Issues and Approaches in Design of Collective Autonomous Agents

Maja J Matarić
Volen Center for Complex Systems
Computer Science Department
Brandeis University
Waltham, MA 02254
tel: (617) 736-2708, fax: (617) 736-2741
maja@cs.brandeis.edu
http://www.ai.mit.edu/people/maja/maja.html

November 30, 1994

Abstract

The problem of synthesizing and analyzing collective autonomous agents has only recently begun to be practically studied by the robotics community. This paper overviews the most prominent directions of research, defines key terms, and summarizes the main issues. Finally, it briefly describes our approach to controlling group behavior and its relation to the field as a whole.

1 Introduction

The problem of synthesizing and analyzing collective autonomous behavior has only recently begun to be practically studied by the robotics community. This paper overviews the most prominent directions of research, defines key terms, and summarizes the main issues. Finally, it briefly describes our approach to controlling group behavior and its relation to the field as a whole.

2 Overview of Work in the Field

2.1 Physical Multi–Robot Systems

The last decade has witnessed a shift in research emphasis toward physical implementations of robotics in general and mobile robotics in particular. Most of the work in robotics so far has focused on control of a single agent, but a few efforts have begun to address multi–robot systems. Fukuda, Nadagawa, Kawauchi & Buss (1989) and subsequent work describe an approach to coordinating multiple homogeneous and heterogeneous mobile robotic units, and demonstrate it on a docking task. Caloud, Choi, Latombe, LePape & Yim (1990) and Noreils (1993) remain faithful to the state–based framework, and apply a traditional planner–based control architecture to a box–moving task implemented with two robots in a master–slave configuration. Kube, Zhang & Wang (1993) and Kube & Zhang (1992) describe a series of simulations of robots performing a collection of simple behaviors that are being incrementally transferred to physical robots. Barman, Kingdon, Mackworth, Pai, Sahota, Wilkinson & Zhang (1993) report on a preliminary testbed for studying control of multiple robots in a soccer–playing task. Parker (1993*b*) and Parker (1993*a*) describe a behavior–based task–sharing architecture for controlling groups of heterogeneous robots, and demonstrates it on a group of four physical robots performing toxic waste cleanup and box pushing. Donald, Jennings & Rus (1993) report on the theoretical grounding for implementing a cooperative manipulation task with a pair of mobile robots. Perhaps closest in philosophy as well as the choice of task is work by Altenburg (1994), describing a variant of the foraging task using a group of LEGO robots controlled in reactive, distributed style, and Beckers, Holland & Deneubourg (1994), demonstrating a group of four robots clustering initially randomly distributed pucks into a single cluster through purely stigmergic communication.

In terms of cooperation and communication, most of the above work has fallen along the two ends of the spectrum: it either uses extensive explicit communication and cooperation, or almost none at all. In systems that are cooperative by design, the robots are aware of each other's existence, and can sense and recognize each other directly or through communication. This type of research explores explicit cooperation, usually through the use of directed communication, and is represented by the work of Caloud et al. (1990), Noreils (1993), and Parker (1993*a*).

The other category includes work on implicit cooperation, in which the robots usually do not recognize each other but merely coexist and indirectly cooperate by having identical or at least compatible goals. Such work includes Dallas (1990) and Kube et al. (1993). Our work, described at the end of the paper, falls nearer this end of the spectrum, but is focused on agents that can discriminate each other from the rest of the world, and can use this ability as a basis for social behavior.

2.2 Simulations of Multi–Agent Systems

With the exception of the work described above, the problem of multi–agent control has been treated mostly in simulation and under two major categories: simulations of situated systems and simulations of abstract agents.

Simulations of situated systems involve some degree of faithfulness to the physical world, to the extent of employing simple models of sensors, effectors, and physical laws. A number of simulations of behavior–style controlled systems have been implemented. For instance, Steels (1989) describes a simulation of simple robots using the principles of self–organization to perform a gathering task. Brooks, Maes, Matarić & Moore (1990) report on a set of simulations in a similar task domain, with a fully decentralized collection of non–communicating robots. Arkin (1992) describes a schema–based approach to designing simple navigation behaviors, used for programming multiple agents working in a simulated environment with future extensions to physical agents; Arkin, Balch & Nitz (1993) apply the approach to a multi–agent retrieval task. Brock, Montana & Ceranowicz (1992) describe SIMNET simulations of large numbers of tank–like robots performing avoidance and formation following. Kube et al. (1993) propose a behavior–arbitration scheme that will be tested on physical robots. Simulations tend to simplify both sensing and actuation. Physically–based simulations, however, using realistic physics models of the agent, allow for generating and testing more realistic behavior. For example, Hodgins & Brogan (1994) describe experiments with fully physically–based simulations of groups of hopping robots.

In contrast to simulations of multiple robots, "swarm intelligence" refers to simulations of abstract agents dealing with more theoretical problems of communication protocols, the design of social rules, and strategies for avoiding conflict and deadlock often in societies with with large numbers of simple agents. Representative work includes Fukuda, Sekiyama, Ueyama & Arai (1993), Dario & Rucci (1993), Dudek, Jenkin, Milios & Wilkes (1993), Huang & Beni (1993), Sandini, Lucarini & Varoli (1993), Kurosu, Furuya & Soeda (1993), Beni & Hackwood (1992), Dario, Ribechini, Genovese & Sandini (1991), and many others. Swarm intelligence work is also related to DAI (see below) but deals with agents of comparatively low cognitive complexity.

