

Kin Recognition, Similarity, and Group Behavior

Maja J Mataric
MIT Artificial Intelligence Laboratory
545 Technology Square #721
Cambridge, MA 02139
phone: (617) 253-8839
fax: (617) 253-0039
maja@ai.mit.edu

Abstract

This paper presents an approach to describing group behavior using simple local interactions among individuals. We propose that for a given domain a set of basic interactions can be defined which describes a large variety of group behaviors. The methodology we present allows for simplified qualitative analysis of group behavior through the use of shared goals, kin recognition, and minimal communication. We also demonstrate how these basic interactions can be simply combined into more complex compound group behaviors.

To validate our approach we implemented an array of basic group behaviors in the domain of spatial interactions among homogeneous agents. We describe some of the experimental results from two distinct domains: a software environment, and a collection of 20 mobile robots. We also describe a compound behavior involving a combination of the basic interactions. Finally, we compare the performance of homogeneous groups to those of dominance hierarchies on the same set of basic behaviors.

Introduction

Our work is based on the belief that intelligence is, at least partially, a social phenomenon. In order to understand and analyze intelligent behavior, we need to study agents situated in social contexts. We examine group behavior by focusing on one of the simplest social contexts: a collection of agents sharing common goals, an analogy to a family of kin.

Group behavior is a result of the local interactions between the members of a group, and their interactions with the environment. Local dynamics between individuals produce consequences at the collective level. Thus, group behavior can neither be analyzed nor, for the purposes of AI and robotics, synthesized by observing a single individual. However, even the behavior of a single agent

is unpredictable, and the the multi-agent case is considerably more complex (Lozano-Pérez, Mason & Taylor 1984, Canny 1988, Brooks 1991).

In this paper we present a methodology for describing group behavior which allows for simplified synthesis and analysis. Our approach consists of observing a group behavior as a collection of basic behaviors consisting of simple local interactions, and combined into more complex aggregates. We demonstrate how the use of shared goals, kin recognition, and limited communication aid in generating such group behaviors. Finally, we validate our methodology on both software and physical agents, and evaluate and discuss the implications of the experimental data.

Inferring Goals

The behavior of a society as a whole is determined by the temporal consequences of the local interaction between the individuals. These local interactions are determined by the agents' goals, the amount of information each agent has about the others, and the agents' ability to act on that information. One of the determinants of the complexity of a society is the amount of variance among the individuals.¹ For the purposes of this work, the variance can be expressed as the difference in individual goals. We have studied the effects of the difference in the goal structure and the ability to communicate those goals on the complexity of the society.

In order for a society to function, it must overcome conditions of persisting inter-agent interference. Even in the simplest society in which all of the agents have identical goals at all times, conflicts, such as competition for resources, can arise. In more diverse societies where agents' goals differ, increasingly complex conflicts can persist, including clobbering of each other's work, deadlock, and

¹Assuming the variance is manifested in the individual behavior.

oscillations.

It may appear necessary for agents to be able to infer each other's goals in order to locally act in such a way as to produce coherent global results. This ability requires a high computational and cognitive overhead (Gasser & N. Huhns 1989, Rosenschein & Genesereth 1985, Axelrod 1984). However, work in both developmental psychology and ethology indicates that inferring the goals of other agents is not necessary for a large repertoire of complex interactions (Tomasello, Kruger & Rother 1992, McFarland 1987, Gould 1982, Rosenthal & Zimmerman 1978). The alternative is to base interactions on observable external behavior and its interpretation. Of course, the interpretation is determined by the amount of knowledge available to the agent. In developmental and ethological work the knowledge is innate and difficult to circumscribe. In contrast, computational and robot experiments allow us to vary the goals and the amount of built-in knowledge. We have performed a number of such experiments. In particular, our work has demonstrated that significant information about an individual's goals is reflected in the behavior, the externally observable state, and can be obtained with minimal if any direct communication.

Homogeneity and Similarity

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, or kin. The agents are *homogeneous* in that they are situated in the same world, embodied with similar abilities, and have similar goal structures. Such similarity has important implications.

Identical and similar agents have innate knowledge of each other. Thus, homogeneity allows for leaving much of the information about the world implicit. Since agents share a common goal structure, their behavior is implicitly or explicitly *predictable* to each other.

