Behavioral Diversity as Multiagent Cooperation

Tucker Balch

Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA, USA

ABSTRACT

In many cases cooperation between robots is implemented using explicit, perhaps complex, coordination protocols. However, research in behavior-based multirobot systems suggests that effective cooperative teams can be composed of agents using simple individual agent behaviors with limited or no communication. In this paper we propose behavioral diversity as an alternative cooperative strategy. Behavioral diversity refers to the extent to which agents assume distinct behavioral roles in a group. By filling diverse roles the team members distribute responsibility for various components of the task. It is not always the case, however, that diversity is advantageous. Results of experiments in robotic soccer and multirobot foraging tasks indicate that the utility of diversity depends on the task. This paper describes behaviorally diverse solutions to these tasks and provides a comparison that suggests why some tasks are suited for behavioral diversity and others are not.

Keywords: multiagent robotic systems, cooperation, diversity

1. INTRODUCTION

Multi-robot team design is challenging because system performance depends significantly on new issues that arise solely from interaction between agents. Cooperation, robot-robot interference and communication, for instance, are not considerations for a single robot, but are crucial in multi-robot systems. All of these issues must be considered in the development of an effective robot team. This article focuses specifically on multiagent cooperation.

The dictionary defines cooperation as "to work or act together toward a common end or purpose." Planning and communication are among the most actively investigated means for implementing cooperation in robot teams.[6] In this paper, we present an additional approach: behavioral diversity. *Behavioral diversity* is defined in terms of behavioral differences between a robot team's individuals. The team is *behavioral homogeneous* if all robots are behaviorally equivalent to the others, and *behaviorally heterogeneous* otherwise.

In earlier work, we have proposed an information theoretic measure of the diversity in a multi-robot system; diversity is a sliding scale between the extreme where all robots are identical and the other where all are different.[3] In this paper we show how behavioral diversity can be used to enable cooperation in two tasks for multirobot teams: foraging and robotic soccer. These results have been reported individually in earlier work.[7,5] In this paper however, we present them together for the purpose of comparing and contrasting the results in each domain.

The remainder of this paper is organized as follows. The next section discusses the foraging and soccer tasks. After that, diverse multirobot behaviors for solving the tasks are presented and analyzed. We close with a discussion of the results.

2. MULTIROBOT TASKS

2.1. Robotic foraging

The forage task involves the collection of objects of interest (attractors) scattered about the environment. In a typical strategy, an agent begins by wandering about the environment looking for attractors. Upon encountering an attractor, the robot moves towards it and grasps it. After attachment, the robot delivers the object to a home base. In some foraging strategies attractors may be handed off to another agent for final delivery. In terms of the taxonomy introduced in[3] this task has the following characteristics: `TIME_LIM` because performance is measured over a fixed period, `RESOURCE_LIM` because as agents collect objects, the the availability of attractors is reduced, `COMP_INT` because robots on the team compete for access to attractors among themselves, and `SINGLE_AGENT` since an individual agent can perform positively, even though multiple agents may provide improved performance.

Correspondence: E-mail: trb@cs.cmu.edu; WWW: http://www.cs.cmu.edu/~trb.

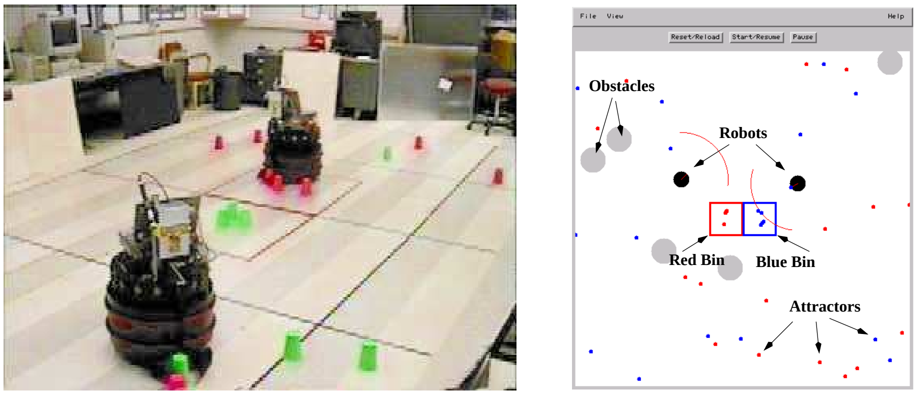

Figure 1. Real and simulated robot foraging. Left: two robots forage for colored attractors in the laboratory; after grasping an object, they deposit it in one of two delivery zones according to color. Right: in simulation, robots are represented as black circles, arcs indicate the robots' visual sensing range, obstacles are drawn as gray circles, the small discs are attractors. The robots deliver the attractors to the color-coded squares representing delivery areas.

