Coordinating Mobile Robot Group Behavior Using a Model of Interaction Dynamics

Dani Goldberg and Maja J Matarić
Computer Science Department
University of Southern California
Los Angeles, CA 90089-0781
{dani,mataric}@usc.edu

Abstract

In this paper we show how various levels of coordinated behavior may be achieved in a group of mobile robots by using a model of the interaction dynamics between a robot and its environment. We present *augmented Markov models* (AMMs) as a tool for capturing such interaction dynamics on-line and in real-time, with little computational and storage overhead. We begin by describing the structure of AMMs and the algorithm for generating them, then verify the approach utilizing data from physical mobile robots performing elements of a foraging task. Finally, we demonstrate the application of the model for resolving group coordination issues arising from three sources: individual performance, group affiliation, and group performance. Corresponding respectively to these are the three experimental examples we present — fault detection, group membership based on ability and experience, and dynamic leader selection.

1 Introduction

Learning models of the environment, other robots, and interactions between them is a very challenging task in mobile robotics. Not only do noisy sensors and actuators pose inherent difficulty, but in the multi-robot domain which is the focus of our work, non-stationarity is an additional challenge. Limited computational and memory resources, in conjunction with often limited amounts of training data, can make brute force approaches to many learning techniques (i.e., hidden Markov models, partially observable Markov decision processes, reinforcement learning) intractable on mobile robots. In order to achieve convergence, it is often necessary to bias the learning system, for example, by providing an appropriate initialization, choosing a tractable search space, and/or making heuristic modifications to the learning algorithm.

We present *augmented Markov models* (AMMs) as an approach to creating behavioral models of robot/environment interaction dynamics, which accommodates the domain challenges and limitations mentioned above. The approach is computationally inexpensive, incrementally generating and modifying parsimonious models in real-time using only a small continuous stream of training data. In addition, it produces models that are available for use almost immediately, even during the early stages of learning. There is also no need for model initialization, besides a specification of the robot's *behavior space* (i.e., its set of behaviors). The basic structure of an AMM is a Markov chain, with each state representing a behavior, and with probabilistic transitions (links) between states. The Markov chain is augmented with statistics on state and link use which are employed in model construction, modification, and utilization.

AMMs have a number of characteristics that make them naturally applicable to the robotics domain: they are compact, have a low computational overhead, and can be generated and used in real-time in one trial.

In this paper we describe how AMMs may be used to coordinate activity in a group of mobile robots. We consider three issues — individual performance, group affiliation, and group performance — that impact the ability of a group to achieve effective coordination, and present applications of AMMs that help negotiate these issues. As an example of the impact of individual performance on group coordination, consider a scenario where a single robot develops a hardware failure and is neither able to complete its portion of the group task, nor to inform the other group members of its failure. If the members do not know to compensate for the incapacitated robot, the group as a whole may fail to complete its task. Monitoring individual robot performance, in this case for fault detection (one of our experimental examples), is thus an important component of group coordination.

The ability of a robot to determine what group it belongs to (i.e., its group affiliation) is another important component of group coordination. Suppose a robot were introduced into an environment containing several groups specializing in different tasks. In order to be able to coordinate its activity with the group it fits into best, it must have some mechanism for determining its group affiliation. One way is to compare its abilities and experience with those of other robots — an approach we present in another of the experimental examples.

A third issue impacting coordination is group performance. Consider a group of robots organized in a hierarchy. As with many hierarchies, the performance of the entire group is strongly dependent on the members in the upper strata. In corporate management, for example, the CEO may be considered most responsible for the overall performance of the company. If the company becomes financially unstable, some form or restructuring may occur, including the replacement of the CEO. This ability to dynamically reorganize the structure of a group (re-coordinate the individuals) is important to performance, especially in the case of hierarchies where significant reliance is placed on the leader. We also present experimental results for this type of dynamic leader selection later in the paper.

The next section presents the details of the AMM modeling approach.

2 Augmented Markov Models

An augmented Markov model (AMM) is a deterministic finite state automaton with probabilities associated with transitions from state to state, i.e., a Markov chain [4]. The Markov chain is expanded with data used for dynamic model construction, modification, and exploitation.

2.1 Structure of AMMs

An AMM is defined by the following five elements:

1. S, a set of symbols $\{s_1, s_2, \ldots, s_M\}$ recognized by the network.
2. A, a set of states (or nodes) $\{a_1, a_2, \ldots, a_N\}$, each of which recognizes one of the symbols in S. Each state also has several statistics associated with it: the mean number of time steps spent in the state; the variance associated with this mean; and the probability of remaining in the state in the next time step.
3. B, an $N \times M$ transition matrix, where $b_i(k)$ contains the value of the state to transition to if the current state is a_i and symbol s_k is observed. If the current observed symbol is identical to the last symbol, then the system remains in the same state.
4. L, a set of directed links $\{l_1, l_2, \ldots, l_P\}$, connecting the states. Each link has two attributes indicating which state it connects *from* and which state it connects *to*. There are also several statistics for each link: the number of times the link has been traversed; the sum of all durations spent in the state that the link connects *to*, after first traversing the link; the sum of squares of these durations; and the probability of using the link given the system is in the state that the link connects *from*.
5. T, a set of elements $\{t_1, t_2, \ldots, t_Q\}$, each of which stores information on a particular two-link traversal sequence entering and leaving a state. Each t_i has two attributes, indicating the link entering the state (in-link) and the link leaving the state (out-link). Additional statistics include: the sum of all durations in the state that the outgoing link connects to; and the sum of the squares of these durations.

