698

# Synthesizing Complex Behaviors by Composing Simple Primitives

Maja J. Matarić and Matthew J. Marjanović
MIT Artificial Intelligence Lab
545 Technology Square #721 and #803
Cambridge, MA 02139
maja@ai.mit.edu, maddog@ai.mit.edu

April 14, 1993

### Abstract

This paper presents an approach to describing and synthesizing group behavior using simple local interactions among individuals. We propose that a set of basic interactions can be defined which can be used to describe a large variety of group behaviors within a given domain. This approach utilizes agent homogeneity and minimal communication to simplify synthesis and analysis of group behavior.

To validate our approach we implemented a battery of basic group behaviors implemented in the domain of physical spatial interactions. We present some of the experimental results from two distinct environments: a software Interaction Modeler and a collection of 20 mobile robots. We also describe an aggregate behavior involving a combination of the primitive behaviors.

# 1 Introduction

Group behavior is a result of the local interactions between the members of the group, and their interactions with the environment. Local dynamics between individuals produce consequences at the global level. Thus, the group behavior of interacting agents can neither be predicted nor analyzed by observing a single individual. Unfortunately, analyzing a collection of agents is much more complex than the already difficult single-agent case.

In this work, we present a methodology for describing group behavior which allows for simplified synthesis and analysis. Our approach consists of viewing group behavior as a collection of basic, primitive behaviors, built from simple local interactions, and combined into more complex aggregates. We apply our methodology to two groups of agents (software and hardware), evaluate and analyze the experimental data, and discuss the ramifications of the results.

# 2 Inference and Interference

In order for a society of independent agents to function, it must overcome situations of persisting interference between its members. Even in the simplest society in which all agents have identical goals[1] at all times, conflicts such as competition for resources can arise. More diverse societies, where agents' goals differ, demonstrate increasingly complex conflicts, including clobbering of others' work, deadlock, and oscillations.

Conflict among agents can be avoided if each agent infers the goals of others. However, the ability to make such inferences requires a high computational and cognitive overhead (Gasser & N. Huhns 1989, Rosenschein & Genesereth 1985, Axelrod 1984). Work in both ethology and developmental psychology suggests that inferring the goals of other agents is not necessary for many complex interactions (Tomasello, Kruger & Rather 1992, McFarland 1987, Gould 1982, Rosenthal & Zimmerman 1978).

The alternative to inferring the goals of others is to base interactions only on observable behavior. The interpretation of the observed behavior is determined by the amount of knowledge available to the agent. In biology, this knowledge is innate and difficult to circumscribe. In contrast, computational and robotic experiments allow for varying the goals and the amount of built-in knowledge. Our work has demonstrated that significant information about individual goals is reflected in behavior, and can be obtained by observation without communication. ??

# 3 Agent Homogeneity

To evaluate how much knowledge and communication is necessary, we have focused on the simplest, but by no means simple, form of a society: one consisting of homogeneous agents. The agents are homogeneous in that they

---

[1] We define a goal to be whatever the agent is "programmed to do." Goals can be expressed explicitly or implicitly.

are situated in the same world, embodied with similar dynamics (i.e. same physical description) and have similar goals (i.e. same program).

Homogeneity has important implications. Identical agents have an innate source of knowledge about each other, which allows for leaving much of the information about the world implicit. Since agents share common goals, their behavior is predictable to each other.

Homogeneity also offers a society flexibility in that agents are interchangeable. They need neither identities nor the ability for individual identification. Furthermore, in sufficiently large groups, irregular behavior of any individual should not seriously affect the group, since no one agent is more important than any other.

Taking advantage of homogeneity depends on a key property: agents must be able to recognize other agents of the same kind. With this ability, which is innate and ubiquitous in nature, even the simplest of local interactions can produce purposive collective behavior. For example, while driving on a two-lane road[2] and faced with an oncoming car, one is confident that the correct behavior is to stay on the right, since the other car will follow the same strategy. However, if instead of a car, a cow is approaching, there is no way of predicting what the cow will do or what the correct response is. Homogeneity greatly reduces individual cognitive requirements. We use this property for simplifying the processes of generating and understanding group behavior.