Goldberg and Matarić have proposed a framework for investigating the relative merits of heterogeneous and homogeneous behavior in foraging tasks.[9] Like the research reported in this paper, their work focuses on mechanically identical, but behaviorally different agents. They propose *interference* as a metric for evaluating a foraging robot team. Interference refers to the situation where two robots attempt to occupy the same place at the same time; it is measured as the amount of time agents spend avoiding one another. Since interference may reduce the efficiency of a robot team, Goldberg suggests arbitration mechanisms for generating efficient behavior and reducing interference. Separately Fontan and Matarić have investigated a territorial heterogeneous foraging strategy where the search area is equally divided between agents.[8] Robots hand off collected attractors from area to area, with the last agent completing delivery to the homebase. Their work indicates that performance degrades if the number of robots is increased beyond a certain maximum.

Most robotic foraging tasks investigated to date involve the collection of attractors of a single type and their delivery to a single destination. This basic task is referred to as *simple foraging*. Simple foraging is an important robotic capability, but many practical industrial and military tasks call for more functionality. Consider, for example, a janitorial robot responsible for collecting and sorting recyclable trash objects into glass, aluminum and paper bins. Similarly, many assembly and construction tasks involve collecting parts or materials and placing them in a specific location. These more complex tasks are referred to as *multi-foraging* tasks. In general, the multi-foraging task calls for several types of objects to be collected and placed in specific locations according to type. Here *multi* refers to the multiple types of object to deliver, not the number of robots engaged in the task. An example team of robots executing a multi-foraging task is presented in Figure 1.

2.2. Robotic soccer

Soccer is a rich task for multiagent research because it requires cooperation between teammates, competition versus an opponent and unpredictable dynamic play. Recent interest has sparked more research in robot soccer. Kitano and Asada promote the Robot World Cup (RoboCup) as a vehicle for multiagent research.[10] They have developed an internationally agreed upon set of rules for a game involving mobile robots and a separate simulation system using the same rules. The simulated robotic soccer task explored in this work calls for a team of four agents to push a ball across an opponent's goal. Of course the opponent team defends its goal and strives to push the ball across the team's goal. The experiments reported in this work were conducted in the JavaSoccer simulation environment (Figure 6).

In terms of the taxonomy introduced in[3] this task has the following characteristics: `COMP_EXT` because robots on the team compete for positive performance (goals) against an (external) opposing team, `COMP_INT` because robots on the team compete for goals amongst themselves, and `MULTI_AGENT` since a single agent is unlikely to net a positive score differential against a multiagent opponent.

behavior	robot activity
wander	Search the environment for attractors.
stay_near_home	Search the home zone for attractors.
acquire_blue	Proceed to the closest blue object and grasp it.
acquire_red	Proceed to the closest red object and grasp it.
deliver_blue	Go the the blue delivery area.
deliver_red	Go to the red delivery area.

Table 1. Behaviors the robots select from in accomplishing the foraging task. The agents sequence from one behavior to another according to a fixed strategy.

3. DIVERSE FORAGING BEHAVIOR

A schema-based reactive control system is used for robot programming. In this approach, the agent is provided several pre-programmed behavioral assemblages that correspond to steps in achieving the task (e.g. *wander*, *acquire*, *deliver*, and so on[1]). Binary perceptual features (also referred to as perceptual triggers) are used to sequence the robot through steps in achieving the task. Selection of the appropriate behavior, given the situation, may be hand-coded or discovered by the robot through reinforcement learning.[3] This work focuses on hand-coded solutions.

The specific multi-foraging task in this investigation is for the robots to search for red and blue objects, collect them and deposit them into one of two color-coded bins. Three strategies for multi-foraging are implemented and evaluated:

- **Behaviorally homogeneous:** all the robots collect all the different types of attractor and deliver them to corresponding color-coded delivery areas.
- **Specialization by color:** half the robots specialize in collecting one type of attractor while the rest specialize in collecting the second type. Specialization by color is a heterogeneous strategy as well.
- **Territorial:** In this scheme one robot, referred to as the *sorting* agent, is responsible for collecting attractors within a three meter circle centered on the homebase. The other agents search at a distance from the homebase. When these robots find an attractor, they hand it off at the boundary of the "home zone." Final delivery is then completed by the sorting agent. This is a behaviorally heterogeneous approach.

