From Local Interactions to Collective Intelligence

Maja J Matarić
MIT Artificial Intelligence Laboratory
545 Technology Square #721
Cambridge, MA 02139
maja@ai.mit.edu

Abstract

We describe an approach to synthesizing and analyzing intelligent group behavior. The goal of our research is to understand the types of simple local interactions which produce complex and purposive group behaviors by taking a synthetic, bottom–up approach. We propose *basic behaviors* as the basis for control and learning. These behaviors serve as a substrate for more complex interactions and structure and simplify the process of both designing and analyzing group behaviors. Finally, they are an efficient basis for learning behavior selection. To illustrate these ideas, we present a behavior repertoire developed and tested on a herd of physical mobile robots.

1 Introduction and Motivation

One of the main goals of Artificial Intelligence (AI) is to gain insight into natural intelligence through a synthetic approach, by generating and analyzing artificial intelligent behavior. In order to glean an understanding of a phenomenon as complex as natural intelligence, we need to study complex behavior in complex environments.

Traditionally, AI has concerned itself with complex agents in relatively simple environments, simple in the sense that they could be precisely modeled and involved little or no noise and uncertainty. In contrast to traditional systems, reactive and behavior–based systems have placed agents with low levels of cognitive complexity into complex, noisy and uncertain environments. Our work attempts to simultaneously scale up along both dimensions: by introducing other agents into the environment, and by introducing learning capabilities into each of the agents (Figure 1).

In order to be understood, individual intelligence must be observed and analyzed within its social and therefore cultural environment. In contrast to traditional AI, which addresses intelligence as an isolated phenomenon, this work is based on the belief that intelligent behavior is inextricably tied to its cultural context and cannot be understood in isolation. This emphasis is similar to the founding principles of ethology, the study of animal behavior in its natural habitat.

The complexity of culture results from the local interactions among individuals. Our work focuses on exploring simple local social interactions which result in purposive group behaviors. Nature abounds with systems whose complex global behaviors emerge from simple local interactions. Such behaviors can be found in many domains and at all scales, from the subatomic (Gutzwiller 1992), to the semantic (Minsky 1986), to the social (Deneubourg, Goss, Franks, Sendova-Franks, Detrain & Chretien 1990). In contrast to these bottom–up effects, most of AI relies on top–down modularity. As expressed by Simon (1969), a system is well designed if it is analyzable through *decomposition* into non–interacting modules.

However, imposing such modularity minimizes precisely the type of interactions that seem to generate complexity in nature. The global behavior of complex systems, such as groups of social agents, is determined by the local interactions of their constituent parts. These interactions merit careful study in order to understand the global behavior of the society. In natural systems, such interactions resulted in the evolution of complex and stable behaviors that do not lend themselves to

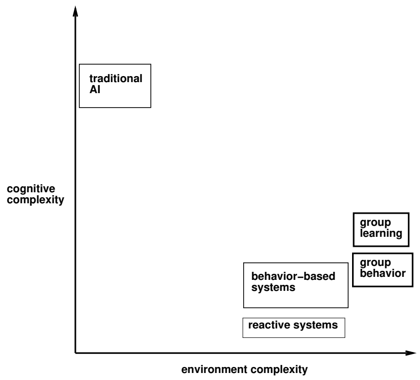

Figure 1: Traditional AI has addressed complex agents in simple environments while reactive and behavior–based approaches have dealt with simple agents in noisy and uncertain worlds. This work attempts to scale up along both dimensions simultaneously, by addressing synthesis and learning of complex group behavior.

traditional, top–down style of analysis. In order to reach that level of complexity synthetically, such behaviors must be generated through a similar, interaction–driven, incrementally refined process. We describe such an approach, a combination of bottom–up experiments and theory with the purpose of of designing, observing, and formalizing group behaviors.

1.1 Analyzing Complex Behavior

Analyzing and predicting the behavior of a single situated agent is an unsolved problem in robotics and AI. Even highly constrained domains are intractable, and realistic worlds usually do not contain the structure, determinism and predictability necessary for strict formal analysis (Canny 1988, Brooks 1990*b*, Brooks 1991).