## 4 Basic Interactions

Regardless of the simplicity of the individual agents, the global consequences of even the simplest local interactions can be arbitrarily complex. In general, it is impossible to predict precisely or even qualitatively what the global-level behavior of such a system with interacting components will be (Mataric 1992, Weisbuch 1991, Wiggins 1990, Nicolis & Prigogine 1989). Societies are by nature complex systems and as such cannot be usefully analyzed with traditional methods.

While it is impossible to predict the behavior of an arbitrary society, we propose that it is possible to perform qualitative analysis if the behavior of the system can be represented as a collection of *basic interactions* whose dynamics are well understood. Basic interactions are behaviors which can be considered typical for a particular society. These behaviors are stable, repeatable, observable at a global level, and determined by the goals and local interactions of the individuals. Basic group behaviors may be observable in physical space (such as flocking, herding, following, traffic jams), or in some abstract interaction space of the society.

Our work is based on the hypothesis that most of group behavior consists of such basic interactions. While the exact behavior of each individual may not be known, the collective behavior is qualitatively predictable and repeatable. In practice, most societies are too complex (in terms of the properties of the individuals as well as their interactions) to be modeled analytically. However,

---

[2] In the United States, for example.

the notion of stable group interactions can be used for qualitative analysis, as well as for designing group behaviors.

We have applied the concept of basic group behaviors to a collection of software and hardware agents, and have focused on their manifested, observable interactions. By placing our experiments in physical space, we demonstrate group behavior with simple examples of physical interactions and spatial patterns. We used the constraints imposed by the environment and the mechanics of the agents to construct a set of basic interactions we call *behavior primitives* (Mataric 1992) which allow for a variety of group behaviors. The next section describes our experimental environments, methodology, and results.

## 5 Experimental Methodology and Results

Biology, sociology, and anthropology provide important inspiration for this work. However, our goal is not to simulate biological systems, but to study group interaction by synthesizing, observing, and analyzing similar phenomena. Since behavior observation is the primary methodology for validating theories in this work, it is important to separate artifacts of the experimental environment from effects intrinsic to the interaction being observed. Toward this end, two very different test environments are used, one in software and one in hardware.

The software environment consists of a Lisp program called the Interaction Modeler which allows for modeling of agents with very simple dynamics and sensing. This modeled environment is not intended to be a testbed for developing strategies for physical robots. Rather, it is used for observing and testing the behaviors of groups of different sizes and compositions, and comparing them to behaviors observed on physical robots. In experiments so far, the agents were programmed to have local sensors for detecting nearby agents and objects in the world. The software agents had no centralized control or global communication capabilities.

The hardware environment consists of a collection of 20 mobile robots. Each robot is a 12"-long four-wheeled vehicle equipped with proximity and bump sensors. In front, each robot has a two-pronged forklift for picking up, carrying, and stacking metal pucks. The robots are also equipped with radio transceivers for limited data communication (seven bytes per second per robot, ignoring a large error rate) and a sonar-based global positioning system. The robots are programmed in the Behavior Language (BL), a behavior-based parallel programming language (Brooks 1990). The sensors and the radio system allow the robots to: explore their environment; detect each other; and find, pick up, and carry pucks. These elementary abilities form the basis of the robots' various tasks and experiments, in which they run autonomously, with all processing and power carried on-board.

An important property of the physical implementation is the inevitable variability among the agents. Even after calibration, characteristics of the sensors and actuators differ between robots. This fundamental property provides a stringent test for the robustness of our algorithms.

## 5.1 Communication and Cooperation Limitations

Communication is a necessary tool for some but not all types of group behavior. In order to determine just how necessary it is, we conducted experiments comparing behaviors which use no direct communication with counterparts that added a local broadcast ability.

