# Learning Multiple Models for Reward Maximization

**Dani Goldberg** DANI@USC.EDU
**Maja J Matarić** MATARIC@USC.EDU
Computer Science Department, University of Southern California, Los Angeles, CA 90089-0781 USA

## Abstract

We present an approach to reward maximization in a non-stationary mobile robot environment. The approach works within the realistic constraints of limited local sensing and limited *a priori* knowledge of the environment. It is based on the use of *augmented Markov models* (AMMs), a general modeling tool we have developed. AMMs are essentially Markov chains having additional statistics associated with states and state transitions. We have developed an algorithm that constructs AMMs on-line and in real-time with little computational and space overhead, making it practical to learn multiple models of the interaction dynamics between a robot and its environment during the execution of a task. For the purposes of reward maximization in a non-stationary environment, these models monitor events at increasing intervals of time and provide statistics used to discard redundant or outdated information while reducing the probability of conforming to noise. We have successfully implemented this approach with a physical mobile robot performing a mine collection task. In the context of this task, we first present experimental results validating our reward maximization criterion in a stationary environment. We then incorporate our algorithm for redundant/outdated information reduction using multiple models and apply the approach to a non-stationary environment with an abrupt change. Finally, we apply the technique to a simulated version of the task with a gradually shifting environment.

## 1. Introduction and Motivation

In certain classes of mobile robot tasks, a robot may be required to perform optimally with respect to the information it possesses about the structure of its environment. Reward maximization may be used as a means of quantifying performance. In that framework, the robot receives reward (e.g., points) in proportion to its performance. Reward maximization in a non-stationary environment requires the robot to be able to estimate the state of the changing environment. There are a number of issues that can compound the difficulty of this problem:

- The robot may have no *a priori* knowledge of the environment and thus also lack a baseline for gauging the non-stationarity of the environment.
- Given only local sensing capabilities, the robot may require a significant amount of time to estimate the state of the environment. Any estimate of state, however, may be outdated in a non-stationary system.
- The nature of the task may be stochastic, with uncertainties large enough to preclude an effective predictive model of environmental state, or dynamics too complex to make the development of such a model feasible or tractable. Alternatively, however potentially simple the system, there may be no *a priori* data with which to instantiate a model.
- Further, in a stochastic system, the variability associated with performing a task (or elements thereof) may be enormous and effectively mask gradual shifts in the environment. Conversely, in a system with very low variability, even minute shifts may be easily detected. Thus, effective estimation of environmental state requires an understanding of the system's variability (as often measured by variances, covariances, etc.).
- Depending on the task or environment, the time scale at which the non-stationarity manifests and thus can be detected may differ. For example, in one task, the environmental change may be almost instantaneous, detectable between one moment and the next. In another task, the change may be slow and incremental, requiring the examination of a large time interval for detection. Hard-coding the robot with a specific time scale to use for state estimation can be problematic. A time scale that is too small makes the robot incapable of detecting the change. Conversely, a time scale that is unnecessarily large increases the time required to detect the change and may be undesirable in time-critical situations.

As a concrete example, consider the task of collecting undetonated land mines in a field. Assume that there are two types of mines, large and small, with destructive power proportional to their size. The robot's goal is to minimize the total destructive power of the mine field as much as possible during a given period of time. When the robot is given points in proportion to the destructive power of the mines it collects, the goal becomes equivalent to reward maximization. To accomplish its goal, the robot must have enough data about its environment (the field) to intelligently decide whether it is best to collect large mines or small ones at each point in time. The difficulty of this task is compounded when the issues mentioned above apply. The task is likely stochastic, with unknown variability. The robot may have no *a priori* information about the numbers of large and small mines in the field, their distributions, or relative proportions. The robot may also lack global sensing of the mines in the field. These limitations relegate the robot to estimating the environmental state while performing the task. With only an estimate, however, the robot may not perform in a globally optimal manner. The heart of this problem, therefore, is to use the best possible estimate of environmental state given the limitations of the system.