Predicting the behavior of a multi–agent system is even more complex than the single–agent case. In general, no satisfactory solution exists for predicting the behavior of a system with nontrivial interacting components. Statistical methods used for analyzing particle systems do not apply as they require minimal interactions among the components (Weisbuch 1991). While systems with large numbers of simple and simply interacting components can be analyzed this way (Wiggins 1990), no tools are available for systems consisting of comparatively few but more complex components with more complex interactions.

While it is difficult to prove properties of the exact behavior of individual agents within a group, we believe that it is possible to evaluate and at least qualitatively predict the behavior of the ensemble as a whole. Instead of attempting to analyze arbitrary complex behaviors, this 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. We will describe a set of local interactions which produce cohesive collective behaviors in this domain, and can be combined into variety of more complex global behaviors. The behaviors we describe emerge from local interactions, but rather than being unexpected, are repeatable and well understood.

2 Basic Principles

The following principles form the basis for the methodology used in this work.

- Agents are homogeneous in terms of hardware and software. No *a priori* leaders or followers are designated.
- Agents do not use any explicit one–to–one communication. Instead, all communication is based only on sensing the external state of nearby agents, and general broadcast of messages that can be received by any of the neighbors.
- Agents do not engage in any explicit cooperation. Instead, cooperation is implicit, and occurs through the world rather than through directed communication.
- Agents have no hidden goals. This work does not deal with the agents' underlying motivations, the possibility of cheating, competition, and ulterior motives. Agents are assumed to have a set of common, similar, or at least overlapping goals.
- Agents are able to detect other agents of the same kind. The ability to categorize the perceptible objects in the world into at least two classes: "others like me" and "everything else" is a necessary condition for intelligent collective behavior. With this ability, which is innate and ubiquitous in nature, even the simplest of interactions can produce purposive collective behavior, as we will demonstrate in the following sections.

3 Synthesis and Analysis of Group Behavior

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 a goal. For example, *obstacle avoidance* is a behavior that maintains the goal of avoiding collisions with obstacles.

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 dynamics of the agent and the environment. From the top–down, the process is constrained by the agent's goals as specified by the task. The combination of the two types of constraints helps to prune the agent's behavior space.

We will use the example of group interactions between situated, embodied agents to illustrate the process of selecting a basic behavior set. The agents are mobile robots, embodied and endowed with specific mechanical, sensory, and effector constraints. We define the high–level goals of the system as consisting of collectively moving objects (pucks) in the environment in an efficient fashion. In this work, efficiency is defined in terms of minimizing energy by minimizing the amount of time required to complete a task or the number of moves required for each of the agents.

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

We propose the following set of basic behaviors:

- *safe–wandering* – minimizes collisions between agents and environment
- *following* – minimizes interference by structuring movement of any two agents
- *aggregation* – gathers the agents
- *dispersion* – dissipates the agents
- *homing* – enables the agent to proceed to a particular location

According to our definition, the above behavior set is *minimal* or *basic* in that its members are not further reducible to each other. The described basic behaviors are defined with respect to the group. Other utility behaviors, such as *grasping* and *dropping*, can also be a part of an agent's repertoire.

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 (Matarić 1994*a*).

Once a basic behavior set is established, it can be implemented with a variety of algorithms. The first step in the verification of basic behavior algorithms is a comparison between the formal behavior specification and the formal correctness of the algorithm. Using the notion of ensemble behavior, we propose group behavior algorithms that utilize a centroid operator that averages the inputs from multiple agents. This operator has statistical properties that allow analyzing and making predictions about the behavior of the group (Matarić 1992). The algorithms are also empirically evaluated based on the following criteria:

- repeatability: how consistent is the behavior over different trials?
- stability: does the behavior oscillate under any conditions?
- robustness: how robust is the behavior in the presence of sensor and effector error and noise?
- scalability: how is the behavior effected by increased and decreased group sizes?

4 Combining Interactions

Basic behaviors are intended as building blocks for achieving higher–level goals. The behaviors are embedded in an architecture that allows two types of combination: direct, by summation, and temporal, by switching (see figure 4).

The following are some of the compound behaviors that can be constructed with the above basis set.

