{1.28})$  for the least restricted societies ( $\alpha = 2, \beta = 2$ ) to  $O(n^{1.45})$  for the most restricted ( $\alpha = 4, \beta = 8$ ). To examine the interaction of the parameters in more detail, we fix the number of agents

in the system and observe how the convergence time is affected by various parameter settings.

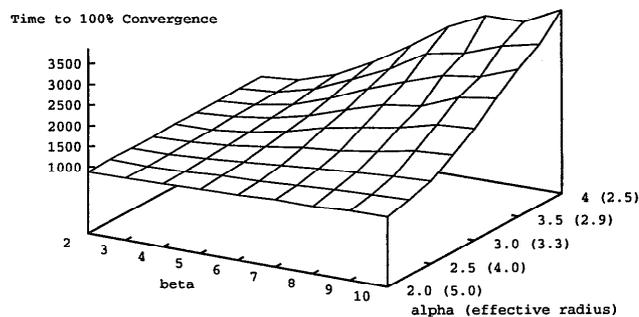


Figure 4: Convergence time for 100 agents on a two-dimensional grid for various parameter settings.

In Figure 4, we see this for a society of 100 agents; the effective radius in numbers of grid units is noted next to each value of  $\alpha$ . We find that the steepness of the drop-off in the interaction probability, controlled by  $\beta$ , becomes more and more significant as  $\alpha$  is increased. To think of it in terms of mobile robots, as the effective radius of a robot's domain is decreased, the rigidity of the boundary of its domain becomes increasingly relevant.

**Trees, Hierarchies, and Authority** We now look at the results of experiments with the tree and hierarchy organizational structures. Initially, we will assume full top-down authority, i.e. parent nodes never pay attention to feedback from their children.

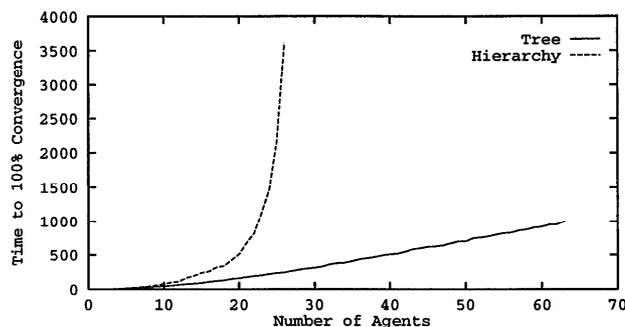


Figure 5: Convergence time vs. number of agents for tree and hierarchy organizations with full top-down authority.

In Figure 5, we see the effects of system size on the time to achieve total coordination. It appears that the convergence time for trees is polynomial in the number of agents, while for hierarchies, the convergence time seems to be exponential in the number of agents. Fitting to curves yields empirical estimates of  $O(e^{0.26n})$  for hierarchies and  $O(n^{1.85})$  for trees.

## Discussion

From the results of experiments with agents on a grid, we might speculate that increased interaction between agents promotes the emergence of conventions for coordinating actions. This is further borne out by the data in Figure 7—as the probability of upward feedback is increased, the amount of interaction effectively increases, and the system converges more rapidly. However, the behavior of trees seems to defy this conjecture: Figure 6 shows that increased interaction on a tree *decreases* the efficiency of convention emergence. Furthermore, although societies with top-down authority have *less* overall interaction when organized as trees rather than as hierarchies, they converge much more readily when tree-structured, as seen in Figure 5. It appears that neither locality nor authority is a sufficient predictor of system performance by itself.

To develop further intuition about the results with trees, hierarchies, and authority, we can observe the time evolution of representative systems. In Figure 8, we see four societies of 63 agents. The systems depicted represent the possible combinations of tree vs. hierarchy and top-down authority vs. no authority. All four systems start from the same initial condition, but they evolve quite differently.