2.3 Artificial Life

The field of Artificial Life (Alife) focuses on bottom–up modeling of various complex systems, including simulations of colonies of ant–like agents, as described by Corbara, Drogoul, Fresneau & Lalande (1993), Colorni, Dorigo & Maniezzo (1992), Drogous, Ferber, Corbara & Fresneau (1992), Travers (1988), and many others. Deneubourg, Goss, Franks, Sendova-Franks, Detrain & Chretien (1990) and related work have experimented with real and simulated ant colonies and examined the role of simple control rules and limited communication in producing trail formation and task sharing. Deneubourg, Theraulax & Beckers (1992) define some key terms in swarm intelligence and discuss issues of relating local and global behavior of a distributed system. Assad & Packard (1992), Hogeweg & Hesper (1985) and other related work also report on a variety of simulations of simple organisms producing complex behaviors emerging from simple interactions. Schmieder (1993) reports on an experiment in which the amount of "knowledge" agents have about each other is increased and decreased based on local encounters. Werner & Dyer (1990) and MacLennan (1990) describe systems that evolve simple communication strategies. On the more theoretical end, Keshet (1993) describes a model of trail formation that fits biological data.

Our work is related to Artificial Life in that both are concerned with exploiting the dynamics of local interactions between agents and the world in order to create complex global behaviors. However, work in Alife does not typically concern itself with agents situated in physically realistic worlds. Additionally, it usually deals with much larger populations sizes that the work presented here. Finally, it most commonly employs genetic techniques for evolving the agents' comparatively simple control systems.

2.4 Distributed Artificial Intelligence

Distributed Artificial Intelligence (DAI) also deals with multi–agent interactions (see Gasser & Huhns (1989) for an overview). DAI focuses on negotiation and coordination of multi–agent environments in which agents can vary from knowledge–based systems to sorting algorithms, and approaches can vary from heuristic search to decision theory. In general, DAI deals with cognitively complex agents compared to those considered by the research areas described so far. However, the types of environments it deals with are relatively simple and low complexity in that they feature no noise or uncertainty and can be accurately characterized.

DAI can be divided into two subfields: Distributed Problem Solving (DPS) and Multi–Agent Systems (MAS) (Rosenschein 1993). DPS deals with centrally designed systems solving global problems and using built–

in cooperation strategies. In contrast, MAS work deals with heterogeneous, not necessarily centrally designed agents faced with the goal of utility–maximizing coexistence.

Examples of DPS work include Decker & Lesser (1993), addressing the task of fast coordination and reorganization of agents on a distributed sensor network, and Hogg & Williams (1993) showing how parallel search performs better with distributed cooperative agents than with independent agents. Examples of MAS work include Ephrati (1992), describing a master–slave scenario between two agents with essentially the same goals, and Miceli & Cesta (1993), using estimates of usefulness of social interactions for agents to select whom to interact with. Along similar lines, Kraus (1993) studies negotiations and contracts between selfish agents and Durfee, Lee & Gmytrasiewicz (1993) discuss game–theoretic and AI approaches to deals among rational agents.

Certain aspects of DAI work are purely theoretical and deal with the difficulty of multi–agent planning and control in abstract environments (e.g. Shoham & Tennenholtz (1992)). Some DAI work draws heavily from mathematical results in the field of parallel distributed systems (e.g. Huberman (1990), Clearwater, Huberman & Hogg (1991), and many others). DAI and Alife merge in the experimental mathematics field that studies computational ecosystems, using simulations of populations of agents with well defined interactions. The research is focused on global effects and the changes in the system as a whole over time. This process of global changes is usually referred to as "co–evolution" (Kephart, Hogg & Huberman 1990). Co–evolution experiments are usually used to find improved search-based optimization techniques (Hillis 1990). Often the systems studied have some similarities to the global effects found in biological ecosystems, but the complex details of biological systems cannot be reasonably addressed.

3 Key Terms and Definitions

Previous section offered a glimpse at the highly varied directions and approaches to studying multi–robot and multi–agent systems. One of the main hurdles in the way of cross–fertilization between research directions is inconsistent vocabulary. This section defines and overviews some of the key terms in order to make the described research accessible.

3.1 Behaviors and Goals

In the last few years the notion of "behavior" as a fundamental building block has been popularized in the AI, control, and learning communities. We define *behavior* to be a control law for reaching and/or maintaining a particular goal. For example, in the robot domain, *following* is a control law that takes inputs from an agent's sensors and uses them to generate actions which will keep the agent moving within a fixed region behind another moving object. This definition specifies that a behavior is a type of an operator that guarantees a particular goal, whatever its type. *Attainment goals* are terminal states; having reached a goal, the agent is finished. Such goals include reaching a home region and picking up an object. *Maintenance goals* persist in time, and are not always representable with terminal states, but rather with dynamic equilibria to be maintained. Examples include avoiding obstacles and minimizing interference. Situated agents can have multiple concurrent goals, including at least one attainment goal, and one or more maintenance goals.

In the scope of our work, *interaction* is mutual influence on behavior, and *ensemble, collective* or *group behavior* is an observer–defined temporal pattern of interactions between multiple agents. Of the innumerably many possible such behaviors for a given domain, only a small subset is relevant and desirable for achieving the agents' goals.