No explicit one-to-one communication between agents is used in any of these experiments. Instead, communication is based on sensing the detectable state of nearby agents, and (in some experiments) broadcasting a simple message to a limited area around the agent. Thus, no explicit cooperation exists between agents. We define *explicit* cooperation as a set of interactions which involve exchanging information or performing actions in order to help another agent. In contrast, *implicit* cooperation includes actions that are a part of the agent's own goal-achieving behavior, but may have effects in the world that help other agents achieve their own goals. In the experiments we describe, cooperation is implicit, as agents affect one another by means of their position and detectable state.

These communication and cooperation constraints were chosen in order to test the limits of implicit communication as advocated by the previously described developmental psychology and ethology theories.

## 5.2 Experimental Group Behaviors

We have developed a collection of simple local rules that implement the following basic behaviors:

**Collision Avoidance:** the ability of an agent to avoid colliding with anything in the world. Two distinct strategies were devised; one for "kin" (other agents of the same kind), and another for everything else that the robot might encounter.

**Following:** the ability to stay behind or along side of another agent without colliding (Figure 1).

**Dispersion:** the ability of a group of agents to spread out over an area in order to maintain some predetermined minimum separation (Figure 2).

**Aggregation:** the ability of a group of agents to gather within some predetermined maximum distance. This behavior is the inverse of dispersion (Figure 3).

**Homing:** the ability of one or a group of agents to reach a goal region or location (Figure 4).

**Flocking:** the ability of a group of agents to move as a coherent aggregate without prespecified leaders and followers. Flocking includes components of collision avoidance, following, dispersion and aggregation (Figure 5).

These behaviors were shown to be repeatable over many time-extended trials. They are stable over a variety of initial conditions, and insensitive to small perturbations in the sensor and actuator operation. Furthermore, most of the behavior primitives were implemented with multiple algorithms, with variations in performance due to the implementation details, but with consistent overall behavior . The details of the robot implementations are

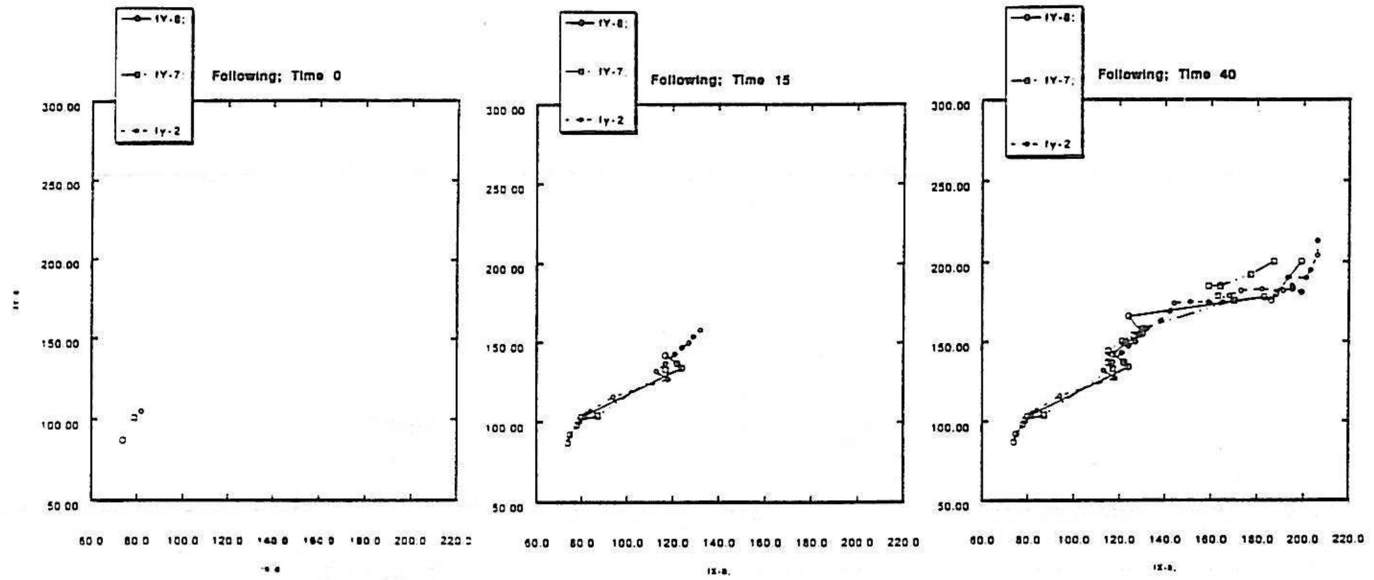

Figure 1: The following behavior as demonstrated by three robots. The points mark the robots' positions based on radio data; the lines are the interpolated path.