If the task environment is stationary, then all of the data the robot gathers may be used to estimate the state, with more data presumably providing a better estimate. Conversely, for non-stationary environmental state estimation, some mechanism must exist for discarding old data. This is a tricky proposition. If too much data are discarded, the estimate may be susceptible to noise and variance; if too little are discarded, the estimate may be skewed and not accurately represent the current state. This is analogous to the issues of overfitting and underfitting generally encountered in machine learning (Mitchell, 1997). The appropriate amount of data to be kept is not necessarily static and pre-determinable, but rather, depends on the variances of the system and the type of non-stationarity exhibited. Low variances require less data (i.e., a smaller time scale) to characterize, as does non-stationarity exemplified by abrupt shifts. Both high variances and gradually shifting non-stationarity require greater amounts of data (i.e., a larger time scale) to characterize. Thus, a mechanism for estimating environmental state must accommodate both the variances and the type of non-stationarity exhibited by the system. Additionally, since multiple types of relevant non-stationarity and variances may exist in the system, a state estimation procedure capable of dynamically compensating is desirable.

In this paper, we propose an algorithm that provides a moving average estimate of the state of a non-stationary system. The algorithm dynamically adjusts the window size used in the moving average to accommodate the variances and type of non-stationarity exhibited by the system, while discarding outdated and redundant data. Our focus in this paper is the application of the algorithm to the problem of reward maximization in a non-stationary environment. The algorithm uses *augmented Markov models* (AMMs), and builds upon our previous work with AMMs (Goldberg & Matarić, 1999b) to demonstrate another, potentially concurrent, application of the model. Multiple AMMs are learned, capturing in real time the dynamics of a robot interacting with its environment in terms of the behaviors it performs. One AMM is learned and maintained at each time scale that is monitored, and statistics about the environment at that time scale are derived from it. The state of the environment is thus estimated indirectly though the robot's interaction with it. As task execution continues, AMMs are dynamically generated to accommodate the increasing time intervals. Sets of statistics from the models are used to determine whether old data and AMMs are redundant/outdated and can be discarded. This approach requires no *a priori* knowledge, uses only local sensing, and captures the notion of time scale. Additionally, it works naturally with stochastic task domains where variations between trials may change the most appropriate amount of data for state estimation.

In the next section, we present some details of AMM structure. We follow with a discussion of Markov chain theory applicable to AMMs and derive the statistics that will be used for reward maximization and the dynamic moving average algorithm we subsequently present. We validate our approach using an implementation of the mine collection task with a real mobile robot and in simulation.

## 2. Structure of AMMs

An augmented Markov model (AMM) is a deterministic finite state automaton having probabilities associated with transitions between states, i.e., a Markov chain (Roberts, 1976). The Markov chain is expanded with data used for dynamic model construction, modification, and exploitation.

An AMM model is defined as a five-tuple $\langle S, A, B, L, T \rangle$:

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 state $a_i$ has four attributes:

- $a_i^s$, the symbol that the state recognizes, i.e., an element of $S$;
- $a_i^\mu$, the average number of time steps that the system remains in $a_i$ whenever it enters that state;
- $a_i^{\sigma^2}$, the variance associated with $a_i^\mu$;
- $a_i^p$, the probability of remaining in $a_i$ in the next time step.

3. $B$, an $N \times M$ transition matrix.

4. $L$, a set of directed links $\{l_1, l_2, \ldots, l_P\}$, connecting the states. Each link $l_i$ includes the following attributes:

- $l_i^f$, indicates the state from which the link begins, $l_i^f \in A$;
- $l_i^t$, indicates the state to which the link connects, $l_i^t \in A$.
- $l_i^p$, the probability of using link $l_i$ from state $l_i^f$.

5. $T$, a set of elements $\{t_1, t_2, \ldots, t_Q\}$, each storing information on a particular two-link traversal sequence entering and leaving a state. This two-link data allows the model to represent second-order Markovian systems.