3.2 Communication and Cooperation

Communication and cooperation have become popular topics in both abstract and applied multi–agent work (Yanco & Stein 1993, Dudek et al. 1993, Altenburg 1994). Communication is the most common means of interaction among intelligent agents. Since any observable behavior and its consequences can be interpreted as a form of communication, we propose a stricter classification.

Direct communication is a purely communicative act, one with the sole purpose of transmitting information, such as a speech act, or a transmission of a radio message. More specifically, *directed communication* is direct communication aimed at a particular receiver. Such communication can be one–to–one or one–to–many, in all cases to identified receivers. In contrast, *indirect communication* is based on the observed behavior, not communication, of other agents. This type of communication is referred to as *stigmergic* in biological literature, where it refers to communication based on modifications of the environment rather than direct message passing.

Cooperation is a form of interaction, usually based on some form of communication. Certain types of cooperative behavior depend on directed communication. Specifically, any cooperative behaviors that require negotiation between agents depend on directed communication in order to assign particular tasks to the participants. Analogously to communication, *explicit cooperation* is defined as a set of interactions which involve exchanging information or performing actions in order to benefit another agent. In contrast, *implicit cooperation* consists of actions that are a part of the agent's own goal–achieving behavior repertoire, but have effects in the world that help other agents achieve their goals.

3.3 Interference and Conflict

All approaches to multi–agent control must deal with *interference*, any influence that opposes or blocks an agents' goal–driven behavior. In societies consisting of agents with identical goals, interference manifests itself as competition for shared resources. In diverse societies, where agents' goals differ, more complex conflicts can arise, including goal clobbering, deadlocks, and oscillations.

Two functionally distinct types of interference are relevant to this work: interference caused by multiplicity, called *resource competition*, and interference caused by goal–related conflict, called *goal competition*.

Resource competition includes any interference resulting from multiple agents competing for common resources, such as space, information, or objects. As the size of the group grows, this type of interference increases, causing the decline in global performance, and presenting an impetus for social rules.

Resource competition manifests itself in homogeneous and heterogeneous groups of coexisting agents. In contrast, goal competition arises between agents with different goals. Such agents may have identical high–level goals (such as, for example, a family has), but individuals may pursue different and potentially interfering subgoals i.e., they can be "functionally heterogeneous." Such heterogeneity does not arise in SIMD–style groups of functionally identical agents in which all are executing exactly the same program at each point in time.

Goal competition is studied primarily by the Distributed AI community (Gasser & Huhns 1989). It usually involves predicting other agents' goals and intentions, thus requiring agents to maintain models of each other (e.g., Huber & Durfee (1993) and Miceli & Cesta (1993)). However, such prediction abilities require computational resources that do not scale well with increased group sizes.

One means of simplifying prediction is through the use of social rules which attempt to eliminate or at least minimize both resource and goal competition. In particular, their purpose is to direct behavior away from individual greediness and toward global efficiency. In certain groups and tasks, agents must give up individual optimality in favor of collective efficiency. In those cases, greedy individualistic strategies perform poorly in group situations because resource competition grows with the size of the group.

Since social rules are designed for optimizing global resources, it is in the interest of each of the individuals to obey them. However, since the connection between individual and collective benefit is rarely direct, societies can harbor deserters who disobey social rules in favor of individual benefit. Game theory offers elaborate studies of the effects of deserters on individual optimality (Axelrod 1984), but domains treated in game theory are much more cleanly constrained than those treated here. In particular, game theory deals with rational agents capable of evaluating the utility of their actions and strategies. In contrast, our work is concerned with situated domains where the agents cannot be assumed to be rational due to incomplete or nonexistent world models and models of other agents, inconsistent reinforcement, and noise and uncertainty.

The goal of our work is to provide methodologies for finding efficient approaches to a variety of related problems. Optimality criteria for agents situated in physical worlds and maintaining long–term achievement and maintenance goals are difficult to characterize and even more difficult to achieve. While in game theory interference is a part of a competing agent's predictable strategy, in the embodied multi–agent domain interference is largely a result of direct resource competition, which can be moderated with relatively simple social rules.

4 Approaches to Multi–Agent Control

The problem of *multi–agent control* can be viewed at the individual agent level and at the collective level. The two levels are interdependent and the design of one is, or should be, strongly influenced by the other. However, multi–agent control has grown out of individual agent control, and this history is often reflected in the control strategies at the collective level. Individual agent control strategies can be classified into reactive, behavior–based, planner–based, and hybrid approaches (see Matarić (1994*a*) and Matarić (1992) for detailed comparisons and discussion).

Extending the planning paradigm from single–agent to multi–agent domains requires expanding the global state space to include the state of each of the agents. Such a global state space is exponential in the number of agents. Specifically, the size of the global state space G is: $|G| = s^a$ where s is the size of the state space of each agent, here assumed to be equal for all agents, or at worst the maximum for all agents, and a is the number of agents. Exponential growth of the state space makes the problem of global on–line planning intractable for all but the smallest group sizes, unless control is synchronized and has SIMD form, i.e. all agents perform the same behavior at the same time. Furthermore, since global planning requires communication between the agents and the controller, the bandwidth can grow with the number of agents. Additionally, the uncertainty in perceiving state grows with the increased complexity of the environment. Consequently, global planner–based approaches to control do not appear well suited for problems involving multiple agents acting in real–time based on uncertain sensory information.