- *flocking* – moves a group as a coherent aggregate without prespecified leaders and followers. A summing combination of *safe–wandering*, *dispersion*, *aggregation*, and *homing*.
- *docking* – parks a group along a boundary. A variation of *safe–wandering*.
- *surrounding* – positions a group of agents around the ouside boundary of the rest. A summing combinations of *following* and *aggregation*.
- *herding* – moves a group as a coherent aggregate guided by the agents on the periphery. A switching combination of *surrounding* and *flocking*.
- *foraging* – group searching and collection of objects. A switching combination of *safe–wandering*, *dispersion*, *following*, and *homing*

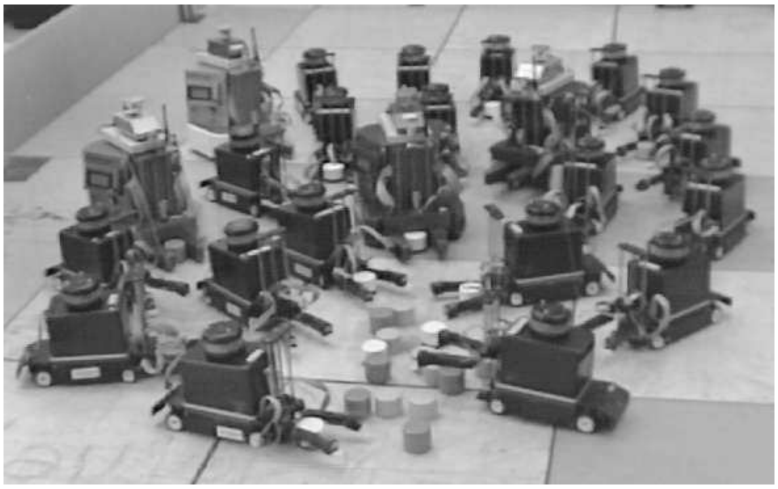

Figure 2: A family photo of the physical experimental agents used to demonstrate and verify our group behavior and learning work.

5 Experimental Methodology

Since complex interactive behavior depends strongly on the particular parameters of the system, in order to duplicate the behavior it is necessary to either construct a very accurate model of the system or to build the system and experimentally reproduce the behavior. Our work relies on the latter approach, a bottom–up experimental methodology, which drives and refines our theories. Although there is a great tendency to design behaviors around solving a problem, or simulating a naturally system, it is necessary to go beyond finding a working solution to evaluating the effects of various possibilities.

Rather than attempt to to simulate biological systems, the goal of this work is to elucidate the phenomena of social interaction by synthesizing, observing, and analyzing similar phenomena. Consequently, our evaluation criteria are based on performing and analyzing a variety of experiments in our synthetic domain.

We applied the described evaluation criteria were applied to the data obtained by performing at least 50 trials of each basic behavior. The experiments were performed on two different multi–agent environments, in order to minimize domain biases. The first environment was a multi–agent simulator featuring up to 50 agents with local sensing and distributed, local control. The second environment was a collection of 24 physical mobile robots equipped with local sensors and local control (Figure 2). Each of the robots is equipped with a suite of infra–red sensors for collision avoidance, puck detection, and stacking, and with micro switches and bump sensors for contact detection. In addition to the local sensors, the robots are equipped with radios and sonars for triangulating their position relative to two stationary beacons, and for broadcasting that position within a limited radius. The radios are used to detect other robots and gather data for local centroid computations.

The robot experiments are run fully autonomously with all of the processing and power on board. The robots are programmed in the Behavior Language, a parallel programming language based on the Subsumption Architecture (Brooks 1990*a*, Brooks 1987). Their control systems consist of a collection of parallel, concurrently active behaviors. Some of the behaviors implement basic