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. In other words, each AMM state corresponds to a particular robot behavior. The AMM states do not explicitly capture the sensor inputs and internal state of the robot that triggers the execution of a behavior. AMM generation is similar to the generation of Markov chains, but with the complication of additional node data and a statistical node-splitting mechanism for accommodating second-order Markovian systems. Details of AMM generation, usage, and structure are available in Goldberg and Matarić (1999a). 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 learned and modified 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 a 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. Finally, an individual behavior $\mathcal{B}_i$ in the behavior space may be composed of several more primitive behaviors that are executed in parallel.

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 (Matarić, 1997; Gat, 1998). 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 how to derive several AMM statistics used later in our moving average algorithm for state estimation.

## 4. AMM Statistics

We wish to derive several useful statistics from an AMM which we will subsequently use in our reward maximization criterion and our dynamic moving average algorithm using multiple AMMs.

One such statistic is the expected number of time steps the system takes to reach a destination state from a given start state. This is known as the *mean first passage*. The theory of Markov chains provides powerful tools for easily calculating such expectations.

First, the stochastic transition matrix, $P$, of the Markov chain designated by the AMM is extracted using the $a_i^p$ and $l_i^p$. In order to continue with the calculation, $P$ must be regular (i.e., any state must be reachable from any other state in $\kappa$ steps). The requirement of regularity is usually met in repetitive tasks such as mine collection. Next, the fixed point probability vector $w$ indicating the proportion of time spent in each state is calculated from $wP = w$. The fundamental matrix, $Z$, is now calculated as $Z = [I - (P - W)]^{-1}$, where each row of $W$ is $w$. Finally, $Z$ is used in the calculation of the mean first passage matrix $E = (I - Z + JZ_{\text{dg}})D$, where J is a matrix of ones, $Z_{\text{dg}}$ is the main diagonal of $Z$, and $D$ is the diagonal matrix with $d_{ii} = 1/w_i$. Each element, $e_{ij}$, of $E$ gives the mean time (number of steps) to get from state $i$ to state $j$. An associated variance and degrees of freedom can also be calculated for each element in

$E$. A more detailed treatment of Markov chain theory is available (Roberts, 1976; Kemeny et al., 1966).

Given the mean first passage values in $E$ and their associated variances and degrees of freedom, we can perform two standard tests of statistical significance: the $t$-test and $F$-test (Freund, 1992). The $t$-test uses Student's $t$ distribution to determine whether the difference between two means is significant. Calculation of this test requires the two mean values and their combined degrees of freedom derived above. The $F$-test uses the $F$ distribution to determine if the variances of two normal distributions are significantly different given their degrees of freedom. Either of these two tests can be used to compare values for mean first passage and variance within the same AMM and across AMMs. See Press et al. (1992) for more details on the computation of these tests. Goldberg and Matarić (2000) provide more details on statistics for AMMs.

In the next section, we use AMMs and the statistics derived above in our dynamic moving average algorithm.

## 5. Dynamic Moving Average Algorithm

To manage the amount of data used to estimate the state of a non-stationary environment, we present an algorithm which essentially computes a moving average with a dynamic window size. The algorithm maintains multiple AMMs for different time intervals, and uses a $t$-test and $F$-test on comparable values from the different AMMs to determine which AMM provides the best information. These tests allow the algorithm to adjust the window size of the moving average to accommodate both the amount of variance in the system and the type of non-stationarity (ranging from abrupt to very gradual). The window size of the moving average is allowed to grow by maintaining and expanding old AMMs, and is shrunk by deleting AMMs. We now present the algorithm.