Since hybrid systems typically employ a planner at the high level, in terms of multi–agent extensions they can be classified into the planner–based category. The collective behavior of a hybrid system would generally be a result of a plan produced by a global controller and distributed over independent possibly partially autonomous modules.

At the other end of the control spectrum, extending the reactive and behavior–based approaches to multi–agent domain results in completely distributed systems with no centralized controller. The systems are identical at the local and global levels: at the global level the systems are a collection of reactive agents each executing task–related rules relying only on local sensing and communication. Since all control in such distributed systems is local, it scales well with the number of agents, does not require global communication, and is more robust to sensor and effector errors. However, global consequences of local interactions between agents are difficult to predict. Thus, centralized approaches have the advantage of potential theoretical analysis while parallel distributed systems typically do not lend themselves to traditional analytical procedure.

4.1 Analysis of Behavior

Multi–agent systems are typically complex, either because they are composed of a large number of elements, or because the inter–element interactions are not simple. Systems of several situated agents with uncertain sensors and effectors display both types of complexity. This section addresses how these properties affect their behavior and its analysis.

The exact behavior of an agent situated in a nondeterministic world, subject to real error and noise, and using even the simplest of algorithms, is impossible to predict exactly. By induction, the exact behavior of each part of a multi–agent system of such nature is also unpredictable. A group of interacting agents is a dynamical system whose behavior is determined by the local interactions between individuals. In natural systems, such interactions result in the evolution of complex and stable behaviors that are difficult to analyze using traditional, top–down approaches. We postulate that in order to reach that level of complexity synthetically, such behaviors must be generated through a similar, interaction–driven, incrementally refined process.

Precise analysis and prediction of the behavior of a single situated agent, specifically, a mobile robot in the physical world, is an unsolved problem in robotics and AI. Previous work has shown that synthesis and analysis of correct plans in the presence of uncertainty can be intractable even in highly constrained domains (Lozano-Pérez, Mason & Taylor 1984, Canny 1988, Erdmann 1989) and even on the simplest of systems (Smithers 1994). Physical environments pose a great challenge as they usually do not contain the structure, determinism, and thus predictability usually required for formal analysis (Brooks 1991). Predicting the behavior of a multi–agent system is more complex than the single–agent case. The difficulty in analyzing comes from two properties intrinsic to complex systems:

1. the actions of an agent depend on the states/actions of other agents,

2. the behavior of the system as a whole is determined by the interactions between the agents rather than by individual behavior.

In general, no mathematical tools are available for predicting the behavior of a system with several, but not numerous, relatively complex interacting components, namely a collection of situated agents. In contrast to physical particle systems, which consist of large numbers of simple elements, multi–agent systems in nature and AI are defined by comparatively small groups of much more complex agents. Statistical methods used for analyzing particle systems do not directly apply as they require minimal interactions between the components (Weisbuch 1991, Wiggins 1990).

The difficulty in analyzing complex multi–agent systems lies in the level of system description. Descriptions used for control are usually low level, detailed, and continuous. In contrast, planning and analysis are usually done at a high level, often using an abstract, discrete model. A more desirable and manageable level may lie in between those two.

Instead of attempting to analyze arbitrary complex behaviors, our work focuses on providing a set of behavior primitives that can be used for synthesizing and analyzing a particular type of complex multi–agent systems. The primitives provide a programming language for designing analyzable control programs and resulting group behaviors. We describe the approach in the next section.

5 The Basis Behavior Approach

Our work is based on the belief that intelligent collective behavior in a decentralized system results from *local interactions* based on simple rules. **Basic behaviors** are proposed as a methodology for structuring those rules through a principled process of synthesis and evaluation. A behavior is a control law that achieves and maintains some goal.

We postulate that, for each domain, a set of behaviors can be found that are **basic** in that they are required for generating other behaviors, as well as being a minimal set the agent needs to reach its goal repertoire. The process of choosing the set of basic behaviors for a domain is dually constrained. From the bottom up, the process is constrained by the agent and environment dynamics. From the top down, the process is constrained by the repertoire of the agent's goals.

An effective set of basic behaviors in the spatial domain should enable the agents to employ a variety of flex-

Figure 1: The mobile robots used to demonstrate and verify our group behavior and learning work. These robots demonstrated group avoidance, following, aggregation, dispersion, flocking, wandering, foraging, docking, and learning to forage.

ible strategies for puck manipulation, collection, and distribution. The effectiveness of such strategies depends on maximizing synergy between agents: achieving the necessary goals while minimizing inter–agent interference.

The following set of basic behaviors is proposed: *avoidance*, *following*, *aggregation*, *dispersion*, *homing*, and *wandering*. According to our definition, the above behavior set is *minimal* or *basic* in that its members are not further reducible to each other, and they are sufficient for achieving our set of pre–specified goals. A number of other utility behaviors can be a part of an agent's repertoire, such as *grasping* and *dropping*, the only other behaviors used in this work.

The basic behavior set is evaluated by giving a formal specification of each of the behaviors, and comparing the collection of those specifications to a formal specification of the set of global tasks required for the group. We have provided detailed specifications and algorithms for each of the basic behaviors, implemented and evaluated them based on the following criteria: *repeatability*, *stability*, *robustness*, and *scalability*.