To ensure a fair comparison between the various foraging systems, a fixed repertoire of behaviors is utilized across all implementations. A range of behaviors was developed to support several foraging strategies and to avoid bias towards any particular approach. The repertoire is suitable for building behaviorally homogeneous foraging teams as well as territorial and other heterogeneous strategies. Agents utilizing distinct strategies differ in the order they activate behaviors. summarized in Table 1. A more detailed description of the component behaviors is available in.[3]

At runtime, the robot behavioral systems utilize perceptual features to switch between behaviors. A perceptual feature is a single, abstracted bit of environmental or sensor state germane to the robot's task (e.g. whether or not the robot is holding an attractor in its gripper). Agents must decide on the basis of these environmental cues which behavior to activate at each point in time. The hand-coded teams are programmed as Finite State Automatons (FSAs) that sequence from one state to another. Each state of the FSA determines which behavior is to be activated, with transitions between behavioral states triggered when particular perceptual features are activated. This approach is called *perceptual sequencing*.[1]

3.1. Homogeneous foraging

We begin with a homogeneous approach, where all agents are programmed with the same strategy. In this approach each agent is able to complete the entire task on its own. It is referred to as *homogeneous* because all the agents are programmed with the same behavior. The homogeneous approach provides fault-tolerance because when one or more agents fail, the remaining robots can still accomplish the task. An FSA representing the homogeneous strategy is shown in Figure 2.

Each agent begins with the *wander* behavior activated, roaming about the environment in search of attractors. When a robot encounters an attractor, either the `red_visible` or `blue_visible` perceptual feature is triggered, causing the agent to transition to the corresponding *acquire_red* or *acquire_blue* behavior. Upon capturing an attractor,

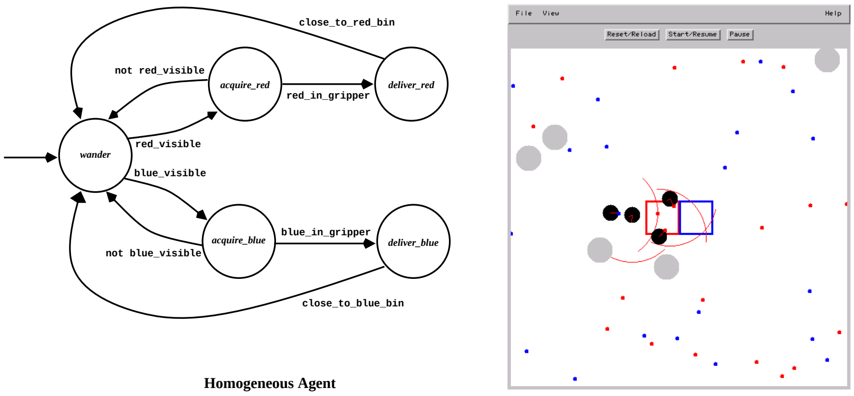

Figure 2. Homogeneous foraging. Left: an FSA representing the homogeneous foraging strategy. This kind of agent can collect all types of attractor. Right: inter-robot interference. In this simulation three robots are attempting to deliver attractors to the left (red) delivery area simultaneously. Interference reduces the rate at which agents can deliver attractors.

a robot returns back to homebase using one of the *deliver* behaviors. Finally, upon reaching the corresponding delivery area, the agent drops the attractor and transitions back to *wander*.

When using the homogeneous approach, multiple robots frequently attempt to deliver attractors to the same delivery area simultaneously; they interfere with one another and degrade performance. Figure 2 shows an example simulation where robots impede each other's progress as they attempt to deliver attractors to the left (red) delivery area. One way to reduce interference and potentially improve performance is to utilize heterogeneous strategies where the task is partitioned so that responsibility for portions of the task is divided among the robots.