1. Let $\mathcal{L}$ be a queue-like list of AMMs, with $\mathcal{L}_0$ as the first element. Initialize $\mathcal{L}$ to contain one AMM.
2. Let $\varpi_t$ and $\varpi_F$ be constants specifying the significance levels for the $t$-test and $F$-test, respectively.
3. Let *mode* be a variable designating the two modes of the algorithm.
4. For each new input symbol do the following:
5. Update each AMM in $\mathcal{L}$ with the new input.
6. If it is time to create a new AMM, then:
7. Create a new AMM and add it to $\mathcal{L}$.
8. Compute the mean first passage matrix for $\mathcal{L}_0$, extract the desired values and calculate their associated variance and degrees of freedom.
9. Do the same for $\mathcal{L}_1$.
10. Perform an $F$-test between the variances calculated for $\mathcal{L}_0$ and $\mathcal{L}_1$.
11. If the significance level returned by the $F$-test is less than $\varpi_F$ (i.e., the variances are different) then set *mode*=1, else set *mode*=2.
12. If *mode*==1, then let $i$ be the index of the first AMM in $\mathcal{L}$ after $\mathcal{L}_0$ that has either significantly different variances or significantly different means (i.e., significance level $< \varpi_F, \varpi_t$). If such an $i$ exists, delete $\mathcal{L}_0$ through $\mathcal{L}_{i-2}$ and use the new $\mathcal{L}_0$ as the best estimate of the state.
13. If *mode*=2, then let $i$ be the index of the first AMM in $\mathcal{L}$ after $\mathcal{L}_0$ that has neither significantly different variances nor means (i.e., significance levels $> \varpi_F, \varpi_t$). If such an $i$ exists, delete $\mathcal{L}_0$ through $\mathcal{L}_{i-1}$ and use the new $\mathcal{L}_0$ as the best estimate of the state.
14. If no such $i$ exists for either value of *mode*, then do not delete any AMMs and use the current $\mathcal{L}_0$ as the best estimate.

There are several characteristics of the algorithm worth noting. The decision criterion used to create a new AMM (line 6) is very general. For example, a new AMM might be created after a certain period of time, a certain number of input symbols, or when a particular input symbol is observed. In the experiments described below, a new AMM is created every time the robot finds an object (puck) to collect.

The algorithm adjusts the amount of data in the moving average to accommodate the variance in the system. This is accomplished by considering deletion of $\mathcal{L}_0$, the AMM representing the largest time window of data, only when its variance is comparable to (i.e., not significantly different from) that of another AMM. When the variability in the system is high, the AMMs require more data (i.e., larger windows) to acquire comparable variances. When variability is low, less data are required to accurately characterize the variances of the system.

The algorithm also has two distinct modes. In the first mode (*mode*=1), the algorithm removes redundant/old data. In systems with very gradual non-stationarity, this mode effectively maintains a good moving average estimate of the state. When there is an abrupt change in the system, the means and variances of adjacent AMMs may become increasingly different as more data are collected, causing the first mode to stall (i.e., not delete old AMMs). The second mode (*mode*=2) solves this problem by comparing non-adjacent AMMs to find two that are not significantly different. The second mode then "jumps over" the intervening AMMs by deleting them.

One final item of note is the importance of the $\varpi_t$ and $\varpi_F$ thresholds for significance level. Both values must be in the interval $[0, 1.0]$, with a significance level of 0.05, or less, generally considered significant. The effect of $\varpi_t$ and $\varpi_F$ is to adjust the size of the moving average window. Extremely large values of both thresholds produce a very large window with excessive smoothing and a potentially skewed estimate of state. Very small values of the thresholds result in a small window and state estimation that is prone to overfitting. Empirical tests suggest that values in $[0.01, 0.1]$ tend to work fairly well in the algorithm, with relatively little sensitivity. It should also be noted that the experiments described later use this algorithm in

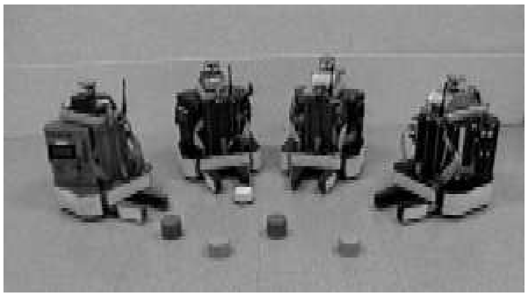

*Figure 1.* Four IS Robotics R2e robots used individually in the experiments.

real time.

In the next section we describe our experimental mine collection task and the associated behavior space.