We applied the above criteria to the data obtained by running a large number of trials (at least 50) of each basic behavior on a collection of over 20 physical mobile robots equipped with local sensors and local control (Figure 1). Each of the robots is equipped with a suite of infra–red sensors for collision avoidance and puck detection, micro switches and bump sensors for contact detection, and radios and sonars for triangulating their position relative to two stationary beacons, and broadcasting word–sized messages within a limited radius. The basic behaviors, each consisting of one rule or a small set of simple rules, generated robust group behaviors that met the prespecified evaluation criteria. The top row of Figure 2 shows a small subset of the robot data. The robots, physically 12 inches long, are scaled down and plotted as black rectangles, with white arrows indicating their heading. The the row of rectangles at the bottom shows the robots that were used in the experiment. Boxes on the lower right indicate frame numbers and the elapsed time in seconds for each of the runs.

Basic behaviors are intended as building blocks for achieving higher–level goals and can be embedded in an architecture that allows two types of combination: direct (by summation) and temporal (by switching). Direct combination execute multiple behaviors concurrently and combine their outputs. In contrast, temporal combinations execute only one behavior at a time. The architecture allows for multiple applications of the two operators to basic behavior subsets. To demonstrate the combination operators, we implemented two higher–level behaviors, flocking and foraging, as illustrated in the bottom row of Figure 2.

We generated simple and robust *flocking* behavior by summing the outputs of *avoidance*, *aggregation* and *wandering*. When *homing* was added, the flock could direct itself toward a particular goal location. In all cases, the flocks had no fixed leaders, and were not vulnerable to failures of individual robots.

A more complex example of a high–level behavior we demonstrated is *foraging*. It was implemented by applying a temporal combination operator to switch between *avoidance*, *dispersion*, *following*, *homing*, and *wandering* under appropriate sensory conditions. Those basic behaviors and the rudimentary abilities to pick up and drop pucks were sufficient to produce a robust and flexible collective foraging behavior that consisted of collecting all of the pucks in the area and depositing them in the home region while avoiding collisions and minimizing interference between the robots.

In addition to empirical testing of the behaviors and their combinations, we compared our methodology to a centralized, "total knowledge" approach. The experimental results showed that the simple, fully distributed strategies, applied to dispersion and aggregation tasks, converged only a constant factor slower than the centralized approach. A detailed description of the basic behavior approach as well as all of the related experiments and data can be found in Matarić (1994*a*).

6 Learning in Complex Group Environments

In the previous section we described our theoretical and experimental demonstration of basic behavior selection and combination into a large repertoire of higher–level competencies. This section briefly overviews our extension of the basic behavior approach to include a methodology for automatically generating higher–level behaviors. Such behaviors can be generated by having the agents learn through their interactions with the world and with other agents, i.e. though unsupervised **reinforcement learning** (RL).

RL has been successfully applied to a variety of domains where the agent–environment interaction can be described as a Markov Decision Process (MDP). However, that assumption does not directly apply to the stochastic, noisy, and uncertain multi–agent environments addressed in this work. We implemented a reformulation of the traditional RL model consisting of states, actions, and reinforcement in order to make it applicable to our domain. Instead of using actions, our system learns at the level of *basic behaviors* that hide low–level control details, and are more general and robust. After actions are replaced with behaviors, states can be replaced with *conditions*, the necessary and sufficient subsets of state required for triggering the behavior set. Conditions are many fewer than states, so their use diminishes the agent's learning space and speeds up any RL algorithm (Matarić 1994*c*).

In addition to the use of behaviors and conditions, we introduced two ways of shaping the reinforcement function to aid the agent–learner in the nondeterministic, noisy, and dynamic environment. We used *heterogeneous reward functions*, which partition the task into subgoals, thus providing more immediate reinforcement. Within a single behavior (i.e., a single goal), *progress estimators* we introduced as functions associated with particular conditions that provided some metric of the learner's performance. Progress estimators decrease the learner's sensitivity to noise, minimize thrashing, and minimize fortuitous rewards by correlating some domain knowledge about progress with appropriate behaviors the agent has taken in the past.

The proposed RL formulation was validated on the task of learning to forage. The behavior space included the foraging subset of basic behaviors described above, augmented with *grasping*, *dropping*, and *resting* (an opportunity for the robots to recharge). The state space was effectively reduced to the power set of the following conditions: *have-puck?*, *at-home?*, *night-time?*, and *near-intruder?*.

We implemented different versions of learning algorithms in our domain and compared their performance over a large number of trials. The popular standard RL Q-learning was implemented and used as a control case, and compared to an algorithm using heterogeneous reward functions, and to one using them in addition to progress estimators. Our approach outperformed the alternatives, and the analysis of the data yielded clean intuition about the varying difficulty levels even within the lifetime of a single foraging task. The use of progress estimators was found to be necessary, and only two such estimators were sufficient for the learning system to consistently converge within 15 minutes. For a detailed description of the learning algorithms and the data see