Given this notation, a Markov chain could be represented by S, A, B and L, without the statistics on durations spent in various states. These statistics provide the augmentation necessary for model construction and utilization.

2.2 Generation of AMMs

The data used for constructing an AMM consist of a continuous stream of symbols belonging to S. In our implemented robot behavior model, the symbols would be the names of behaviors that the robot performs, sampled at some appropriately high frequency to avoid aliasing. Construction of an AMM proceeds according to the following algorithm:

- Initialize the system by creating the initial state, $a_\emptyset$.
- If the current input symbol has never been seen before, add it to S and create a state that recognizes this symbol. Create a link from the current state to the new state and make the transition. Add this transition to B and create the corresponding new entry for T.
- If the current input symbol is the same as the last input symbol, then remain in the current state. Update the appropriate values in A, L, and T for mean values and length of time in the state. Recalculate the transition probabilities associated with that state and its links.
- If the current input symbol has been seen before but is different from the last symbol, transition to a state that accepts the new symbol. If the link for this transition does not exists, create it.
- When transitioning from one state to another, update the variance for the state being transitioned from, and the appropriate sum of squares values in L and T.
- When about to transition from one state to another, calculate the binomial mean and variance for transitions in T with the same in-link. If the mean falls outside c times the standard deviation, where usually $2 \leq c \leq 3$, then *split* the current state, and attach the current in-link (with its associated out-links, as indicated in T) to the new state. Using T, make all appropriate changes to the two states and their related links, in order to keep all global probabilities consistent. Update T appropriately.

The above rules do not provide the complete semantics, but capture the general flavor of the model construction process. The final rule of the algorithm describes *node splitting* and deserves further explanation. Since an AMM is constructed incrementally as training data become available, it is important that there be some mechanism for model modification when new data invalidate the current structure of the model. This mechanism is provided by node splitting which utilizes data from T and many of the statistics maintained in the AMM, ensuring that the model remains consistent with the training data.

Note that there is relatively little computation involved in model construction. The computational complexity per input symbol is $O(N^2)$ (for a fully connected network) when a state is being split, and at most $O(N)$ otherwise. In addition, the space complexity is $O(N^3)$ for N fully connected states. In practice, the space complexity seems likely to fall closer to $O(N)$ or $O(N^2)$ for something less than a fully connected graph. This may be considered parsimonious since arbitrarily large amounts of training data need not be additionally stored. [8] provides the full details of AMM structure and the construction algorithm.

2.3 Operations on AMMs

Aside from construction, there are two general operations that may be performed on a given AMM: generating data, and calculating the probability of a symbol sequence.

2.3.1 Data Generation

Data generation using a given AMM is analogous to simulating the data generating process that the AMM is modeling. If the AMM is a good model of that process, the original data and the generated data will "look" alike. The algorithm for data generation is as follows:

1. Begin at some initial state.
2. Generate the symbol that the current state accepts.
3. Select a transition to the next state according the probability distribution for the current state and all links leading from that state. Make the transition.
4. Return to step 2.

2.3.2 Probability of a Symbol Sequence

Given an AMM M and a sequence of symbols $O = \{O_1, O_2, \ldots, O_k\}$, where $O_i \in S$, we wish to calculate $P(O|M)$. This may be done in a straightforward manner using the following algorithm:

1. Repeat the following for each state $a_i \in A$ that accepts the symbol O_1, considering it the initial state:
   (a) Initialize the probability $p_i = 1$.
   (b) For each symbol from $O_2, ..., O_k$ in turn: make a transition from the current state to the state that accepts that symbol, multiply p_i by the probability of that transition, and set the current state to be the new state. If no such transition exists, then set $p_i = 0$.
2. $P(O|M) = \max_i p_i$.

Next we consider in more detail the application of AMMs to the mobile robotics domain.

3 Model Use with Mobile Robots

Due to their low computational and storage overhead, AMMs can be constructed and maintained as a mobile robot is performing a task, in order to capture the dynamics of its interaction with the environment. At each time step, the data used for generating and updating a model consist simply of the symbol indicating which behavior is currently active. There is no need to heuristically modify the AMM learning algorithm or provide any application-dependent initialization to the system. What is required, however, is the determination of the *behavior space*, i.e., the set of mutually-exclusive behaviors, which continuously describe the robot's activity, and are uniquely labeled. One of these labels or symbols (together comprising S) is sent to the model generation algorithm at each time step. It is important to note that the algorithm only accepts one symbol per time step. If the robot's control system is structured so as to utilize simultaneous execution of two or more behaviors, the AMM algorithm will only be able to consider one of their symbols as input. Unless that one symbol is consistently used to represent the parallel execution of both behaviors, the result will be a model unrepresentative of the actual behavior dynamics. The complementary situation where no input symbol is available also results in ill-defined models.