## 6. Mine Collection Task

To validate our dynamic moving average algorithm and its application to reward maximization, we use a task analogous to the mine collection example of the Introduction. The first two experiments presented are implemented on a real mobile robot.

### 6.1 The Robot

We use one IS Robotics R2e robot (Figure 1) in the experiments. The robot 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. The robot is programmed in the Behavior Language (Brooks, 1990). Experiments are performed in an 11 $\times$ 14 foot rectangular enclosure (the Corrall). The Corrall had up to 36 small plastic cylinders (pucks) of two different colors: clear (representing large mines) and black (representing small ones). The pucks were evenly distributed throughout the Corrall, except in the drop-off area, called Home, a ninety degree sector of a circle with a radius of two feet, located in one corner (Figure 2).

### 6.2 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 other behaviors, and sending outputs to actuators and other behaviors. A set of these behaviors can often trivially map to the behavior space for an AMM, as is the case with our experimental mine collection task. The following 9 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 gripper 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**.
- **color detecting**: if **puck detecting** is successful, detect the color of the puck. If it is the desired color, then perform **homing**, else perform **leave puck**.
- **leave puck**: drop the puck and continue searching for more, using **avoiding**, **wandering** and **puck detecting**.
- **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.
- **reverse homing**: move away from the Home region.

The **color detecting**, **homing**, **avoiding**, and **creeping** behaviors are all qualified to indicate the color of puck the robot has found. Control of the robot's drive motors was the basis for selecting the constituent members of this behavior space. 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. We now present the validation experiment and results for the reward maximization criterion.

## 7. Experiment 1: Validation of the Reward Maximization Criterion

Before demonstrating reward maximization in a non-stationary version of the mine collection task, we first validate the reward maximization criterion in a stationary environment. This validation is necessary to ensure that our subsequent results are not biased by an invalid assumption about the value of reward maximization. If it were shown that our reward maximization criterion did not improve performance over random behavior, then the utility of our dynamic moving average algorithm in the non-stationary version of the task would be suspect.

We now more formally define the reward maximization criterion. Let $\mathcal{R}_s$ and $\mathcal{R}_l$ be the rewards for small and large mines, respectively. Let $\tau_s^f$ be the expected time required for the robot to find a small mine, and let $\tau_s^d$ be the expected time to deliver the mine to the goal location once it has been found. Similarly, $\tau_l^f$ and $\tau_l^d$ represent these times for large mines. The robot maximizes its reward by deciding for each mine found whether to deliver it or leave it in search of a higher valued mine. The action chosen is the one that maximizes the expected reward per unit time (and thus the overall expected reward). If the robot finds a small mine, and the inequality

$$\frac{\mathcal{R}_s}{\tau_s^d} > \frac{\mathcal{R}_l}{\tau_l^f + \tau_l^d}$$

holds, then delivering the small mine maximizes reward. Otherwise, the small mine should be left and a large mine sought. The complementary inequality is used when a large mine is found.

The main issue in evaluating the inequality is calculating $\tau^d$ and $\tau^f$ for each mine type. One could maintain internal variables that record these values. Our approach, however, is to calculate these values from the robot's AMM. As discussed previously, each element, $e_{ij}$, of $E$ gives the mean first passage from state $i$ to state $j$. $E$ also contains the values for $\tau^f$ and $\tau^d$. $\tau^f$ is simply the entry in $E$ associated with the minimum mean time from a **wandering** state to a **puck color** state, and $\tau^d$ is the minimum mean time from a **puck color** state to a **wandering** state. Since these states are qualified by the color of puck the robot possesses, $\tau_s$ and $\tau_l$ can be distinguished.