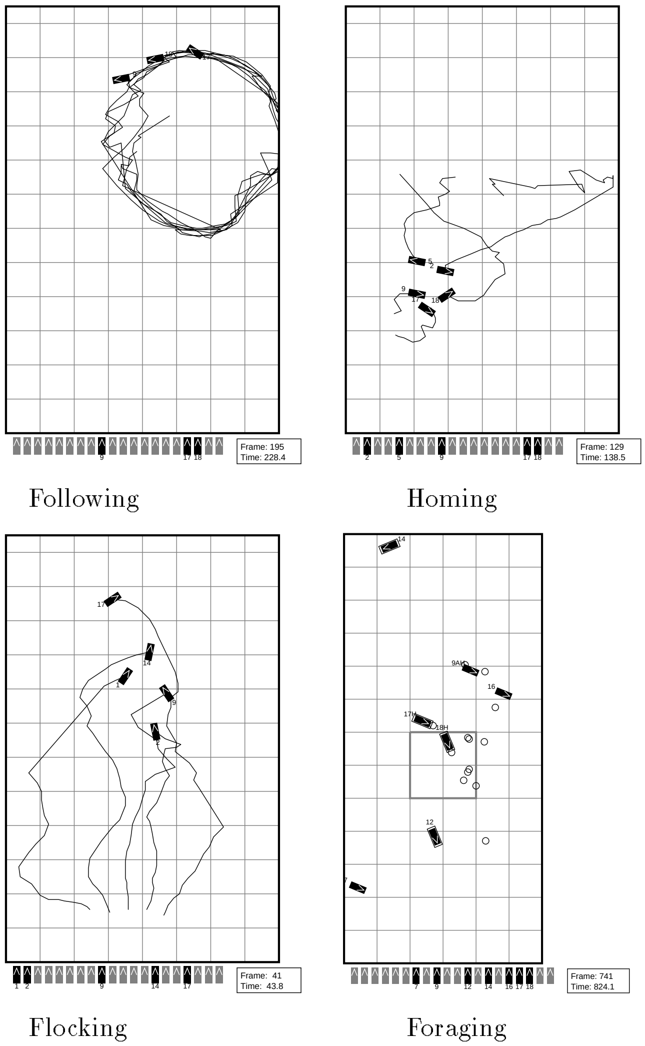

Figure 2: This figure shows real robot data for basic and composite behaviors. Small boxes on the right indicate the elapsed time in seconds for each of the runs. The top row shows following and homing data. The second row shows robot data for the composite behaviors: flocking and foraging. The foraging behavior of 7 robots is shown after 13.7 minutes of running; collected pucks are in the box.

Matarić (1994*c*).

7 Summary

This paper has reviewed the key terms, issues, and approaches in multi–robot and situated multi–agent control. We described the challenges of principled synthesis and analysis of collective behavior. Furthermore, we proposed a methodology for structuring the process of designing group behaviors for multi–robot systems.

The basic behavior approach is general and biologically rooted (Matarić 1994*a*). Thus we believe that it is applicable to various domains of multi–agent interaction featuring complex dynamics, unpredictability, and uncertainly in sensing and action. The methodology is invariant to group size and interaction type. We have so far demonstrated it on up to 20 agents situated in the spatial domain, and have also applied it to smaller groups of more heterogeneous agents (Matarić 1994*b*, Matarić, Nilsson & Simsarian 1994), and intend to test it on even larger group sizes and more abstract domains.

This work is intended as a foundation in a continuing effort toward studying and synthesizing increasingly more complex behavior. The work on basic behaviors distills a general approach to control, planning, and learning. The work also empirically demonstrates some challenging problems and offers some effective solutions to group behavior and learning.

Acknowledgements

The research reported here was done at the MIT Artificial Intelligence Laboratory, and supported 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

Altenburg, K. (1994), Adaptive Resource Allocation for a Multiple Mobile Robot System using Communication, Technical Report NDSU–CSOR–TR–9404, North Dakota State Univeristy.

Arkin, R. C. (1992), 'Cooperation without Communication: Multiagent Schema Based Robot Navigation', *Journal of Robotic Systems*.

Arkin, R. C., Balch, T. & Nitz, E. (1993), Communication of Behavioral State in Multi-agent Retrieval Tasks, *in* 'IEEE International Conference on Robotics and Automation', pp. 588–594.

Assad, A. & Packard, N. (1992), Emergent Colonization in an Artificial Ecology, *in* F. Varela & P. Bourgine, eds, 'Toward A Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life', The MIT Press, pp. 143–152.

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

Barman, R. A., Kingdon, J. J., Mackworth, A. K., Pai, D. K., Sahota, M. K., Wilkinson, H. & Zhang, Y. (1993), Dynamite: A Testbed for Multiple Mobile Robots, *in* 'Proceedings, IJCAI-93 Workshop on Dynamically Interacting Robots', Chambery, France, pp. 38–44.

Beckers, R., Holland, O. E. & Deneubourg, J. L. (1994), From Local Actions to Global Tasks: Stigmergy and Collective Robotics, *in* R. Brooks & P. Maes, eds, 'Artificial Life IV, Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems', The MIT Press.

Beni, G. & Hackwood, S. (1992), The Maximum Entropy Principle and Sensing in Swarm Intelligence, *in* F. Varela & P. Bourgine, eds, 'Toward A Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life', The MIT Press, pp. 153–160.

Brock, D. L., Montana, D. J. & Ceranowicz, A. Z. (1992), Coordination and Control of Multiple Autonomous Vehicles, *in* 'IEEE International Conference on Robotics and Automation', pp. 2725–2730.

Brooks, R. A. (1991), Intelligence Without Reason, *in* 'Proceedings, IJCAI-91'.

Brooks, R. A., Maes, P., Matarić, M. J. & Moore, G. (1990), Lunar Base Construction Robots, *in* 'IEEE International Workshop on Intelligent Robots and Systems (IROS-90)', Tokyo, pp. 389–392.

