Reward Maximization in a Non-Stationary Mobile Robot Environment

Dani Goldberg and Maja J Matarić

Robotics Research Laboratories
University of Southern California
Los Angeles, California 90089-0781

{dani, mataric}@usc.edu

ABSTRACT

In this paper, we present an approach to reward maximization in a non-stationary mobile robot environment. The approach works within the constraints of limited local sensing and limited *a prior* knowledge of the environment. It is based on the use of *augmented Markov models* (AMMs), which 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 maintain 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these models monitor events at increasing intervals of time and provide statistics used to discard redundant or outdated information while reducing the probability of conforming to noise. This approach has been successfully implemented with a real 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 We then incorporate our algorithm for redundant/outdated information reduction using multiple models and apply the approach to a non-stationary environment with an abrupt change.

1. INTRODUCTION AND MOTIVATION

In certain classes of mobile robot tasks, it may be necessary for a robot to perform optimally with respect to the information it possesses about the structure of its environment. Reward maximization may be used as one quantification of such optimality in which 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 environment as it changes. 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 Any estimate of state, however, may be outdated in a non-stationary system.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 there of) 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 goal of the robot 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 [13]. 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, it is desirable that the state estimation procedure be able to dynamically compensate.

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 [7] to demonstrate another, potentially concurrent, application of the model. Multiple AMMs are used to capture, in real time, the dynamics of a robot interacting with its environment in terms of the behaviors it performs. One AMM is created 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 details of AMM structure and the AMM construction algorithm. 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.

2. STRUCTURE OF AMMS

An augmented Markov model (AMM) is a deterministic finite state automaton with probabilities associated with transitions between states, i.e., a Markov chain [15]. 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^μ, 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^μ;
- and a_i^p, the probability of remaining in a_i in the next time step.

A special state, $a_\emptyset$, represents the initial state of the system and possesses none of the four attributes above.

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 $a_i^s = s_k$, then $b_i(k) = a_i$, i.e., if the observed symbol is identical to the last symbol observed, then the system remains in the current state.

4. L, a set of directed links $\{l_1, l_2, \ldots, l_P\}$, connecting the states. Each link l_i has the following six 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, $l_i^t \in A$. The following constraints apply: a link cannot start and end at the same state, $l_i^f \neq l_i^t$; and two links from the same state cannot go to states that accept the same symbol, $\forall i, j$ s.t. $l_i^f = l_j^f$, $a^s_{l_i^t} \neq a^s_{l_j^t}$;
- l_i^δ, stores the number of times the link l_i has been traversed;
- l_i^Σ, stores the total number of time steps that the system has been in state l_i^t, after first having traversed the link l_i;
- $l_i^{\Sigma^2}$, contains the sum of squares of all the durations that comprise l_i^Σ;
- and l_i^p is the probability of using the link l_i at each time step, given the system is in state l_i^f.

Because no two links can have the same value for both their *from* and *to* attributes, they cannot represent the same directed transition. Thus, $N-1 \leq P \leq N(N-1)$: at least

$N-1$ links are needed to connect the non-initial states, and for a fully connected network there are $N(N-1)$ links between the non-initial states. The link from $a_\emptyset$ is a special link designated by $l_\emptyset$, with $l_\emptyset^f = a_\emptyset$. $l_\emptyset$ is traversed exactly one time, and thus $l_i^\delta = 1$ and we may set $l_i^p = 0$.

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. Each element t_i has five attributes:

- t_i^e, the link entering the state, where $t_i^e \in L$, except that there is exactly one t with $t^e = \emptyset$;
- t_i^l, the link leaving the state, where $t_i^l \in L$ and $t_i^e \neq t_i^l$;
- t_i^δ, the number of times the two-link combination has been traversed;
- t_i^Σ, the total number of time steps that the system has been in the state that link t_i^l connects to, after first having traversed t_i;
- $t_i^{\Sigma^2}$, the sum of the squares of all the durations that comprise t_i^Σ.

The bounds of Q are given by:

$$\left.\begin{array}{ll} 0, & \text{if } P < 2 \\ P-1, & \text{if } P \geq 2 \end{array}\right\} \leq Q \leq N(N-1)^2.$$

In order for a two-link transition to exist there must be at least two links. If more than two links exist, the fewest two-link transitions ($P-1$ of them) are created when an Euler path exists and is followed through the network. In a fully connected network, each of the $P = N(N-1)$ links has a transition to $N-1$ other links, giving us the upper bound. Additionally, there is one special two-link transition, designated $t_\emptyset$, with $t_\emptyset^e = l_\emptyset$ and $t_\emptyset^\delta = 1$.

2.1 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 at most $O(N^2)$ (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 would probably tend to fall between $O(N)$ and $O(N^2)$ for something less than a fully connected graph.

Further details of AMM generation, usage, and structure are available in [6]. 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.

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 [11, 5]. 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 the statistics from an AMM which we will subsequently use in the reward maximization experiments.

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, and the number of time steps spent in each of the intervening states. This is known as the *mean first passage*. The theory of Markov chains provides powerful tools for easily calculating such expectations. We apply these tools to AMMs, then use the results to calculate two other statistics: the total variance associated with the mean first passage, and the accompanying degrees of freedom. (A more detailed treatment of Markov chain theory, including proofs and derivations, is available in [15] and [8]).