We performed two sets of experiments in this validation: one control set where the robot collected both types of pucks without discrimination, and one set with reward maximization allowing the robot to decide which pucks to collect. For these experiments, reward values were set in proportion to a mine's explosive power, thus making reward maximization identical to minimizing the mine field's destructive potential. The setup for both sets of experiments was identical, with 18 clear pucks (large mines) and 18 black pucks (small mines) evenly distributed in the Corrall. Each clear puck had a reward of 4 points, while each black had a reward of 1 point. In addition, the environment was kept stationary by replenishing collected pucks. We performed five one-hour trials for each experiment. Using the reward maximizing criterion, the robot accrued an average of 46.6 points, while without reward maximization the robot averaged 37 points. A hypothesis test based on Student's $t$ indicates that the means are different at a significance level of 5%, and validates our reward maximization criterion.

We now present experimental results applying the dynamic moving average algorithm to an abruptly changing non-stationary version of this task.

## 8. Experiment 2: Abruptly Changing Environment

In this set of experiments, we aim to show that, in an abruptly changing non-stationary version of the mine collection task, reward maximization with the addition of the dynamic moving average algorithm is superior to reward maximization alone. The hypothesis is that when the environment changes, average values of $\tau^f$ and $\tau^d$ become inaccurate and thus not effective for reward maximization. Using a dynamic moving average allows quicker adaptation to change.

The environment was first initialized to contain only 18 clear pucks. We ran one trial of approximately 20 minutes in this environment, during which the robot collected 4 clear pucks. This result is extremely variable: the actual time to collect 4 pucks could easily range between 10 and 30 minutes. In order to reduce the variability, the data from this one trial were used as a primer for all of the subsequent experiments, in which the clear pucks (large mines) were replaced with black (small) ones. Doing so allowed us to focus on the key experimental parameter: the time required to adapt to the new environment. We considered the robot to have adapted to the new environment when it consistently began collecting black pucks.

We ran two experiments using reward maximization with the dynamic moving average algorithm and three experiments using reward maximization alone. In these experiments, the reward for a clear puck was 7 points and the reward for a black puck was 1 point. For reward maximization alone, the mean time to adaptation was 47 minutes, while the mean time with the algorithm was 18.3 minutes. A $t$-test indicates that these means are different at a significance level of 0.01. This strongly supports our hypothesis that the dynamic moving average algorithm allows quicker adaptation to abrupt non-stationarities.

## 9. Experiment 3: Gradually Shifting Environment

In this set of experiments, we test the effectiveness of the dynamic moving average algorithm for reward maximization in a gradually shifting environment. Instead of abruptly changing the color of the pucks, as in the previous experiment, here the environment shifts slowly as the robot collects pucks. Due to the high degree of variance in the mine collection task using physical robots, we anticipated that the number of experiments we would have to conduct in order to obtain statistical significance in the gradually shifting environment would pose a practical impossibility. The experiments described in this section were therefore conducted in a simulation of the mine collection task.

The simulation was initialized to contain 18 clear pucks worth 10 points each, and 18 black pucks worth 1 point each (Figure 2). Three experimental scenarios

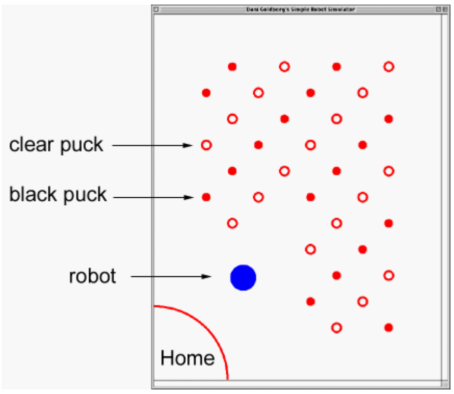


*Figure 2.* The simulated mine collection task: setup for reward maximization in a gradually shifting non-stationary environment.

were examined:

1. **random**: the robot collects any puck it encounters regardless of the color.

2. **control**: the robot uses the reward maximization criterion, but does not employ the dynamic moving average algorithm used to compensate for non-stationarity.

3. **algorithm**: the robot uses both reward maximization and the dynamic moving average algorithm for state estimation.