Caloud, P., Choi, W., Latombe, J., LePape, C. & Yim, M. (1990), Indoor Automation with Many Mobile Robots, *in* 'IROS-90', Tsuchiura, Japan, pp. 67–72.

Canny, J. F. (1988), *The Complexity of Robot Motion Planning*, The MIT Press, Cambridge, Massachusetts.

Clearwater, S. H., Huberman, B. A. & Hogg, T. (1991), 'Cooperative Solution of Constraint Satisfaction Problems', *Science* **254**, 1181–1183.

Colorni, A., Dorigo, M. & Maniezzo, V. (1992), Distributed Optimization by Ant Colonies, *in* F. Varela & P. Bourgine, eds, 'Toward A Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life', The MIT Press, pp. 134–142.

Corbara, B., Drogoul, A., Fresneau, D. & Lalande, S. (1993), Simulating the Sociogenesis Process in Ant Colonies with MANTA, *in* 'Toward A Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life', pp. 224–235.

Dallas, J. (1990), Co-operative Search Behavior in a Group of Lego Robots, Master's thesis, University of Edinburgh.

Dario, P. & Rucci, M. (1993), An Approach to Disassembly Problem in Robotics, *in* 'IEEE/TSJ International Conference on Intelligent Robots and Systems', Yokohama, Japan, pp. 460–468.

Dario, P., Ribechini, F., Genovese, V. & Sandini, G. (1991), Instinctive Behaviors and Personalities in Societies of Cellular Robots, *in* 'IEEE International Conference on Robotics and Automation', pp. 1927–1932.

Decker, K. & Lesser, V. (1993), A One-shot Dynamics Coordination Algorithm for Distributed Sensor Networks, *in* 'Proceedings, AAAI-93', Washington, DC, pp. 210–216.

Deneubourg, J. L., Goss, S., Franks, N., Sendova-Franks, A., Detrain, C. & Chretien, L. (1990), The Dynamics of Collective Sorting, *in* 'From Animals to Animats: International Conference on Simulation of Adaptive Behavior', The MIT Press, pp. 356–363.

Deneubourg, J. L., Theraulax, G. & Beckers, R. (1992), Swarm-Made Architectures, *in* F. Varela & P. Bourgine, eds, 'Toward A Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life', The MIT Press, pp. 123–133.

Donald, B. R., Jennings, J. & Rus, D. (1993), Experimental Information Invariants for Cooperating Autonomous Mobile Robots, *in* 'Proceedings, International Symposium on Robotics Research', Hidden valley, PA.

Drogous, A., Ferber, J., Corbara, B. & Fresneau, D. (1992), A Behavioral Simulation Model for the Study of Emergent Social Structures, *in* F. Varela & P. Bourgine, eds, 'Toward A Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life', The MIT Press, pp. 161–170.

Dudek, G., Jenkin, M., Milios, E. & Wilkes, D. (1993), A Taxonomy for Swarm Robotics, *in* 'IEEE/TSJ International Conference on Intelligent Robots and Systems', Yokohama, Japan, pp. 441–447.

Durfee, E. H., Lee, J. & Gmytrasiewicz, P. J. (1993), Overeager Reciprocal Rationality and Mixed Strategy Equilibria, *in* 'Proceedings, AAAI–93', Washington, DC, pp. 225–230.

Ephrati, E. (1992), Constrained Intelligent Action: Planning Under the Influence of a Master Agent, *in* 'Proceedings, AAAI-92', San Jose, California, pp. 263–268.

Erdmann, M. (1989), On Probabilistic Strategies for Robot Tasks, PhD thesis, MIT.

Fukuda, T., Nadagawa, S., Kawauchi, Y. & Buss, M. (1989), Structure Decision for Self Organizing Robots Based on Cell Structures - CEBOT, *in* 'IEEE International Conference on Robotics and Automation', Scottsdale, Arizona, pp. 695–700.

Fukuda, T., Sekiyama, K., Ueyama, T. & Arai, F. (1993), Efficient Communication Method in the Cellular Robotics System, *in* 'IEEE/TSJ International Conference on Intelligent Robots and Systems', Yokohama, Japan, pp. 1091–1096.

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

Hillis, W. D. (1990), 'Co-evolving Parasites Improve Simulated Evolution as an Optimization Procedure', *Physica D* **42**, 228–234.

Hodgins, J. K. & Brogan, D. C. (1994), Robot Herds: Group Behaviors from Systems with Significant Dynamics, *in* R. Brooks & P. Maes, eds, 'Artificial Life IV, Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems', The MIT Press.

Hogeweg, P. & Hesper, B. (1985), 'Socioinformatic Processes: MIRROR Modelling Methodology', *Journal of Theoretical Biology* **113**, 311–330.

Hogg, T. & Williams, C. P. (1993), Solving the Really Hard Problems with Cooperative Search, *in* 'Proceedings, AAAI-93', Washington, DC, pp. 231–236.

Huang, Q. & Beni, G. (1993), Stationary Waves in 2–Dimensional Cyclic Swarms, *in* 'IEEE/TSJ International Conference on Intelligent Robots and Systems', Yokohama, Japan, pp. 433–440.

Huber, M. J. & Durfee, E. H. (1993), Observational Uncertainty in Plan Recognition Among Interacting Robots, *in* 'Proceedings, IJCAI-93 Workshop on Dynamically Interacting Robots', Chambery, France, pp. 68–75.