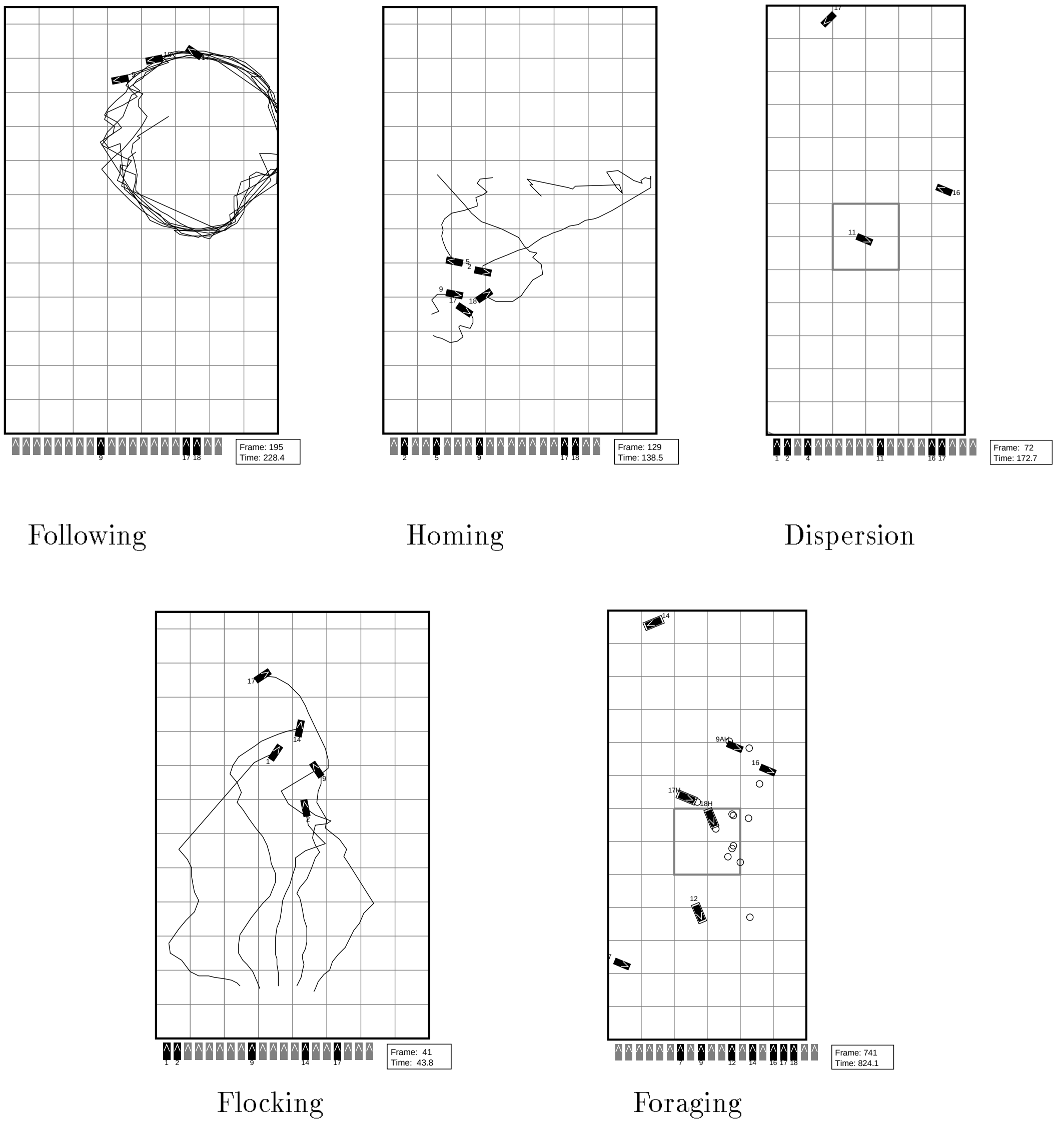

Figure 3: The first row shows examples of real robot data for three basic behaviors: *following*, *homing*, and *dispersion*. The robots, physically 12 inches long, are scaled down and plotted as black rectangles, with white arrows indicating their heading. The the row of rectangles at the bottom shows 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. The second row shows examples of robot data for two different composite group behaviors: *flocking* and *foraging*. *Foraging* behavior of 7 robots is shown after 13.7 minutes of running; collected pucks are in the box.

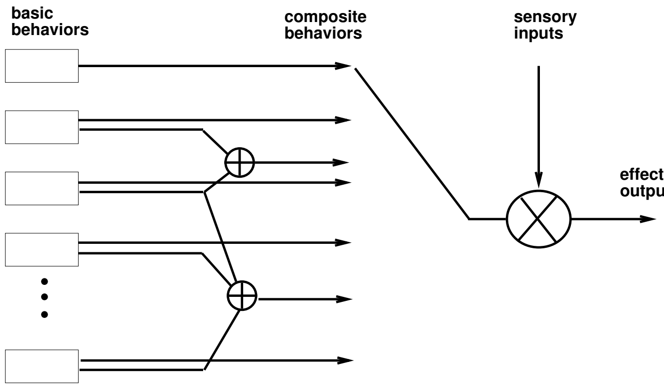

Figure 4: The control architecture for generating group behaviors consists of direct and temporal combinations (i.e. sums and switches) of subsets from a fixed basic behavior set. Direct combinations are marked with $\bigoplus$, temporal combinations with $\bigotimes$.

survival instincts while others implement more complex social interactions.