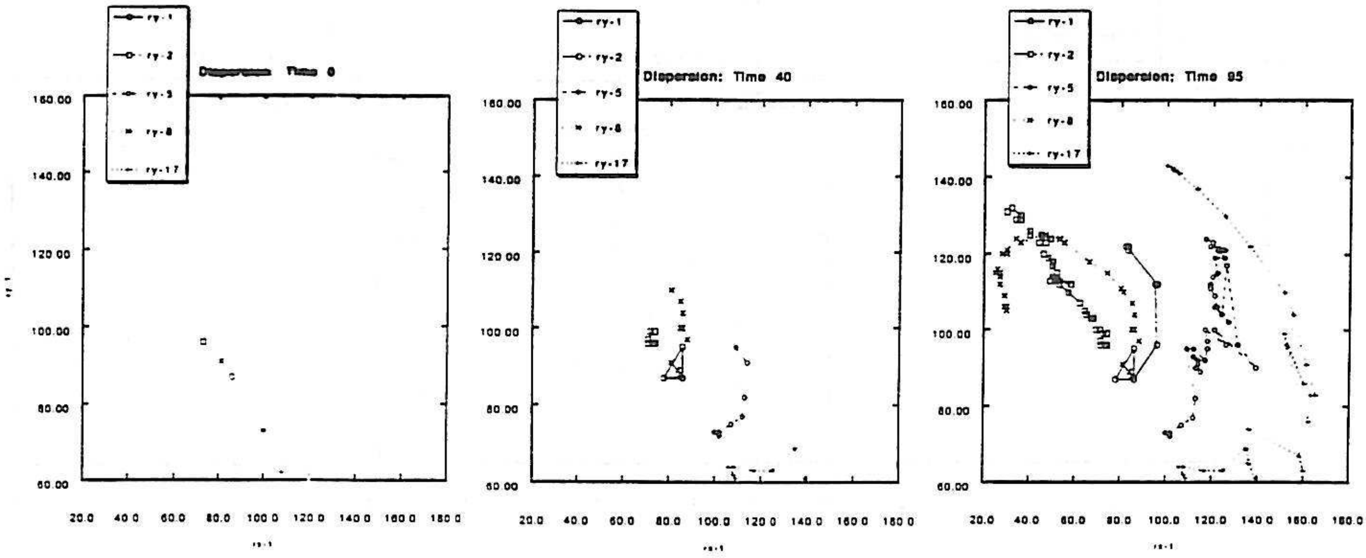

Figure 2: The dispersion behavior. The irregularities in the path are due to radio transmission errors generating flawed position data.