3.2. Heterogeneous foraging strategies

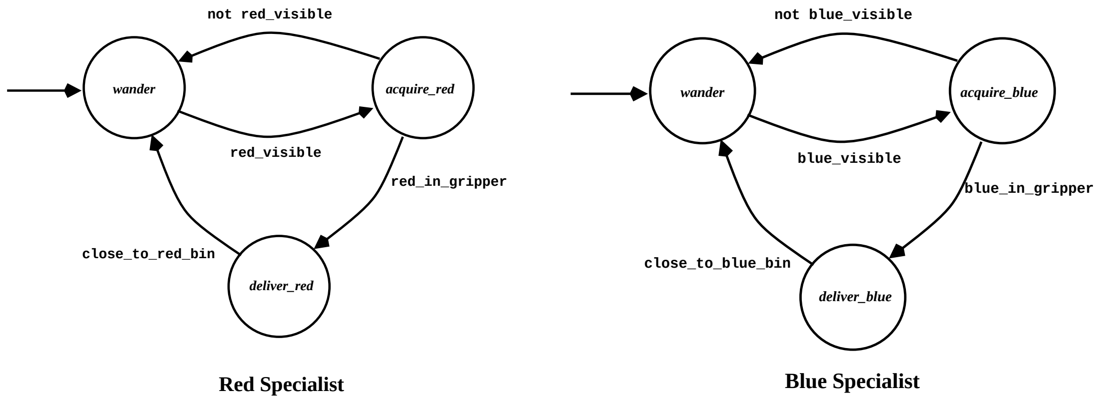

Figure 3. FSAs representing specialized behaviors for foraging. The FSA on the left shows the sequence behaviors are activated in for an agent specializing in collecting red attractors. The FSA on the right shows the sequence for blue specialists.

In the specialize by color heterogeneous approach, half of the agents are responsible for collecting the red objects and the other half for blue. This way, the chance of a "traffic jam" at either delivery area is reduced. FSAs for these specialized agents are illustrated in Figure 3. All agents start with the *wander* behavior activated. They begin to search the environment for attractors. What follows depends on whether the agent is a red specialist or a blue specialist. Red specialists ignore blue attractors, but when they encounter a red attractor while in the *wander* phase, they immediately transition to the *acquire_red* behavior. Similarly, blue specialists ignore red attractors. After acquiring an attractor, the agents deliver it to the appropriate delivery area using one of the *deliver*behaviors, then they switch back to *wander* to search for new items. An additional transition is provided for situations where an agent loses sight of the attractor. In that event it transitions back to *wander*.

3.3. Territorial specialization

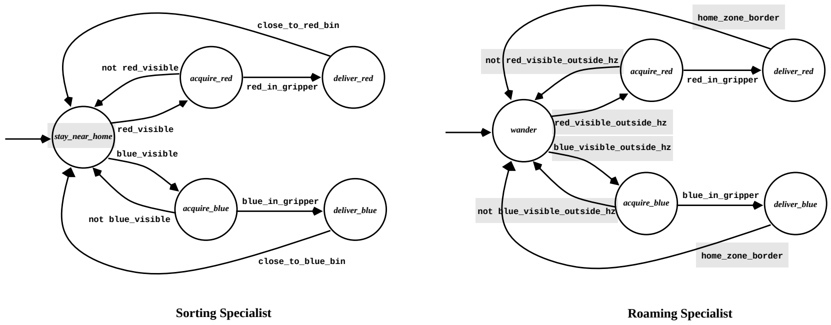

Figure 4. The Territorial behavioral sequences for foraging. FSA on the left shows the sequence of behaviors for an agent that remains close to the homebase, completing the delivery of the attractors. Agents using the strategy on the right search for attractors away from the home zone and deliver them to the home zone boundary. Differences from the homogeneous strategy are highlighted.

We now consider an alternative territorial strategy. In this scheme, one robot remains close to the homebase in the "home zone," delivering attractors that other agents deposit on the zone's boundary. Figure 4 shows the FSAs for robots in this system. The sequencing strategies for the agents are similar to the approach for homogeneous foragers (Figure 2). One significant difference is that the sorting agent utilizes a *stay_near_home* behavior rather than *wander* while searching for attractors. This results in the agent staying close to the delivery areas for sorting. The other agents are also similar to the homogeneous strategy, except they are triggered to drop attractors at the boundary of the home zone instead of depositing them in the delivery areas.

3.4. Discussion of foraging strategies and performance

The JavaBots system was utilized for simulation and mobile robot experimentation.[4,2] Behaviors coded in JavaBots may be run in simulation, and without modification, on Nomadic Technologies' Nomad 150 mobile robots. Statistical results are gathered by running the robot behaviors in thousands of simulation trials. Each strategy was tested with one to eight simulated robots, and one and two mobile robots.