In addition to predictability, homogeneity offers the society flexibility in that agents are interchangeable. Given their similarity, agents do not need identities and thus do not require abilities for identification. Further, irregular behavior of any individual should not seriously affect the group, since no particular agent or group of agents is critical.

Taking advantage of homogeneity depends on a fundamental property: agents must be able to recognize kin, 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.

The following example describes the use of homogeneity and kin recognition. If when driving on a two-lane road we encounter an oncoming car, we are confident that the right behavior is to stay in our lane, since the other car will follow the same strategy.² However, if instead of a car an elephant is approaching, there is no clearly right behavior since there is no way of predicting what the elephant will do.³ As homogeneity and similarity greatly reduce individual cognitive requirements, we use them for simplifying the process of both generating and understanding group behavior.

Group Behavior from Basic Interactions

Irrespective of the simplicity of the individuals, the global consequences of even the simplest local interactions can be arbitrarily complex. In general, 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 do not lend themselves to any traditional methods of analysis.

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 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 are observable in the interaction space of the particular society. The most obvious ones take place in physical space (such as flocking, herding, following, traffic jams). More complex ones operate in informational space. For example, individual competitions for authority repeatedly lead to the formation of dominance hierarchies. Similarly, certain patterns of stock trading lead to stock market crashes. In both cases, the interactions between individuals are simple but the resulting global behavior is complex.

Our work is based on the belief that most of

²Incidentally, we use a similar strategy with our robots

³Example due to Rod Brooks.

group behavior consists of such simple basic interactions. Consequently, while the exact behavior of each individual may not be known, the collective behavior is qualitatively predictable and repeatable. While in practice most societies are too complex (in terms of individuals' behavior as well as their interactions) to be modeled analytically, stable group interactions can be used for qualitative analysis,⁴ 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. Consequently, in this work interaction means action. By placing our experiments in physical space, we can demonstrate group behavior on 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 methodology, environments, and results.

Experimental Methodology and Results

The goal of our work is to elucidate social interaction by synthesizing, observing, and analyzing phenomena similar to those observed in biology, sociology, and anthropology. Since behavior observation is the primary methodology for testing our theories, it is important to attempt to identify the effects of the experimental environment from those intrinsic to the interaction being observed. Toward this end, we use two very different test environments, one in software and one in hardware.

The software environment consists of an interaction modeler which allows for implementing a variety of agent types and group sizes situated in a simplified version of the physics of the world. The main purpose of the modeled environment is observing and testing a variety of group behaviors and comparing them to those observed on biological and synthetic agents.

The hardware environment consists of a collection of 20 foot-long mobile robots capable of detecting each other, and equipped with a forklift for picking up, carrying, and stacking objects. These basic abilities are used to construct various tasks

⁴An analogous notion can be found in complex dynamics. Basic group interactions correspond to attractors, the regions of phase space in which the behavior of the system is stable.

and experiments, in which the robots are run autonomously.

While the various sensory, mechanical, and computational limitations of the hardware environment severely limit the types of experiments that can be implemented, the environment offers some unique features. In particular, a physical implementation introduces the unavoidable variance among the agents which has an effect on the resulting group behavior. In spite of the identical software, the robots behave differently due to their slightly varied physical properties. This variance provides a stringent test for our group behavior strategies.

Communication and Cooperation

In order to focus on the simplest local interactions, as well as control the effects of communication on group behavior, we imposed some limitations on the type and amount of communication available to our agents. No explicit one-to-one communication between the agents was used in any of the experiments. 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* consists of 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 goals. In the described experiments, cooperation is implicit, as agents affect one another through their external state and actions.

We conducted a number of tests of various behaviors using no direct communication, and relying entirely on the agents sensing the external state of others. We subsequently added a local broadcast ability for some behaviors, allowing the agents to transmit a simple message within a small radius. Consequently, our work so far has used no explicit cooperation between agents.

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.

Experiments

Our experiments to date are based on physical interactions among agents. Consequently, the demonstrated group behaviors are the various manifested collective spatial and temporal patterns. We have developed a collection of simple local rules that implement the following basic group behaviors: collision avoidance, following, disper-

sion, aggregation, homing, and flocking.