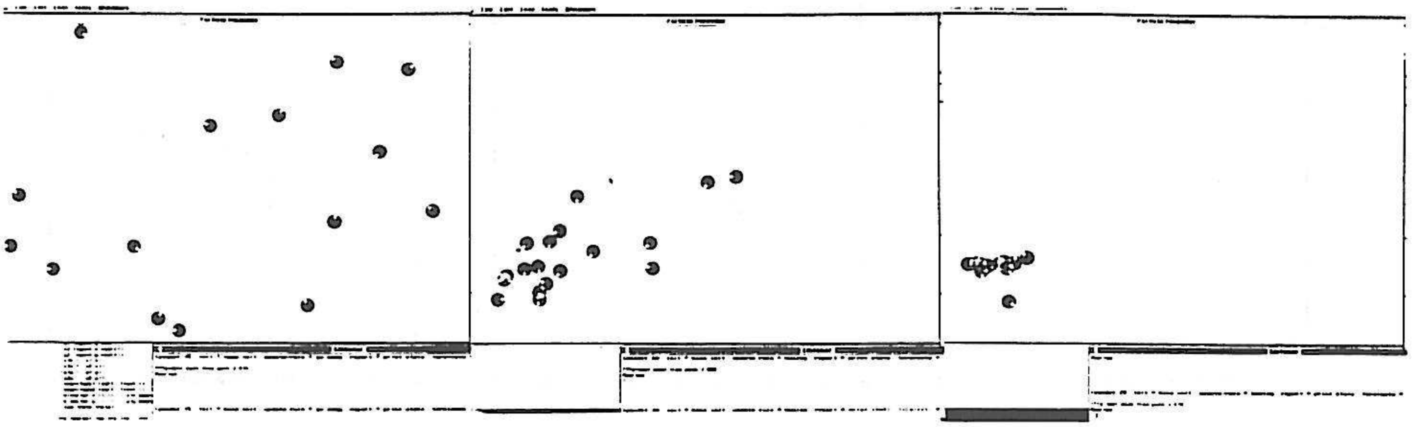

Figure 3: The aggregation behavior as demonstrated by 15 Interaction Modeler agents.

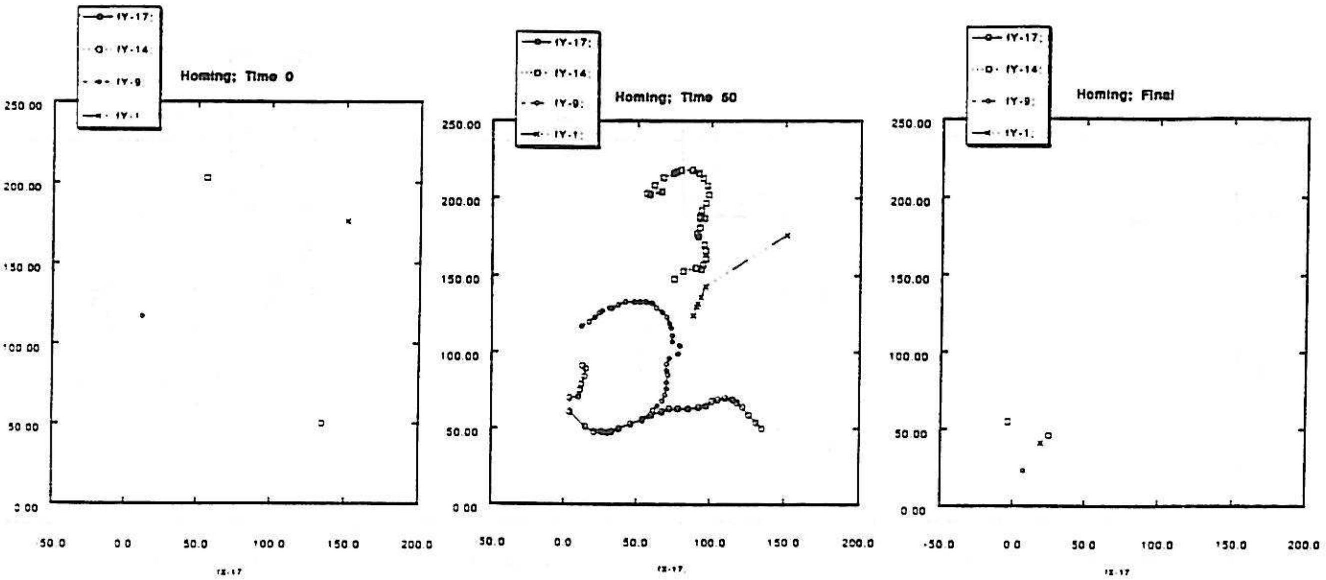

Figure 4: The homing behavior demonstrated by four robots. Home is located in the (0,0-50-50) region. The scale shown is in centimeters.