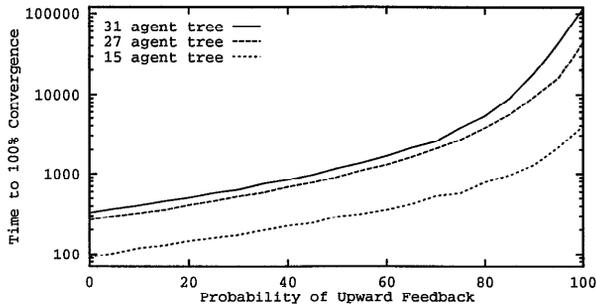


Figure 6: Effect of decreasing authority (increased upward feedback) on convergence time for tree organization.

Now we increase the probability of upward feedback, reducing the authority of agents over their descendants in the tree. In Figure 6, the convergence time is plotted against the probability of upward feedback for three tree-structured systems of different sizes. We see that the time for total coordination to be achieved increases exponentially (the y-axis is logarithmic) with increasing upward feedback, until a probability of about 75% is reached, at which point the convergence time increases even more dramatically.

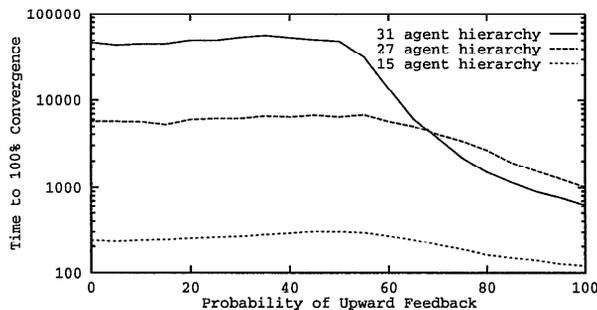


Figure 7: Effect of decreasing authority (increased upward feedback) on convergence time for hierarchy organization.

In Figure 7, the convergence time is plotted against the probability of upward feedback for the same three system sizes, now organized as hierarchies. In this case, the convergence time increases slightly with decreasing authority until a probability of about 50% is reached, at which point the society begins achieving coordination ever more rapidly. It appears that while authority is useful in trees, agents in hierarchies should listen to their subordinates.

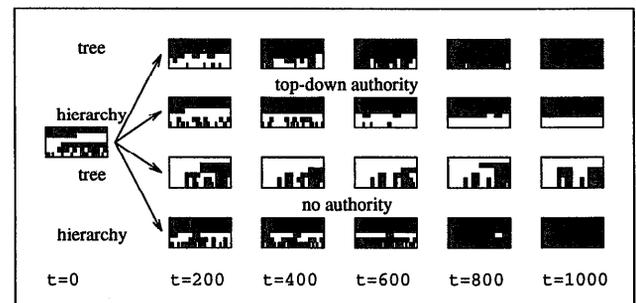


Figure 8: Time evolution example for trees and hierarchies. In this diagram, parent nodes are drawn twice as wide as their children. Thus, coordinated subtrees appear as vertical rectangles.

In the authoritarian tree, we see that there is a strong directional pressure “pushing” a convention through the tree. Each node receives feedback only from its parent, and will quickly adopt its parent’s strategy. We can think of this pressure as defining a “flow of information” through the society. This contrasts with the authoritarian hierarchy, in which each level of the organization is completely connected internally, but only weakly connected to the next level. The deeper a node is in the graph, the less likely it is to interact with its parent, and the flow of information is diluted. It becomes possible for adjacent levels to independently converge upon *different* conventions, hence we see a horizontal division in strategies at  $t = 1000$ .

When we eliminate authority by having upward feedback occur with 100% probability, we increase the po-

tential for inter-level interaction. In trees, this causes the flow of information to become incoherent and we are left with a sprawling, weakly connected society. Subsocieties emerge, but now they develop on subtrees, rather than across levels. For hierarchies, we saw that reducing top-down authority causes the convergence time to decrease, and in the bottom row of Figure 8, we see one of the reasons this happens: a level which might have otherwise converged to an independent convention is rapidly brought into line with the rest of the system.