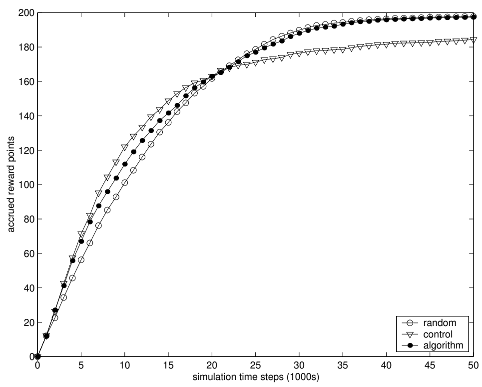


*Figure 3.* Average accrued reward for the three experimental scenarios.

We conducted trials of 50,000 simulation steps for each scenario: 200 trials of the **random** scenario, 40 of **control**, and 100 of **algorithm**, with the actual number determined by the desired level of statistical significance. The data gathered included the time at which each puck was collected, allowing us to calculate the accrued number of reward points. Figure 3 presents the average number of reward points accrued at 1000-time-step intervals for each of the three scenarios. The maximum possible accrued reward is 198 points, corresponding to the collection of all 36 pucks. Both the **random** and **algorithm** scenarios essentially reach this maximum by the end of each trial (with average accrued points of 197.8 and 197.5, respectively). The **control** scenario outperforms the other two until around 20,000 time steps, then quickly declines ending with an average reward of 184.4. The discrepancy between the **algorithm** and **control** cases illustrates the importance of eliminating outdated information, and the effectiveness of our algorithm in doing so. As a further comparison, we calculate the number of reward points that the robot is expected to have accrued on average during the course of a trial: 151.6 for **random**, 149.3 for **control**, and 153.8 for **algorithm**. These data are significantly different at p-value$< 0.02$ using a two-tail version of Student's $t$ test. The superiority of the **algorithm** case illustrates the effectiveness of our moving average algorithm for reward maximization in a gradually shifting environment.

It should be noted that the **random** scenario used above lies along a continuum of possible experiments distinguished by the probability with which the robot collects (or discards) each puck it encounters. For comparison with the **control** and **algorithm** scenarios, we used the random case in which the robot collects 100% of the pucks it encounters. Intuitively, one would expect this to be the best performing of the random cases. When the probability of collecting a found puck is less than 1.0, the robot wastes more time searching for pucks but does not improve its reward since it discards both high-valued and low-valued pucks with equal probability. To verify our intuition, we conducted 80 trials for each of three lower probabilities of collection: 0.75, 0.50, 0.25. As expected, the data show that there is a continual significant decrease in accrued reward as the probability of collection decreases. The random cases with lower collection probabilities might seem to be more appropriate for comparison with the **algorithm** and **control** scenarios. It is, however, the collection probability of 1.0 that is the most difficult challenge for our algorithm. We thus used this version in the comparison experiments described here as well as those in Section 7.

Finally, we have conducted experiments demonstrating that when the point values of the pucks are closer together (black=1, clear=4) and thus the difference in reward for collecting the "wrong" versus the "right" colored puck is very small, our algorithm still significantly outperforms the **random** and **control** cases. In these experiments, we compared data from 140 trials of the **control** and **algorithm** scenarios, and 200 trials of the **random** scenario. All trials were run to 50,000 simulation steps.

## 10. Related Work

We confine our review to a small selection of relevant work. Various model learning approaches have been employed on mobile robot platforms to date. Cassandra et al. (1994) used a partially observable Markov decision process (POMDP) to model uncertainty of location in a robot navigation task for an office environment. Koenig and Simmons (1996) employed a POMDP to model sensor, actuator, and metric uncertainties in a similar office navigation task. Michaud and Matarić (1998) presented a learning technique that uses a tree-like behavioral model of a robot's interaction with the environment. Košecká and Bajcsy (1993) presented Discrete Event Systems (DES) with emphasis on applications to mobile robotics. Unlike much of the above work which requires multiple trials to achieve convergence, the focus in our AMM work in mobile robotics is real-time response, i.e., results within one trial.