Since neither of these situations is desirable, it is necessary for the behavior space to be composed of mutually exclusive behaviors which account for all of the robot's activity. In other words, in behavior space $\mathcal{B} = \{\mathcal{B}_1, \mathcal{B}_2, \ldots, \mathcal{B}_K\}$, no two behaviors $\mathcal{B}_i$ and $\mathcal{B}_j$ may be simultaneously active. In addition, if $\mathcal{T}_{\text{tot}}$ is the total time that the robot is active, and $\mathcal{T}_{\mathcal{B}_i}$ is the total time that behavior $\mathcal{B}_i$ is active, then $\mathcal{T}_{\text{tot}} = \sum_{i=1}^{K} \mathcal{T}_{\mathcal{B}_i}$. The behavior space $\mathcal{B}$ need not contain all of the behaviors that the robot may exhibit, only some representative subset of those. In addition, an individual behavior $\mathcal{B}_i$ in the behavior space may be composed of several more primitive behaviors that are executed in parallel.

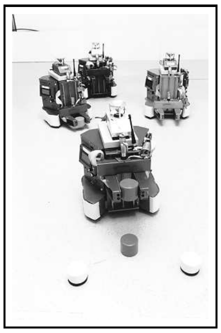

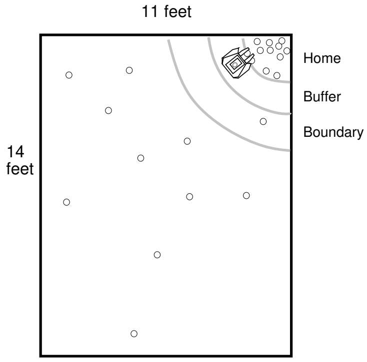

Figure 1: (Left) The Don Group: the four IS Robotics R2e robots used in the experiments; (Right) The foraging task environment.

In the unlikely worst case scenario where all of the behaviors on a robot are executed in parallel for $\mathcal{T}_{\text{tot}}$, or that there is only one single behavior, $\mathcal{B}$ also consists of exactly one behavior. The AMM generated in this degenerate case has one non-initial state that it never leaves, and consequently is of no practical use. Fortunately, the large majority of current approaches to robot control, ranging from hybrid to behavior-based system, utilize sequences and priorities over the actions and behaviors executed on the robot [6, 12]. Thus, as long as the behavior of the robot can be decomposed into a sequence of mutually-exclusive behavior symbols to be provided to the AMM, this model can be applied.

In the next section we describe our experimental example demonstrating elements of foraging, the associated behavior space, and the resulting AMMs.

4 Experimental Example

For testing the construction and utilization of AMMs, we employed behavior data from real mobile robots performing complete and partial versions of a foraging (collection) task, a prototype for various applications (distributed solutions to de-mining, toxic waste clean-up, terrain mapping, etc.). The data were used off-line to validate the modeling approach, and on-line to demonstrate potential applications. The foraging experiments presented here builds on our previous work in the same task domain [7].

4.1 The Robots

We used four IS Robotics R2e robots (Figure 1: Left), programmed in the Behavior Language [2]. Each is a differentially steered base equipped with two drive motors and a two-fingered gripper. The sensing capabilities include piezo-electric contact sensors around the base and in the gripper, five infrared (IR) sensors around the body and one on each finger, a color sensor in the gripper, a radio transmitter/receiver for communication and data gathering, and an ultrasound triangulation system for positioning. Experiments were performed in an 11 × 14 foot rectangular enclosure (the Corrall). For the complete foraging task, 27 small metal cylinders (pucks) were randomly but evenly distributed throughout, except in the drop-off area, called *Home*, a ninety degree sector of a circle with a radius of two feet, located in a corner of the Corrall (Figure 1: Right).

4.2 Behaviors

The Behavior Language, used for programming the robots, provides a natural structure for implementing behavior-based controllers. Behaviors are defined as sets of concurrent, real-time, asynchronous rules receiving inputs from sensors and behaviors, and sending outputs to actuators and behaviors. A set of these behaviors can often trivially map to the behavior space for an AMM, as is the case with our experimental foraging task. The following 7 behaviors constitute the behavior space for our task:

- **avoiding**: avoid any object (detected by IR and contact sensors) deemed to be in the path of the robot.
- **wandering**: move forward and, at random intervals, turn left or right through some random arc.
- **puck detecting**: if avoiding is not active, and if an object is detected by the front IRs, lift up the fingers to determine whether the object is short enough to be a puck. If it is, approach the object and try to place it between the fingers and pick it up. If unsuccessful, perform **avoiding**.
- **homing**: if carrying a puck, move towards the designated goal, Home.
- **creeping**: when near Home, perform a slower, more accurate homing behavior.
- **exiting**: if in the home region, drop puck and exit home, and perform **reverse homing**.
- **reverse homing**: move away from the home region.