## Conclusion

We have seen that locality and authority in a multi-agent society can profoundly affect the evolution of conventions for behavior coordination. We draw three (tentative) conclusions. First, in weakly connected systems, there is a tendency for sub-societies to form, hampering convergence.<sup>4</sup> Conversely, systems with greater overall interaction tend to converge more rapidly. Finally, if our agents are weakly connected, whether by design or necessity, it appears best to have a directional flow of feedback—strong authority—to ensure the rapid spread of conventions throughout the society.

There are numerous ways this work can be extended. In addition to exploring other forms of locality and authority, we can investigate noise, more complex learning algorithms, and explicit communication. We have already begun experimentation with emergent cooperation, and preliminary results seem to indicate that organizing structures which promote the emergence of conventions for coordination do not necessarily serve the objective of emergent cooperation. Ultimately, we would like to develop an analytic theory of convention emergence; experiments such as these serve both to guide exploration and to test theoretical results.

Research into various forms of emergent behavior in multi-agent systems relates to our investigations, particularly work on coordination and cooperation among agents. While our model incorporates adaptive agents that learn purely from experience, many researchers have taken the view that agents should be treated as “rational” in the game-theoretic sense, choosing their actions based on some expected outcome ((Stary 1993) is an overview of these approaches). In (Glance & Huberman 1993), Glance and Huberman approach closest to our work, exploring the effects of a dynamic social structure on cooperation between agents. In their model, an agent decides whether or not to cooperate by examining the behavior of other agents; if the agent perceives enough cooperation in its environment, then it decides to cooperate as well. This differs significantly from our model, in which the actual

<sup>4</sup>If agents from different components interact only very rarely, it may not be important for the subsocieties to have the same convention.

results of an agent's actions cause it to adapt its behavior. Glance and Huberman implement locality by weighting agents' perceptions of one another according to their proximity within the organization. Our notion of locality differs because it is not based on an expectation of interaction: it is based on the actual occurrence or non-occurrence of interactions.

This research bears a close resemblance to work in economics and game theory (Kandori, Mailath, & Rob 1991). One of our current goals is to gain a better understanding of this relationship. More generally, research on adaptive multi-agent systems appears to have ties to work in artificial life (Lindgren 1992), population genetics (Mettler, Gregg, & Schaffer 1988), and quantitative sociology (Weidlich & Haag 1983). However, while systems from these diverse fields share characteristics such as distributed components and complex dynamics, their particulars remain unreconciled.

## References

- Glance, N. S., and Huberman, B. 1993. Organizational fluidity and sustainable cooperation. In *Proceedings of the Modeling Autonomous Agents in a Multi-Agent World conference*. In press.
- Kandori, M.; Mailath, G.; and Rob, R. 1991. Learning, Mutation and Long Equilibria in Games. Mimeo. University of Pennsylvania.
- Lindgren, K. 1992. Evolutionary phenomena in simple dynamics. In *Artificial Life II*. Santa Fe Institute.
- Mettler, L. E.; Gregg, T. G.; and Schaffer, H. E. 1988. *Population Genetics and Evolution*. Prentice Hall, second edition.
- Shoham, Y., and Tennenholtz, M. 1992a. Emergent conventions in multi-agent systems: initial experimental results and observations. In *KR-92*.
- Shoham, Y., and Tennenholtz, M. 1992b. On the synthesis of useful social laws for artificial agent societies. In *Proceedings of the Tenth National Conference on Artificial Intelligence*. AAAI.
- Shoham, Y., and Tennenholtz, M. 1993. Co-learning and the evolution of social activity. Submitted for publication.
- Stary, C. 1993. Dynamic modelling of collaboration among rational agents: redefining the research agenda. In *IFIP Transactions A (Computer Science and Technology)*, volume A-24. Human, Organizational and Social Dimensions of Information Systems Development. IFIP WG8.2 Working Group.
- Watkins, C. 1989. *Learning from Delayed Rewards*. Ph.D. Dissertation, King's College.
- Weidlich, W., and Haag, G. 1983. *Concepts and Models of a Quantitative Sociology; The Dynamics of Interacting Populations*. Springer-Verlag.