As previously mentioned, inter-robot interference is a concern for homogeneous systems, while the heterogeneous strategies were specifically designed to reduce interference. Interference does occur during the delivery phase in both homogeneous and specialize-by-color strategies (Figure 5, left). This study does not include a quantitative measure of interference, but qualitative observations are consistent with results reported in[9]; namely that there is more frequent interference between agents in the homogeneous strategy than in other systems.

In the case of territorial foraging, robot-robot interference occurs much less frequently, but another factor limits performance. In most trials, the roaming agents quickly deliver a large number of the attractors to the boundary of the home zone, but the single sorting robot cannot always keep up. In simulations with seven and eight agents it is not uncommon for undelivered attractors to remain in a ring around the home zone at the end of the trial (Figure 5, center). Even though the number of sorters is constant (one), the territorial foraging experiments illustrate how the ratio of sorters to roamers impacts performance. The ratio varies from 1:0 to 1:7 as the number of agents goes from 1 to 8. In the 1:0 case, the sorting agent "starves" because it quickly finds all the nearby attractors. Conversely, in the 1:7 case, the sorter is overworked; there are many more attractors for it to deliver than it is able to.

A different sort of problem crops up for the specialize-by-color teams. Towards the end of trials for these agents there are often uncollected red attractors on the right side of the field and uncollected blue attractors on the left (Figure 5, right). This occurs because the agents inadvertently segregate themselves geographically to the left or right depending on whether they collect red or blue attractors. After delivering a blue attractor, for instance, a blue-collecting agent is more likely to remain on the same (right) side of the field as the blue bin. Because of this there are no red-collecting agents on the right side of the field and vice-versa. In large robot teams the robot-robot repulsion employed as part of the *wander* behavior serves as an additional force preventing the agents from diffusing from one side to the other. In addition to segregation, the specialize-by-color agents occasionally interfere with

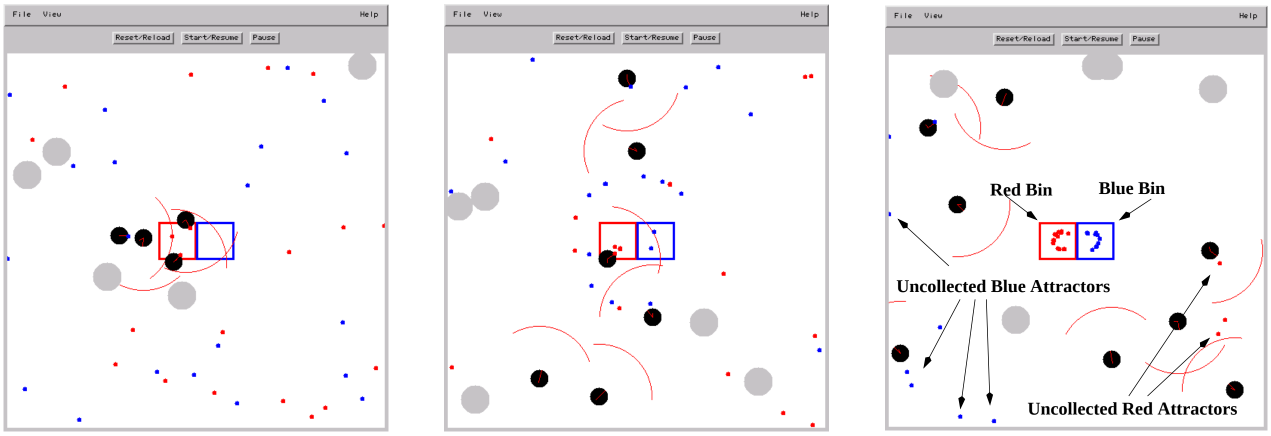

Figure 5. Simulations highlighting some of the factors that impact performance in hand-coded teams. From left to right: interference in a homogeneous team, an over-worked sorting robot, attractors left uncollected by agents specializing in one color of attractor.

behavior	robot activity
move_to_ball (mtb)	The robot moves directly to the ball.
get_behind_ball (gbb)	The robot moves to a position between the ball and the defended goal while dodging.
move_to_back_field (mtbf)	The robot moves to the back third of the field while being simultaneously attracted to the ball.

Table 2. Behaviors the robots select from in the soccer task.

one another when delivering attractors to the same delivery area simultaneously. These factors combine to drive performance down in the specialize-by-color teams.