The basis for selecting the constituent members of this behavior space was the control of the robot's drive motors. When active, each of the behaviors has exclusive control of the motors, and together they account for all activity (or inactivity) of the motors for the duration of the task.

4.3 Results: Model Construction and Validation

Using the data generated by a robot performing the foraging task, the AMM algorithm constructed the model shown in graphical notation in Figure 2 (Left). Many of the numerical details are omitted, but the main elements are present. The states with their recognized symbols (e.g., "avoiding", "wandering") are indicated, as well as the initial state $a_\emptyset$ shown with the symbol $\emptyset$. Links between states are shown, as are the two-link transitions, represented by dashed lines inside states connecting an in-link and an out-link. For the construction of this particular model, we used approximately 1750 data points sampled at 5 Hz over the course of approximately 7 minutes. All together, this model required 35,000 flops to construct and 2700 bytes to store, which is both parsimonious and computationally inexpensive. Similar results were obtained with data from our other robots performing the same task.

The generated models were validated by visual comparison of the robot's behavior (while performing the task) to the structure of the models. Numerous hours of video footage of the foraging task support the validity of the AMMs. One issue that requires special mention, however, is node splitting. Are models generated with node splitting potentially more correct than those generated without?

To ascertain the validity of node splitting, we used the foraging data to construct AMMs with and without node splitting. Figure 2 (Center) shows the AMM generated without node splitting for the same data as were used in Figure 2 (Left). The result is that the state that recognizes "avoiding" is not split into three states. Examining the two-link transitions inside the state, we note that the transition from "wandering" to "avoiding" is followed by a transition from "avoiding" back to "wandering" in all cases. A similar situation holds for "reverse homing" and "homing." If "avoiding" were really one state, intuitively we would expect transitions between "wandering," "homing" and "reverse homing" that pass through "avoiding." These transitions, however, would be inappropriate in terms of the foraging task. Transitioning, for example, from "wandering" through "avoiding" to "homing" is incorrect without first having picked up a puck. Node splitting provides a principled statistical mechanism for deciding when the current representation of a state does not conform to the data and therefore should be changed, i.e., split. Node splitting facilitates the generation of models that better represent the data and thus are truer to the actual behavior dynamics of the system.

As an additional form of validation, we generated new data using the AMMs constructed with and without node splitting, then used this new data to construct AMMs. As one would expect, the data generated by the AMMs without node splitting produced AMMs with transitions that were non-existent in the original data, whereas the AMMs with node splitting did not.

In the next section, we apply AMMs to three experimental examples exploring issues in group coordination.

5 Group Coordination Results

As mentioned previously, issues such as individual performance, group affiliation, and group performance impact how a mobile robot group coordinates its activity. For each of these issues, we present an example application (fault detection, membership based on ability and experience, and dynamic leader selection) in the context of our group foraging task, along with experimental results.

5.1 Individual Performance: Fault Detection

For this application, we limit our consideration of faults to those that would keep the robot in one behavior for an inordinate period of time. Such faults may include sensor and actuator failures, as well as the robot becoming physically stuck.

To detect a potential fault, we modified the AMM construction algorithm to compare, at each time step, the total time the robot has spent in the current model state a_i to the mean and variance calculated from previous data for that state. A simple confidence estimate on the upper bound of the mean can be used to make the comparison:

$$\mu = a_i(mean) + c\sqrt{a_i(var)}$$

where c is a small positive constant (e.g., $1 \leq c \leq 3$). If the current time spent in the state exceeds μ, the algorithm signals that there might be a fault.

We tested this algorithm on-line by having the robots perform wandering and avoiding behaviors of the foraging task. Figure 2 (Right) shows a typical AMM that was constructed. If the model detected that the robot had been in

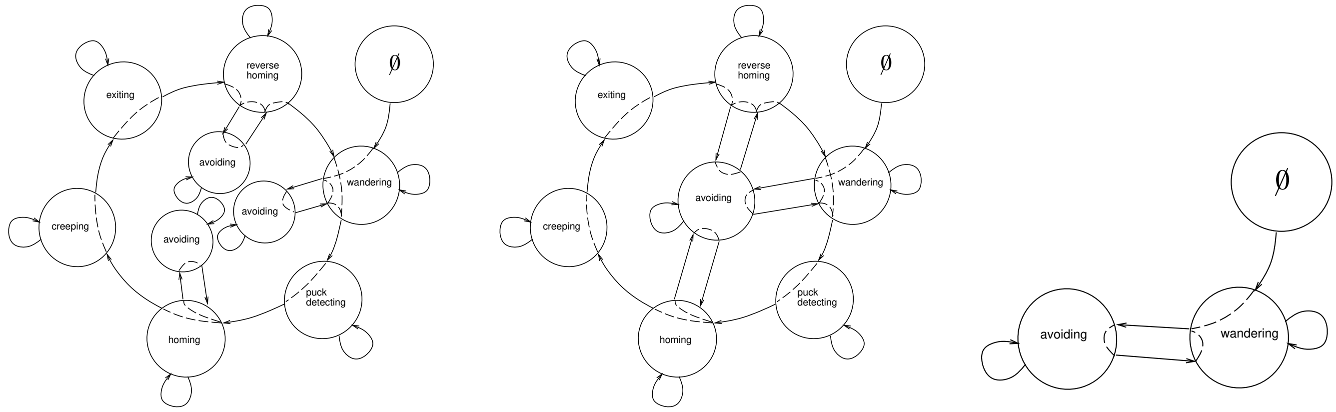

Figure 2: (Left) The AMM constructed from foraging behavior data; (Center) The same AMM, constructed without node splitting; (Right) Sample AMM constructed from the wandering and avoiding behaviors of the foraging task.