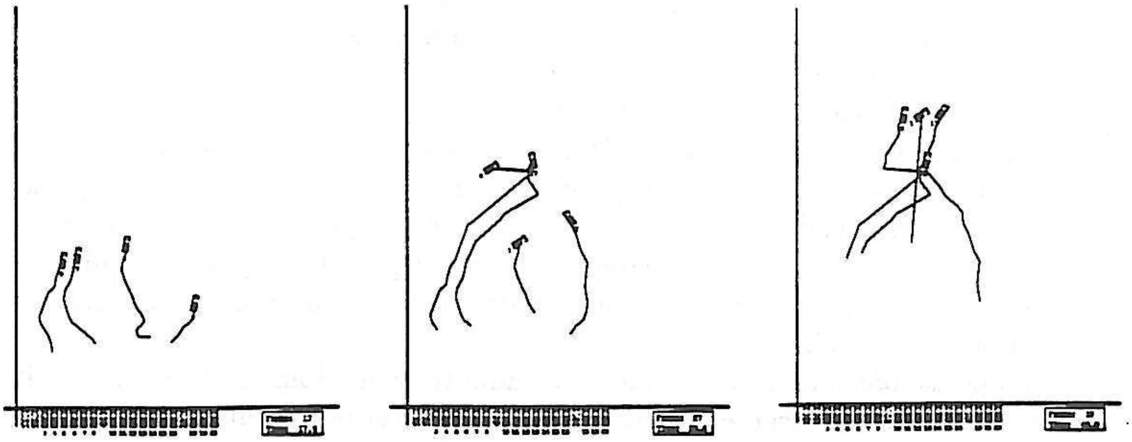

Figure 5: The flocking behavior as demonstrated by four robots. The graph plots actual robot paths. The room and robots are drawn to scale.

described in Mataric (1992). Additional experimental data is available in the form of both video tapes and radio transcripts.

# 6 Combining Basic Behaviors

We are currently experimenting with a foraging behavior, which combines the described basic behaviors. In foraging, the robots collect and take home pucks which they find dispersed around the workspace. In addition to the basic interactive behaviors, the individual agents also have primitives for recognition, grabbing, and dropping of pucks. They have no model of the environment, nor a global view of it, except for the preprogrammed coordinates of home.

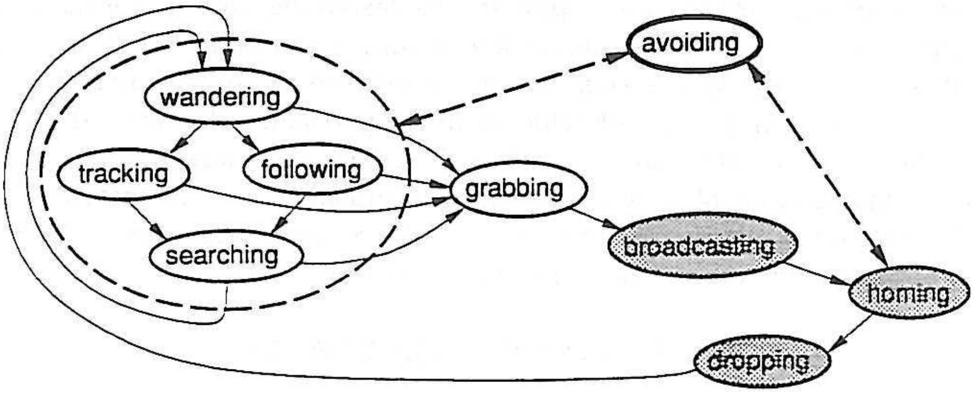

Figure 6: The behavior structure of the foraging agents. The shaded states indicate that the agent is carrying a puck.

Each robot runs the same BL program, which resembles a finite state machine that activates and deactivates the various primitive behaviors (Figure 6). The default state is *wandering*, in which a robot randomly moves around the room. Avoidance behaviors are allowed to interrupt *wandering*, acting as survival instincts which keep the robots from getting stuck against walls or each other. If, while *wandering*, the robot detects a puck in its path, it will enter a *grabbing* state, in which it picks up the puck, and then a *broadcasting* state, in which it transmits a radio message telling other robots that it found a puck. The robot then heads home (*homing*), drops off the puck (*dropping*), and resumes *wandering*. A single, isolated robot will repeat this cycle of randomly searching for pucks and bringing them home. When in a group, however, two variations are possible.