Even though interference is reduced in the diverse foraging teams, the homogeneous teams still perform better. This observations is consistent with results reported in[9]; namely that there is more frequent interference between agents in the homogeneous strategy than in other systems, but performance is best in homogeneous systems. Full details regarding the performance of these systems is provided in.[3]

4. DIVERSE SOCCER BEHAVIOR

Behavior-based approaches are well suited for robot soccer since they excel in dynamic and uncertain environments. As in foraging, a motor schema-based approach is used for robot programming. Experiments in soccer are conducted by engaging an *experimental* team against a fixed opponent *control* team in soccer contests. We begin by describing the control team's behavioral configuration.

Since the experimental team's performance is significantly impacted by the skill of its opponent, it is important to provide an opponent "control" team to ensure consistent results. The control team will always follow a fixed policy against the teams under evaluation. The control team's design is based on the following observations. First, points are scored by bumping the ball across the opponent's goal. Second, robots must avoid bumping the ball in the wrong direction, lest they score against their own team. A reasonable approach then, is for the robot to first ensure it is behind the ball, then move towards it to bump it towards the opponent's goal. Alternately, a defensive robot may opt to remain in the backfield to block an opponent's scoring attempt.

As described above for foraging, each robot selects from a set of behavioral assemblages to complete the task. Instead of viewing the behavioral sequence as an FSA however, in this application we utilize a lookup-table representation (see Figure 3). The behavioral assemblages developed for these experiments are summarized in Table 2. A detailed description of these soccer behaviors is provided in.[3]

Control Team Forward

perceptual feature	assemblage *mtb*	*gbb*	*mtb f*
not behind_ball	0	1	0
behind_ball	1	0	0

Control Team Goalie

perceptual feature	assemblage *mtb*	*gbb*	*mtb f*
not behind_ball	0	1	0
behind_ball	0	0	1

Table 3. The control team's policy summarized as look-up tables. The 1 in each row indicates the behavioral assemblage selected by the robot for the perceived situation indicated on the left. The abbreviations for the assemblages are introduced in the text.

The overall system is completed by sequencing the assemblages with a selector that activates an appropriate skill depending on the robot's situation. This is accomplished by combining a boolean perceptual feature, *behind_ball* with a selection operator. The selector picks one of the three assemblages for activation, depending on the current value of *behind_ball*.

The control team includes three "forwards" and one "goalie." The forwards and goalie are distinguished by the assemblage they activate when they find themselves behind the ball: the forwards move to the ball while the goalie remains in the backfield. Both types of player will try to get behind the ball when they find themselves in front of it.

Given the behaviors and perceptual trigger introduced above there are 9 possible policies for each soccer robot, or 9^4 possible team team strategies. The space of team strategies was explored using reinforcement learning. In brief, the experiments evaluated the impact of three different reward functions on the diversity and performance of resulting robot teams. This report is not concerned with reinforcement learning techniques but with the various homogeneous and heterogeneous teams that result from the learning experiments. Here, we examine two example teams developed in the experiments. For more detail on the learning soccer team experiments please see.[3]

4.1. Homogeneous and heterogeneous soccer strategies

One learned strategy developed experimentally is illustrated on the left in Figure 6. Each robot has converged on a policy of going directly to the ball (*mtb*) when *behind-ball* is true, and activating *get-behind-ball* when *behind-ball* is false. This strategy corresponds to the "forward" policy used by the control team.

In a separate experiment, a diverse team emerged (Figure 6 right). In this team, three of the robots converged to the same "goalie" policy used by the control team (Figure 2), while the remaining agent utilizes a forward behavior. This was, in fact, the best performing team found in the experiments.

Experimental data were gathered by simulating thousands of soccer games and monitoring robot performance. The learning robots are evaluated on three criteria: task performance (score), policy convergence, and diversity of behavior. A series of 100 10-point games were played using each learning strategy. Information on policy convergence and score was recorded after each game. The robots retain their learning set between games. An experiment is composed of 10 runs, or a total of 1000 10-point games. Each run uses the same initial parameters but different pseudo-random number seeds.

Performance results show that teams converging to the homogeneous strategies described above lose most of the time by a score of 4 to 6 (games stop after 10 points are scored). In contrast, the heterogeneous teams win by a score of 6 to 4.

5. DISCUSSION AND CONCLUSION

In this study we have seen that behavioral diversity is beneficial in both foraging and soccer tasks. In foraging, the benefit is limited to reducing inter-agent interference; overall performance (number of attractors delivered) is actually maximized in the homogeneous foraging strategy. In soccer, however, we find that performance is maximized in heterogeneous teams.