one of the behaviors too long, it would send a signal to the robot, which would in turn beep, thereby indicating a potential fault. We simulated a fault (the robot getting stuck on a rock) by lifting the drive wheels off the ground. During a dozen trials we conducted, the robot never failed to detect the fault.

In choosing a particular confidence interval, it is important that one be aware of how rapidly the interval narrows as new data are incorporated into the model. If the interval narrows too quickly, the result will be many false positives, i.e., the robot indicating a fault when none exists. An example of such a confidence interval that narrows too quickly is

$$\mu = a_i(mean) + t_{.001}\sqrt{\frac{a_i(var)}{n}}$$

where $t_{.001}$ is Student's t distribution leaving 0.001 probability in the tail, and n is the total number of times the system has entered a_i. The first confidence interval may be interpreted as this one, but with $t_p = \sqrt{n}$. This allows us to test for a fault with increasing confidence (i.e., with decreasing values of p) as the number of data points n increases, thus effectively reducing the narrowing of the confidence interval to the rate of variance decrease, and greatly reducing false positives.

Unfortunately, due to limitations in the programming environment of the robots, the models generated for fault detection could not be constructed on the robots themselves. Instead, behavior data were transmitted via radio modem to a Power Macintosh that constructed the models and communicated model information back to the robots. Due to other system limitations, the communication throughput was approximately 2.5 bytes/sec/robot. Even at such rates, and with lost packets, the system was able to maintain useful models simultaneously for multiple robots. This example of AMM robustness provides evidence of its utility to mobile robotics.

5.2 Group Affiliation: Membership through Ability and Experience

Coordinating activity in a behaviorally heterogeneous group of robots may require the robots to know their sub-group affiliations. In a learning system where the robot's final behavior is not predetermined, group affiliation is not designated *a priori*. AMMs provide a mechanism for determining group affiliation. Two robots that wish to ascertain whether they belong to the same group can transmit data generated by their AMMs, then determine the probability of the other robot's data on their respective AMMs. If each AMM accepts data generated by the other's AMM (with probability >0), then the robots are designated as members of the same group. They are considered to have the same *ability*, or capacity for performing a particular task. In the case of a complex task, such as our foraging example, it may take a significant amount of time for two robot that "should" belong the same group to explore the same behavior/environment interaction dynamics and be determined to actually belong to that group. In the context of our experimental example, there might be several different groups in an area, with only one performing foraging. When a new foraging robot is introduced into the environment, it must determine that its abilities coincide with those of the other foraging robots, and join their group.

In addition to this coarse "don't accept"/"accept", or ability-based, determination of group affiliation, a more refined categorization can be made by considering the actual probabilities of symbol sequences. To test this notion, we ran 2 trials for each of 3 robots performing the wandering-avoiding behaviors. In one trial, the Corrall was empty, in the other, 18 pucks were distributed evenly, though sparsely, on the floor. The robots occasionally avoided the pucks in addition to normally avoiding the walls of the Corrall. For each of the six three-minute trials, an AMM of the robot's behavior was constructed. One thousand data points were generated by each of the six AMMs, and the probability of each data set was calculated on each of the AMMs, resulting in 36 probability values. Our hypothesis was that a data set from an AMM generated in one of the two environments should produce higher probabilities on the remaining two AMMs from that environment than on the three from the other environment. For each data set, we tested all combinations of two AMMs, one taken from each environment, to determine how often the higher probability would be from the same environment. Of the 36 such combinations, in 26 (or 72%) of them, the same environment had the higher probability. These results, produced from little training data and almost identical environments, suggest that AMMs can be used to make subtle behavioral distinctions. These distinctions can be thought of as experience-based. Since the robots are able to and do perform the same task, it is their specific individual experiences that differ, and are the basis for distinction.

5.3 Group Performance: Dynamic Leader Selection

Due to inherent variations in sensors and actuators, or inexperience with a specific robotic platform, it may be difficult to accurately assess the ability of a robot at performing a novel task. Alternatively, even if performance history is available, there is no guarantee that future performance will neither improve nor degrade. The ability of individuals may change over time, and it is important that the performance of the group remain as high as possible. To achieve this, some mechanism for dynamic restructuring based on performance is necessary, especially in social structures such as hierarchies where significant reliance is placed on the most dominant individuals. We present dynamic leader selection using AMMs as an example of such a mechanism.