In addition to empirical testing of the behaviors and their combinations, the proposed methodology for generating decentralized group behavior was compared 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 (Matarić 1993).

A small subset of the robot data is shown in Figure 3, using the Real Time Viewer[1], a software package for displaying and replaying each of the robots runs, plotting their positions over time, and displaying each frame and the elapsed time for each experiment.

6 Learning in Complex Group Environments

We have also developed a methodology for automatically combining basic behaviors into higher–level ones, though unsupervised **reinforcement learning** based on the agents' interactions with the environment.

In reinforcement learning (RL) approaches the agent learns from external scalar reward and punishment. RL has been successfully applied to a variety of domains that have largely been modeled as Markovian, where the agent–environment interaction can be described as a Markov Decision Process (MDP). However, the MDP assumption does not directly apply to the noisy and uncertain multi–agent environments. Nonetheless, since external and internal feedback are the most natural sources of information for learning in situated agents, methods for applying RL to such complex domains are needed.

The traditional formulation of RL problems in terms of states, actions, and reinforcement required a reformulation in order to be applied to our domain. The notion of state as a monolithic descriptor of the agent and the environment scales poorly in multi–agent domains. The traditional notion of discrete actions is too low level and has effects too unpredictable and noisy to be useful to a situated learning algorithm. Finally, delayed reinforcement and reward discounting appear insufficient for learning in our domain.

[1]Written by Matthew Marjanović.

To make learning possible we proposed a reformulation that elevates the level of system description from states and actions to conditions and behaviors. *Behaviors* are control laws that achieve goals but hide low–level control details. Using the notion of *basic behaviors*, a small basis set can be defined as used as a substrate for learning. When 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, thus greatly diminishing the agent's learning space and speeding up any RL algorithm.

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

The proposed formulation was validated on the task of learning to associate the conditions and behaviors for the *foraging* task described above. Three learning algorithms were implemented and compared, including standard Q-learning. The performance of each of the three algorithms was averaged over 20 trials, and it clearly showed a significant increase in both viability and efficiency of situated learning of foraging when *heterogeneous reward functions* and *progress estimators* were introduced. Without them, Q failed to converge, while with them the robots learned to forage in under 15 minutes. The analysis of the learning performance and its general implications for reinforcement learning in situated, noisy, dynamic and uncertain domains is described in Matarić (1994*b*).

7 Summary

We have described a research program for studying intelligent group behavior with the goal of understanding the types of simple local interactions which produce complex and purposive group behaviors. We described a synthetic, bottom–up approach using *basic behaviors* that serve as efficient and effective substrates for control and learning.

As we continue research on this topic we hope that synthesizing and observing group behavior will offer insights into understanding social interactions in biology, as well as help derive methods for principled behavior control in robotics and AI.

Acknowledgements

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

Brooks, R. A. (1987), A Hardware Retargetable Distributed Layered Architecture for Mobile Robot Control, *in* 'IEEE International Conference on Robotics and Automation', Raleigh, NC, pp. 106–110.

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

Brooks, R. A. (1990*b*), Challenges for Complete Creature Architectures, *in* J. A. Meyer & S. Wilson, eds, 'From Animals to Animats: International Conference on Simulation of Adaptive Behavior', The MIT Press.

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

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

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.

Gutzwiller, M. (1992), 'Quantum Chaos', *Scientific American.*

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

Matarić, M. J. (1993), Kin Recognition, Similarity, and Group Behavior, *in* 'Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society', Boulder, Colorado, pp. 705–710.

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

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

Minsky, M. L. (1986), *The Society of Mind*, Simon and Schuster, New York.

Simon, H. (1969), *The Sciences of the Artificial*, The MIT Press.

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.