Related to our approach in this paper is other work in regime detection, and information representation at different time scales. Weigend et al. (1995) used a network of nonlinear gated experts to detect regime changes in a time series. Sericola (1996) considered stationary regime detection in a finite Markov process as part of the larger problem of availability analysis in a computer system. Sutton et al. (1998) discussed the importance of representing knowledge at multiple time scales in planning and decision making, and used an approach based on Markov decision processes.

## 11. Conclusions

In this paper we have presented an algorithm for moving average state estimation in a non-stationary environment. The algorithm learns multiple augmented Markov models (AMMs) and uses statistics from these to dynamically adjust the window size of the moving average. It requires no *a priori* knowledge of the environment, and is designed to compensate for both the variability and the manifested types of non-stationarity (from abrupt to gradual shifts) in a stochastic system. We have applied this algorithm to the problem of reward maximization in a situated mine collection task for robots, and shown that it improves performance in environments with abrupt and gradual non-stationarities.

## 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 was supported by the Office of Naval Research Grant N00014-95-1-0759.

## References

Brooks, R. A. (1990). *The behavior language; user's guide* (Technical Report AIM-1227). MIT AI Lab, Cambridge, MA.

Cassandra, A. R., Kaelbling, L. P., & Littman, M. L. (1994). Acting optimally in partially observable stochastic domains. *Proceedings of The Eleventh National Conference on Artificial Intelligence* (pp. 1023–1028). Menlo Park, CA: AAAI Press.

Freund, J. E. (1992). *Mathematical statistics.* Upper Saddle River, NJ: Prentice Hall. Fifth edition.

Gat, E. (1998). On three-layer architectures. In D. Kortenkamp, R. P. Bonnasso and R. Murphy (Eds.), *Artificial intelligence and mobile robotics.* Menlo Park, CA: AAAI Press.

Goldberg, D., & Matarić, M. J. (1999a). *Augmented markov models* (Technical Report IRIS-99-367). University of Southern California, Los Angeles, CA.

Goldberg, D., & Matarić, M. J. (1999b). Coordinating mobile robot group behavior using a model of interaction dynamics. *Proceedings of The Third International Conference on Autonomous Agents* (pp. 100–107). New York: ACM Press.

Goldberg, D., & Matarić, M. J. (2000). Reward maximization in a non-stationary mobile robot environment. *Proceedings of the Fourth International Conference on Autonomous Agents.* New York: ACM Press.

Kemeny, J. G., Snell, J. L., & Knapp, A. W. (1966). *Denumerable markov chains.* Princeton, NJ: Van Nostrand.

Koenig, S., & Simmons, R. G. (1996). Unsupervised learning of probabilistic models for robot navigation. *Proceedings of the IEEE International Conference on Robotics and Automation* (pp. 2301–2308). Piscataway, NJ: IEEE Press.

Košecká, J., & Bajcsy, R. (1993). Discrete event systems for autonomous mobile agents. *Proceedings of the Symposium on Intelligent Robotic Systems* (pp. 21–31).

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

Michaud, F., & Matarić, M. J. (1998). Learning from history for behavior-based mobile robots in non-stationary conditions. *Autonomous Robots*, *5*, 335–354.

Mitchell, T. M. (1997). *Machine learning.* Boston: McGraw-Hill.

Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. (1992). *Numerical recipes in c: The art of scientific computing.* New York: Cambridge University Press.

Roberts, F. S. (1976). *Discrete mathematical models: With applications to social, biological, and environmental problems.* Englewood Cliffs, NJ: Prentice-Hall.

Sericola, B. (1996). *Availability analysis and stationary regime detection of markov processes* (Technical Report 2886). INRIA, France.

Sutton, R. S., Precup, D., & Singh, S. (1998). *Between mdps and semi-mdps: Learning, planning, and representing knowledge at multiple temporal scales* (Technical Report 98-74). Department of Computer Science, Univeristy of Massachusetts, Amherst, MA.

Weigend, A. S., Mangeas, M., & Srivastava, A. N. (1995). Nonlinear gated experts for time series: Discovering regimes and avoiding overfitting. *International Journal of Neural Systems*, *6*, 373–399.