These behaviors were demonstrated to be repeatable and reliable over multiple time-extended trials. The behaviors were shown to be stable over a variety of initial conditions, and insensitive to small perturbations in the various perceptual and effector variables. Further, most of the behavior primitives were implemented with multiple algorithms, whose variations in performance were due to the implementational details, but whose overall behavior was consistent. The details of the robot implementation are described in Mataric (1992). Additional experimental data is available both numerically and on video tape.

Combining Behaviors




Figure 1: The behavior structure of the foraging agents. Shaded states indicate carrying an object. Basic behaviors are italicized.

The basic behaviors generated by the interaction primitives demonstrate how very simple local interactions can generate a variety of useful global behaviors. Furthermore, they constitute useful building blocks for compound behaviors, which consist of spatial and temporal combinations of the basic behaviors. In our experimental spatial domain, these behaviors include herding (consisting of flocking and homing), convoying (avoidance and following), foraging and gathering, etc.

We are currently demonstrating foraging and gathering, which combines the described basic behaviors. The high-level goal of the group is to collect objects found anywhere in the environment and deliver them home. In addition to having the

basic social behavior repertoire, individual agents are also able to recognize and manipulate objects. However, they have no model of the environment, nor a global view of it.

The internal behavior structure of each of the agents is identical, and specifies the conditions triggering each of the basic social behaviors (Figure 1). Foraging is initiated by dispersion, followed by a search for objects. Finding an object triggers homing. Encountering another agent with a different immediate goal (as manifested by its external state, e.g. not carrying an object), induces avoidance. Conversely, encountering kin, another agent with the same external state (carrying something, or approaching an area with free objects) triggers following. Three or more followers initiate flocking. Reaching home and depositing the object restarts dispersion.

Foraging demonstrates how basic behaviors can be combined into a higher-level compound behavior. The combination is simple in that conflicts between two or more interacting agents, each potentially executing a different behavior, are resolved uniformly due to agent homogeneity. Since all of the agents share the same goal structure, they will all respond to the environmental condition consistently. For example, if a group of agents is flocking toward home and it encounters a few agents dispersing, the difference in the agents' external state will either induce kin following or non-kin avoidance, thus dividing the group again.

Our experiments have served to demonstrate two points: 1) coherent group behaviors of spatially interacting physical agents can result from simple local interactions and 2) complex, time-extended collective behaviors of such agents can result from simple combinations of relatively few basic group behaviors. We are now looking for similar simple interactions producing global consequences in other, non-spatial domains.

Heterogeneous Groups

In the experiments described so far the agents were fully homogeneous. As a control study, as well as an attempt to address a larger variety of social interactions, we have introduced dominance hierarchies into our agent societies.

In particular, we have tested the performance of dominance hierarchies using hierarchical control strategies on two of the basic behaviors described above: aggregation and dispersion. These two behaviors were chosen because simple performance evaluating criteria could be applied. Specifically, given sufficient space, aggregation and dispersion

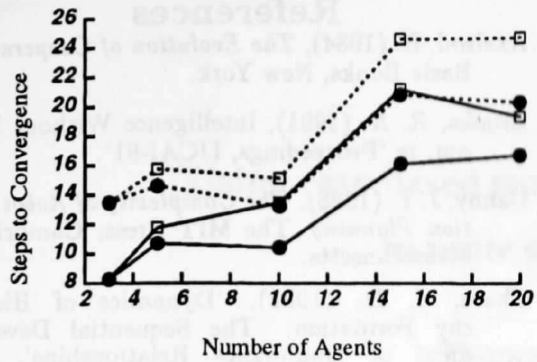


Figure 2: The performance of two different aggregation algorithms based on time to reach static state. Two termination conditions were tested: a single aggregate (data points shown with boxes) and a few stable groups (data points shown with dots). The performance of hierarchical algorithms is interpolated with solid lines while the homogeneous ones are interpolated with dotted lines.

can reach a static state, i.e. obtain the desired distance between the agents. Consequently, it is relatively simple to compare the performance of different aggregation and dispersion algorithms based on the time each required to reach static termination conditions. In contrast, following and flocking are not as simply evaluated due to their dynamic nature.