In the following experiments, the four R2e robots had to perform the foraging task described previously. Collecting 10 of the 27 pucks in the Corrall constituted completion of the task, with a shorter completion time corresponding to better group performance. The robots were organized in a strict dominance hierarchy such that whenever two or more robots simultaneously had pucks to deliver to the goal, the most dominant individual was allowed to proceed, while the less dominant individual(s) each waited their turn. The four robots, however, were not equally efficient at performing the task. The code for each robot was identical, except that the maximum speed was limited to different values, as follows: Robot0 "full-speed" ($\approx$ 0.5 ft/sec); Robot1 "two-thirds-speed" ($\approx$ 0.33 ft/sec); Robot2 "half-speed" ($\approx$ 0.25 ft/sec); and Robot3 "one-third-speed" ($\approx$ 0.17 ft/sec).

We conducted three sets of experiments, two with fixed hierarchies as baselines of comparison to the third, which allowed hierarchy restructuring through the use of AMMs. The experiments were designated as follows:

1. **Control**: The robots were members of a fixed hierarchy with the relative dominance of each inversely proportional to its maximum speed. Thus, Robot3 (the slowest) was the most dominant, and Robot0 (the fastest) was the least dominant.
2. **Optimal**: Complementary to Control, these experiments had the robots arranged in a fixed, optimal hierarchy, with the fastest as most dominant, and slowest as least dominant.
3. **Dynamic Leader Selection (DLS)**: The hierarchy was initialized to be identical to that of the Control experiments, but allowed hierarchy restructuring to improve performance.

In the DLS experiments, with no *a priori* information about a robot's speed provided, an AMM for each robot was constructed at run-time and used to evaluate performance. The metric of evaluation employed was

$$\text{performance} = \frac{\text{number of transitions to } \textit{exiting} \text{ state}}{\text{total time spent } \textit{homing}}$$

where the statistics for the exiting and homing behaviors came directly from the AMMs. The numerator gives a count of the number of pucks that were delivered to the goal, while the denominator measures the total time spent delivering those. The ratio gives the number of pucks per unit time that a robot is able to deliver: the higher this value, the faster the robot delivers pucks, and the better its performance. Each robot began a trial with its performance value initialized to zero. As it executed the task, its AMM was continuously updated, as was the performance value derived from it. The robot's position in the hierarchy was also updated so that it was more dominant than all other robots with lower performance values.

	Control	DLS	Optimal
Mean time to completion	27.2	23.4	22.4
Standard deviation	1.1	1.3	1.1

Table 1: Mean time to completion for the Control, Dynamic Leader Selection, and Optimal experiments.

Hierarchy Positions: Dynamic Leader Selection Experiments				
Robot ID	0	1	2	3
Mean position	2.6	1.8	1.4	0.2
Standard deviation	0.9	0.4	1.1	0.4

Table 2: Mean positions in the hierarchy at the end of the Dynamic Leader Selection experiments.

We ran five trials of each experiment (Control, Dynamic Leader Selection, and Optimal) and used a statistical hypothesis test based on Student's t distribution [5] to ascertain the significance of our results. (Heretofore, the word "significant" as applied to numerical results implies statistical significance with p-value$<$ 0.05. The figures provide p-values as indicated by asterisks.)

Table 1 and Figure 3 (Left) present the average time to completion (i.e., group performance) for the three experiments. In the experiments using dynamic leader selection we see a significant improvement in the time to completion over the Control experiments, thus indicating a successful restructuring of the hierarchies to a more optimal configuration. The Optimal time is slightly, though not significantly, lower than the DLS time. This difference may be attributed to the fact that the DLS experiments are initially configured with the less efficient Control hierarchy structure.

The successful restructuring of the hierarchies is evident in Figure 3 (Center) and Table 2. We see that the final hierarchy positions in the Control and Optimal experiments are unchanged from the initial positions since no hierarchy restructuring was allowed to take place. Even though the initial hierarchy positions in the DLS experiments are identical to those of the Control experiments, we see in Figure 3 (Center) that by the end of the trials, the positions are almost identical to the Optimal ones, though always lying between the Control and Optimal.

The fact that the average DLS positions are not exactly equal to the Optimal positions (indicating non-optimal restructuring) requires further explanation. At the beginning of each DLS trial, when all of the robots have performance values of zero and the hierarchies are still identical to the Control, it is very likely that one or more of the slower, more dominant robots will find a puck and be allowed to deliver first, thereby attaining a non-zero performance value. This helps establish its position in the hierarchy and makes it