Huberman, B. A. (1990), 'The Performance of Cooperative Processes', *Physica D* **42**, 38–47.

Kephart, J. O., Hogg, T. & Huberman, B. A. (1990), 'Collective Behavior of Predictive Agents', *Physica D* **42**, 48–65.

Keshet, L. E. (1993), Trail Following as an Adaptable Mechanism for Population Behaviour, *in* 'Toward A Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life', pp. 326–346.

Kraus, S. (1993), Agents Contracting Tasks in Non-Collaborative Environments, *in* 'Proceedings, AAAI-93', Washington, DC, pp. 243–248.

Kube, C. R. & Zhang, H. (1992), Collective Robotic Intelligence, *in* 'From Animals to Animats: International Conference on Simulation of Adaptive Behavior', pp. 460–468.

Kube, C. R., Zhang, H. & Wang, X. (1993), Controlling Collective Tasks With an ALN, *in* 'IEEE/TSJ International Conference on Intelligent Robots and Systems', Yokohama, Japan, pp. 289–293.

Kurosu, K., Furuya, T. & Soeda, M. (1993), Fuzzy Control of Group With a Leader and Their Behaviors, *in* 'IEEE/TSJ International Conference on Intelligent Robots and Systems', Yokohama, Japan, pp. 1105–1109.

Lozano-Pérez, T., Mason, M. T. & Taylor, R. H. (1984), 'Automatic Synthesis of Fine Motion Strategies for Robots', *International Journal of Robotics Research* **3**(1), 3–24.

MacLennan, B. J. (1990), Evolution of Communication in a Population of Simple Machines, Technical Report Computer Science Department Technical Report CS–90–104, University of Tennessee.

Matarić, M. J. (1992), Behavior-Based Systems: Key Properties and Implications, *in* 'IEEE International Conference on Robotics and Automation, Workshop on Architectures for Intelligent Control Systems', Nice, France, pp. 46–54.

Matarić, M. J. (1994*a*), Interaction and Intelligent Behavior, Technical Report AI-TR-1495, MIT Artificial Intelligence Lab.

Matarić, M. J. (1994*b*), Learning to Behave Socially, *in* D. Cliff, P. Husbands, J.-A. Meyer & S. Wilson, eds, 'From Animals to Animats: International Conference on Simulation of Adaptive Behavior', pp. 453–462.

Matarić, M. J. (1994*c*), Reward Functions for Accelerated Learning, *in* W. W. Cohen & H. Hirsh, eds, 'Proceedings of the Eleventh International Conference on Machine Learning (ML-94)', Morgan Kauffman Publishers, Inc., New Brunswick, New Jersey, pp. 181–189.

Matarić, M. J., Nilsson, M. & Simsarian, K. T. (1994), Cooperative Multi–Robot Box–Pushing, *in* 'submitted to IROS-95', Pittsburgh, PA.

Miceli, M. & Cesta, A. (1993), Strategic Social Planning: Looking for Willingness in Multi-Agent Domains, *in* 'Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society', Boulder, Colorado, pp. 741–746.

Noreils, F. R. (1993), 'Toward a Robot Architecture Integrating Cooperation between Mobile Robots: Application to Indoor Environment', *The International Journal of Robotics Research* **12**(1), 79–98.

Parker, L. E. (1993*a*), An Experiment in Mobile Robotic Cooperation, *in* 'Proceedings, Robotics for Challenging Environment', Albuquerque, New Mexico.

Parker, L. E. (1993*b*), Learning in Cooperative Robot Teams, *in* 'Proceedings, IJCAI-93 Workshop on Dynamically Interacting Robots', Chambery, France, pp. 12–23.

Rosenschein, J. S. (1993), Consenting Agents: Negotiation Mechanisms for Multi-Agent Systems, *in* 'IJCAI-93', pp. 792–799.

Sandini, G., Lucarini, G. & Varoli, M. (1993), Gradient Driven Self–Organizing Systems, *in* 'IEEE/TSJ International Conference on Intelligent Robots and Systems', Yokohama, Japan, pp. 429–432.

Schmieder, R. W. (1993), A Knowledge-Tracking Algorithm for Generating Collective Behavior in Individual-Based Populations, *in* 'Toward A Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life', pp. 980–989.

Shoham, Y. & Tennenholtz, M. (1992), On the synthesis of useful social laws for artificial agent societies, *in* 'Proceedings, AAAI-92', San Jose, California, pp. 276–281.

Smithers, T. (1994), On Why Better Robots Make it Harder, *in* 'The Third International Conference on Simulation of Adaptive Behavior'.

Steels, L. (1989), Cooperation Between Distributed Agents Through Self-Organization, *in* 'Workshop on Multi-Agent Cooperation', North Holland, Cambridge, UK.

Travers, M. (1988), Animal Construction Kits, *in* C. Langton, ed., 'Artificial Life', Addison–Wesley.

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

Werner, G. M. & Dyer, M. G. (1990), Evolution of Communication in Artificial Organisms, Technical Report UCLA–AI–90–06, University of California, Los Angeles.

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

Yanco, H. & Stein, L. A. (1993), An Adaptive Communication Protocol for Cooperating Mobile Robots, *in* 'From Animals to Animats: International Conference on Simulation of Adaptive Behavior', The MIT Press, pp. 478–485.