In this set of experiments the collection of agents was classified into a total order, based on a randomly assigned unique ID number, thus simulating an established pecking order in the group (Chase 1982, Chase & Rohwer 1987). Unlike the homogeneous algorithms, in which all agents moved simultaneously according to identical local rules, in the hierarchical case the ID number determined which agents were allowed to move while others waited. (In all cases, a simple precedence order was established such that within a small radius the agent with the highest ID got to move.) As in the homogeneous case, we tested multiple hierarchical algorithms for each of the group behaviors.

Using the software environment, we conducted 20 experiments with each group size (3, 5, 10, 15, and 20 agents) and each of the algorithms. Additionally, we tested the algorithms on two different degrees of task difficulty. For aggregation we tested two terminating conditions: a single aggregate containing all of the agents, and a small number of stable aggregates. The former terminating condition is more difficult. Similarly, for disper-




Figure 3: The performance of two different dispersion algorithms based on time to reach static state. Two initial states were tested: a random distribution (data points shown with stars) and a packed distribution (data points shown with crosses). The performance of the hierarchical algorithms is interpolated with solid lines while the homogeneous ones are interpolated with dotted lines.

sion we tested two initial conditions: a random distribution of initial positions, and a packed distribution in which all of the agents start out in one half of the available space. The latter condition is more difficult.

We found that, in the case of aggregation, hierarchical strategies performed slightly better than totally homogeneous ones. Figure 2 plots the average number of moves an agent takes in the aggregation task against the different group sizes and the two different terminating conditions: a single aggregate and a few stable groups. Both hierarchical and homogeneous algorithms behaved as expected, improving on the simpler of the two terminating conditions. Their performance declined consistently with the growing group size.

Unlike aggregation, in the case of dispersion, homogeneous strategies outperformed hierarchical ones. Figure 3 plots the average number of moves an agent makes in the dispersion task for the different group sizes on two different initial conditions: a random distribution, and a packed initial state. Again, both hierarchical and homogeneous algorithms improved with the easier initial conditions.

Although the performance difference between the homogeneous and hierarchical algorithms was repeatable and consistent, it was small, and its magnitude barely surpassed the standard deviation among individual trials for each of the algorithms and group sizes. The standard deviation was particularly significant in the case of small (3

and 5) group sizes. We believe that the difference between the two strategies would be negligible on physical agents.

The experiments comparing simple hierarchical and homogeneous algorithms demonstrate that for the described domain simple hierarchical strategies do not improve the global performance. In our experiments, hierarchies were assigned randomly since all of the agents are identical. More complex hierarchical strategies could be devised, but would require an increased perceptual and cognitive overhead. From our data we can hypothesize that for simple spatial domains 1) the simplest, homogeneous solution works well, and 2) quite a bit more knowledge and processing is required to significantly improve it.

Conclusions

In this paper we discussed an approach to describing and synthesizing group behavior. We proposed that for a given domain a set of basic interactions can be identified for describing a large variety of group behaviors. To validate our approach we have demonstrated an array of such basic group behaviors in the domain of spatial interactions of mobile agents. We implemented and tested the basic behavior set on two distinct domains and are currently using it to test compound behaviors with various goal structures.

Our work to date has demonstrated the simplifying advantages provided by agent homogeneity and similarity. We have also shown that for simple spatial behaviors no dominance hierarchies are necessary, or indeed helpful. We are currently extending our approach to domains involving information trading, imitation and social learning. In order to gain further insight into the dynamics of group behavior, we are continuing to pursue a synthetic approach, by generating, testing, and evaluating behaviors in varying environments and contexts.

Acknowledgements

Matt Marjanovic wrote the flocking software and the tools for recording and displaying robot data. Stanley Wang and Owen Wessling continually fixed the hardware and patiently ran experiments. This research was supported in part by the JPL contract 959333 and in part by the Advanced Research Projects Agency under ONR grant N00014-91-J-4038.

References

- Axelrod, R. (1984), *The Evolution of Cooperation*, Basic Books, New York.
- 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.
- Chase, I. D. (1982), 'Dynamics of Hierarchy Formation: The Sequential Development of Dominance Relationships', *Behaviour* 80, 218-240.
- Chase, I. D. & Rohwer, S. (1987), 'Two Methods for Quantifying the Development of Dominance Hierarchies in Large Groups with Application to Harris' Sparrows', *Animal Behavior* 35, 1113-1128.
- 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.
- 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.
- 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.