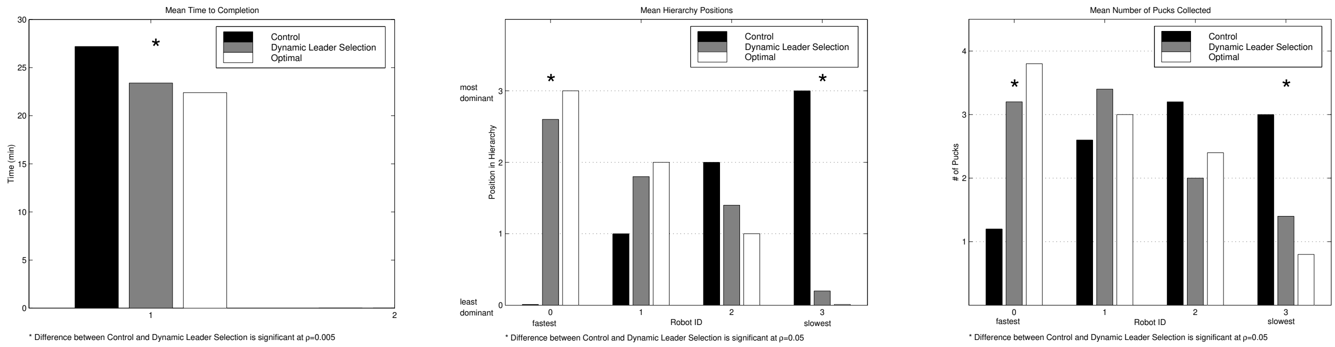

Figure 3: (Left) Mean time to completion. (Center) Average hierarchy positions at the completion of the Control, Dynamic Leader Selection, and Optimal experiments; (Right) Mean number of pucks collected.

Pucks Collected												
	Control Experiments				*DLS Experiments*				*Optimal Experiments*			
Robot ID	0	1	2	3	0	1	2	3	0	1	2	3
Mean number of pucks collected	1.2	2.6	3.2	3.0	3.2	3.4	2.0	1.4	3.8	3.2	2.2	0.8
Standard deviation	0.4	0.5	0.8	1.0	1.3	0.9	1.2	0.5	0.4	0.8	1.1	0.4

Table 3: Mean number of pucks collected in the Control, DLS, and Optimal experiments.

more difficult for the faster, less dominant robots to get a chance to deliver a puck and get a better performance value. In addition, the foraging task is quite stochastic. On average, all of the robots will find pucks the same distance from the goal, but in any one trial, a slower, less capable robot may find pucks very close to the goal and deliver them in little time, thereby gaining a higher performance value than a more capable robot finding pucks further away. Similar to the fault detection experiments, the AMMs in these experiments were also generated on a PowerMac with data transmitted to and from the robots via radio modem. Lost data packets and noisy signals, in addition to occasional robot failures, also hindered optimal restructuring. With all of these complications, the success of the Dynamic Leader Selection experiments is an indication of the potential robustness of AMMs.

Figure 3 (Right) and Table 3 present the average number of pucks that each robot collected in each of the three experiments. We note by a comparison with Figure 3 (Center) that in all three experiments, the number of pucks a robot collected is proportional to its final position in the hierarchy. Intuitively, this is consistent with the notion of a dominance hierarchy: the less dominant a robot is, the less often it will be allowed to deliver a puck, since it must wait to be the most dominant individual ready to deliver a puck. In a large hierarchy, this may never happen.

There are two slight anomalies in the puck data that, although not significant, bear some consideration. In the Control experiments, Robot3 (the most dominant and slowest) did not collect as many pucks as Robot2, although one might expect it to collect more. One explanation for this inconsistency is that the performance of Robot3 degraded more than could be compensated for by its high position in the hierarchy. Robot3 was simply poor at finding pucks. There is a complementary inconsistency in the DLS data where Robot0 (the fastest, and on average most dominant at the end of the trials) collected fewer pucks than Robot1. This seems due to the fact that Robot0 began the DLS trials as the least dominant robot, and consequently at times was not allowed to deliver a puck and establish its performance value until relatively late in a trial.

Finally, it is worth noting that none of the results for the Optimal experiments were significantly different from the dynamic leader selection experiments. One implication of this result (and some of the others mentioned above) is that in order to test the significance of minor differences in the data gathered, additional, possibly numerous, trials would be necessary — likely an impossibility due to time and robot robustness constraints. The complementary implication is that, on average, dynamic leader selection using AMMs yields performance virtually identical to a prespecified optimal hierarchy, even starting with the most (or any) sub-optimal hierarchy.

6 Related Work

We confine our review to the most relevant work applied to the mobile robotics domain. Various models have been employed on mobile robot platforms to date, a few examples of which we consider here. [3] uses a partially observable Markov decision process (POMDP) to model uncertainty of location in a robot navigation task for an office environment. This example of model construction on, and use by, a mobile robot is similar to our effort. The high computational complexity and large data requirements for POMDP generation, however, made it necessary to use a number of heuristics to reduce the problem size and facilitate learning, which is not the case for our approach.

In more recent work, [10] uses a POMDP to model sensor, actuator, and metric uncertainties in a similar office navigation task. For tractability, the learning system on the robot is initialized with a POMDP compiled off-line using sensor models and a topological map with metric uncertainties. A modified POMDP learning algorithm is used to passively fine-tune the probability distributions. Both POMDPs and AMMs represent uncertainties using probability distributions, but POMDPs also learn control policies. Such policies may require sensor and environment state information, making the search space large and requiring more (possibly heuristic) computation. AMMs do not specify a control policy, but passively capture the interaction dynamics between a robot and the environment at the level of behavioral activity. In addition, the number of states in a POMDP is pre-specified, whereas in an AMM it is learned.