While *wandering*, a robot listens to radio transmissions in its vicinity.[3] If it receives an announcement that a fellow robot has picked up a puck, it will enter the *tracking* state and home in on the origin of the announcement, with the assumption that one puck found means that more are nearby. Once it arrives, the robot will begin *searching*, a more localized form of *wandering*. If no puck is found within a fixed time period, the robot resumes *wandering*.

While *tracking*, a robot transmits a message which, when received by nearby *wandering* robots, induces them to enter the *following* state and pursue the *tracking* robot. When the tracker stops *tracking* and begins *searching*, the followers do the same. With this minimal level of communication, a group of robots maintains some ephemeral knowledge of its world (i.e. where pucks are) and can cooperate to achieve a goal (i.e. bring those pucks home).

Our experiments have demonstrated two points: 1) coherent group behaviors in robots can result from simple local interactions; and 2) complex, time-extended collective behaviors can be generated by the combination of a few basic primitives. We plan to observe and analyze similar simple interactions which produce global consequences in other group domains.

## 7 Summary

This paper has discussed an approach to describing and synthesizing group behavior using simple local interactions among individuals. We have proposed that a set of basic interactions can be found which can be used to describe a large variety of group behavior within a particular domain. To validate our theory we created an array of basic group behaviors in the domain of spatial interactions of mobile agents. We implemented and tested the basic behaviors in two distinct environments and are continuing to experiment with composition and combination of these behaviors.

## Acknowledgements

Many thanks to Stan Wang and Owen Wessling for making the unruly robot herd behave, and helping its collective intelligence emerge in spite of many

---

[3] To enforce a degree of locality, robots ignore messages received from beyond a fixed radius.

hardware problems. Thanks to Tim Smithers for a detailed and insightful review of an earlier draft of this paper.

The research reported here was done at the MIT Artificial Intelligence Laboratory. Support for this research was provided in part by the Jet Propulsion Laboratory contract 959333 and in part by the Advanced Research Projects Agency under Office of Naval Research grant N00014–91–J–4038.

## References

Axelrod, R. (1984), *The Evolution of Cooperation*, Basic Books, New York.

Brooks, R. A. (1990), The Behavior Language; User's Guide, Technical Report AIM-1127, MIT Artificial Intelligence Lab.

Gasser, L. & N. Huhns, e. (1989), *Distributed Artificial Intelligence*, Pitman, London.

Gould, J. L. (1982), *Ethology; The Mechanisms and Evolution of Behavior*, W. W. Norton & Co., New York.

Mataric, M. J. (1992), Designing Emergent Behaviors: From Local Interactions to Collective Intelligence, *in* 'From Animals to Animats: International Conference on Simulation of Adaptive Behavior'.

McFarland, D. (1987), The Oxford Companion to Animal Behavior, *in* 'Oxford, University Press'.

Nicolis, G. & Prigogine, I. (1989), *Exploring Complexity*, W. H. Freeman and Co., New York.

Rosenschein, G. & Genesereth, M. R. (1985), Deals Among Rational Agents, *in* 'IJCAI-85', pp. 91–99.

Rosenthal, T. L. & Zimmerman, B. J. (1978), *Social Learning and Cognition*, Academic Press, New York.

Tomasello, M., Kruger, A. C. & Rather, H. H. (1992), 'Cultural Learning', *to appear in The Journal of Brain and Behavior Sciences*.

Weisbuch, G. (1991), Complex System Dynamics, *in* 'Lecture Notes Vol. II, Santa Fe Institute Studies in the Sciences of Complexity', Addison-Wesley, New York.

Wiggins, S. (1990), *Introduction to Applied Nonlinear Dynamical Systems and Chaos*, Springer-Verlag, New York.