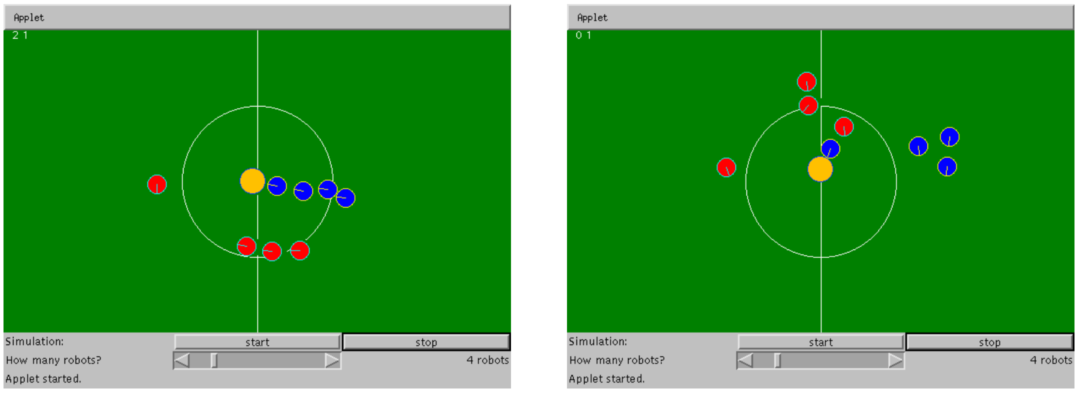

Figure 6. Examples of homo- and heterogeneous soccer teams. In both cases the learning team (dark) defends the goal on the right. The agents try to propel the ball across the opponent's goal by bumping it. A homogeneous team (left image) has converged to four identical behaviors which in this case causes them to group together as they move towards the ball. A heterogeneous team (right) has settled on diverse policies which spread them apart into the forward and middle of the field.

These observations beg the question "when is diverse behavior useful?" Task characterization is important for answering questions regarding how diversity can lead to different performance levels in different tasks. Without answering "how are soccer and foraging different?" for instance, we can't answer "why is diversity good in soccer but bad in foraging?" To address this, a taxonomy of multi-robot tasks was developed and presented in.[3]

The task classifications provide a clue: soccer and cooperative movement are both classified as `MULTI_AGENT`, while foraging is a `SINGLE_AGENT` task (this distinction is based on whether an individual agent could reasonably perform the task alone). We speculate that homogeneous teams excel in `SINGLE_AGENT` tasks while heterogeneity is preferred in `MULTI_AGENT` tasks.

REFERENCES

1. R. Arkin and D. MacKenzie. Temporal coordination of perceptual algorithms for mobile robot navigation. *IEEE Transactions on Robotics and Automation*, 10(3):276–286, June 1994.
2. T Balch. Javabots. `www.cc.gatech.edu/~tucker/JavaBots`.
3. T. Balch. *Behavioral Diversity in Learning Robot Teams*. PhD thesis, College of Computing, Georgia Institute of Technology, 1998.
4. T. Balch. Integrating robotics research with javabots. In *AAAI-98 Spring Symposium*, March 1998.
5. T. Balch. The impact of diversity on performance in multirobot foraging. In *Proc. Autonomous Agents 99*, Seattle, WA, 1999.
6. T. Balch and R.C. Arkin. Communication in reactive multiagent robotic systems. *Autonomous Robots*, 1(1), 1995.
7. Tucker Balch. Learning roles: Behavioral diversity in robot teams. In *AAAI-97 Workshop on Multiagent Learning*, Providence, R.I., 1997. AAAI.
8. M. Fontan and M. Mataric. A study of territoriality: The role of critical mass in adaptive task division. In *From Animals to Animats 4: Proceedings of the Fourth International Conference of Simulation of Adaptive Behavior*, pages 553–561. MIT Press, 1997.
9. D. Goldberg and M. Mataric. Interference as a tool for designing and evaluating multi-robot controllers. In *Proceedings, AAAI-97*, pages 637–642, July 1997.
10. M. Kitano, M. Asada, Y. Kuniyoushi, I. Noda, and E. Osawa. Robocup: The robot world cup initiative. In *Proc. Autonomous Agents 97*, Marina Del Rey, California, 1997.