[13] also presents a learning technique that uses a behavioral model of a robot's interaction with the environment. The model takes the form of a tree capturing the history of behavior use, i.e., specific sequences of behaviors, with nodes representing executed behaviors, and links representing transitions between them. Their approach and our AMMs are both constructive, being generated and modified as needed. Unlike AMMs, which can form arbitrary graphs with alternative behavior choices represented by probability distributions over transitions, their approach stores potentially long linear sequences of behaviors in a tree with branches indicating alternative choices. The result is that the tree representation may require many trials to collect useful statistics, whereas AMMs can be generate them during the course of one trial.

[11] presents Discrete Event Systems (DES) with emphasis on applications to mobile robotics. The structure of this DES approach is also related to AMMs, both being represented as directed graphs with behaviors as states. Unlike AMMs, DES require *a priori* specification of all possible states, events, and the full transition function. This specification endows DES with control theoretic properties, though a practical specification of these parameters for mobile robots would seem to require heuristic engineering.

The basic structure of augmented Markov models is very similar to that of hidden Markov models (HMMs) [14]. The difference between the two models is that in the AMM, it is assumed that the state of the system is known, thereby removing the hidden state assumption. We set each state to only accept/generate a particular unique symbol, and thereby significantly reduce the computational complexity.

The notion of splitting states based on local estimates of mean and variance is related to work by Hanson on the stochastic delta rule, used in updating mean and variance estimates on the weights of feed–forward neural networks, and Meiosis Networks, which use these mean and variance estimates to decide when to split nodes in the hidden layer [9]. Other analogous approaches to node splitting have also been explored. In general, the AMM approach is related to non-parametric methods, in which a parameterization of the problem is not specified *a priori*. Other work in modeling the interaction between an agent (robot) and its environment has used a dynamical systems approach [1, 15].

7 Conclusions and Future Work

We have presented augmented Markov models (AMMs) as a parsimonious, computationally inexpensive real-time approach to capturing and utilizing the behavior dynamics generated by a mobile robot's interaction with the environment. The AMM approach has low computational and space overhead, and can generate and utilize models in real-time with only small amounts of data. These characteristics make it particularly amenable to the mobile robotics domain. We have also presented examples of model construction and use for achieving group coordination at three levels (individual performance, group affiliation, and group performance) and demonstrated reliable fault detection, experience-based or ability-based membership determination, and dynamic leader selection. Our immediate goal is to explore further robotics applications of AMMs, one possibility being regime detection: the robot's ability to detect and adapt to fundamental changes in its environment, including its interactions with other agents.

Acknowledgments The research reported here was conducted at the Interaction Lab, part of the Robotics Research Lab at the University of Southern California Computer Science Department. The work is supported by the Office of Naval Research Grant N00014-95-1-0759 and the National Science Foundation Infrastructure Grant CDA-9512448.

References

[1] Beer, R. D. A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence 72* (1993), 173–215.

[2] Brooks, R. A. The behavior language; user's guide. Tech. Rep. AIM-1227, MIT AI Lab, 1990.

[3] Cassandra, A. R., Kaelbling, L. P., and Littman, M. L. Acting optimally in partially observable stochastic domains. In *Proceedings, AAAI-94* (Seattle, WA, 1994).

[4] Doyle, J., and Sacks, E. P. *Stochastic Analysis of Qualitative Dynamics*. MIT/LCS/TM-418, MIT, December 1989.

[5] Freund, J. E. *Mathematical Statistics*, fifth ed. Prentice Hall, 1992.

[6] Gat, E. On three-layer architectures. In *Artificial Intelligence and Mobile Robotics*, D. Kortenkamp, R. P. Bonnasso, and R. Murphy, Eds. AAAI Press, 1998.

[7] Goldberg, D., and Matarić, M. J. Interference as a tool for designing and evaluating multi-robot controllers. In *Proceedings, AAAI-97* (Providence, Rhode Island, 1997).

[8] Goldberg, D., and Matarić, M. J. Augmented markov models. Tech. Rep. IRIS-99-367, University of Southern California, 1999.

[9] Hanson, S. J. Meiosis networks. *Advances in Neural Information Processing Systems 2* (1990).

[10] Koenig, S., and Simmons, R. G. Unsupervised learning of probabilistic models for robot navigation. In *Proceedings of the IEEE International Conference on Robotics and Automation* (1996).

[11] Košecká, J., and Bajcsy, R. Discrete event systems for autonomous mobile agents. In *Proceedings, Workshop on Intelligent Control* (July 1993).

[12] Matarić, M. J. Behavior-based control: Examples from navigation, learning, and group behavior. *Journal of Experimental and Theoretical Artificial Intelligence 9*, 2–3 (1997), 323–336.

[13] Michaud, F., and Matarić, M. J. Learning from history for behavior-based mobile robots in non-stationary conditions. *Autonomous Robots 5*, 3–4 (Jul/Aug 1998).

[14] Rabiner, L. R. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE 77*, 2 (Feb. 1989), 257–285.

[15] Smithers, T. What the dynamics of adaptive behavior and cognition might look like in agent-environment interaction systems. In *Practice and Future of Autonomous Agents*. Mt. Verita, Switzerland, 1995.