4.1 Mean First Passage

The first step in calculating the mean first passage is to extract the $N \times N$ transition matrix P of an AMM M. P is a matrix such that each element p_{ij} is the probability of transitioning directly to state j given the system is in state i. Each element of the matrix is given by:

$$p_{ij} = \begin{cases} a_i^p, & \text{if } i = j \\ l_k^p, & \text{if } \exists k \text{ s.t. } l_k^f = a_i \text{ and } l_k^t = a_j \\ 0, & \text{otherwise} \end{cases}$$

where a_i, a_i^p, l_k^p, l_k^f and l_k^t are extracted from M.

In order to be able to continue the calculation of the mean first passage, P must represent a *regular* Markov chain. To be regular, every state in the underlying directed graph of P must be able to reach every other state in κ steps. We use a theorem that states if P is *ergodic* (i.e., every state can reach every other state) and it has a non-zero main diagonal, then it is also regular [15]. We can determine that a directed graph is ergodic by showing that it has one strongly connected component (see [3] for an algorithm).

Given that P is regular, we calculate the stationary vector $w = w_1, w_2, \ldots, w_N$ giving the probability of being in each state $a_1, a_2, \ldots, a_N$ of M in the limit. w is found by solving the system of equations $wP = w$ with the constraint that $w_1 + w_2 + \cdots + w_N = 1$. Using w, we calculate the *fundamental matrix* $Z = [I - (P - W)]^{-1}$, where W is the square matrix having w as each row. W now enters the calculation of the mean first passage matrix $E = (I - Z + JZ_{\text{dg}})D$, where J is a matrix of 1's, Z_{dg} is the diagonal matrix containing the main diagonal of Z, and D is the diagonal matrix with $d_{ii} = 1/w_i$. Each element e_{ij} of E represents the mean first passage from state a_i to a_j.

In terms of AMMs, we are more often interested in the mean first passage from an input symbol s_i to another input symbol s_j. Since each input symbol may be recognized by several states in the AMM, the mean first passage between two input symbols may not be unique. In general, if n states recognize s_i and m states recognize s_j, then nm entries in E represent the mean first passage between these symbols. In such a case, our interest is in the *minimum mean first passage between two input symbols* and the corresponding states that give this value in E. For symbols s_i and s_j, let ε_{ij} be this minimum value from E. Let a_α and a_β be the states such that $e_{\alpha\beta} = \varepsilon_{ij}$.

Once we know the states associated with the minimum mean first passage, we can calculate the expected amount of time spent in each state before reaching a_β. To do so, we first convert our regular Markov chain into an absorbing chain with a_β as an absorbing state. Essentially, this means modifying P so that once the system enters state a_β it does not leave it, i.e., $a_\beta^p = 1.0$. The transition matrix P is converted into a new matrix Q containing only transitions to and from the $N-1$ non-absorbing states by simply deleting the β-th row and column from P. Note that $\alpha \neq \beta$. The fundamental matrix N for the absorbing chain is given by $N = (I - Q)^{-1}$, where each element n_{ij} represents the expected amount of time spent in the j-th non-absorbing state, given that the system starts in the i-th non-absorbing state, and

$$e_{\alpha\beta} = \begin{cases} \sum_{j=1}^{N-1} n_{\alpha j}, & if \alpha < \beta \\ \sum_{j=1}^{N-1} n_{(\alpha-1)j}, & if \alpha > \beta. \end{cases}$$

4.2 Variance of Mean First Passage

In order to calculate the variance associated with a value in the mean first passage matrix, we first note that the underlying Markov chain of an AMM may be interpreted as a set of random variables, one for each state of the chain. Generally, it is assumed that these random variables are independent, identically distributed, and follow a normal distribution. Our calculation of variance also assumes independence and a normal distribution, but to be more general it allows non-identical distributions. Given a set of independent random variables $\{X_1, X_2, \ldots, X_N\}$ with each X_i having an associated population mean μ_i and variance σ_i^2, we wish to calculate the mean and variance for the linear combination $Y = c_1X_1 + c_2X_2 + \cdots + c_nX_n$. Basic results from statistics tell us that the mean of Y is $\mu_Y = \sum_{i=1}^{N} c_i\mu_i$

and the variance of Y is $\sigma_Y^2 = \sum_{i=1}^{N} c_i^2 \sigma_i^2$. In our case, we do not know the population means and variances and instead use the sample means and variances calculated for each state of the AMM, i.e., a_i^μ and $a_i^{\sigma^2}$.

Given $e_{\alpha\beta}$, a_α, a_β, and the fundamental matrix N generated by making a_β an absorbing state, we are ready to calculate the variance associated with $e_{\alpha\beta}$. As shown previously, $e_{\alpha\beta}$ is equivalent to the sum of the row of N associated with a_α. Equivalently, $e_{\alpha\beta}$ may be expressed as a linear combination of the a^μ extracted from the AMM:

$$e_{\alpha\beta} = \sum_{i=0}^{N} c_i a_i^\mu$$

where

$$c_i = \begin{cases} n_{\alpha i}/a_\alpha^\mu, & \text{if } \alpha < \beta \text{ and } i < \beta \\ n_{\alpha-1,i}/a_\alpha^\mu, & \text{if } \alpha > \beta \text{ and } i < \beta \\ n_{\alpha,i-1}/a_\alpha^\mu, & \text{if } \alpha < \beta \text{ and } i > \beta \\ n_{\alpha-1,i-1}/a_\alpha^\mu, & \text{if } \alpha > \beta \text{ and } i > \beta \\ 0, & \text{if } i = \beta \end{cases}$$

Now that we have expressed $e_{\alpha\beta}$ as a linear combination of the means of the random variables composing the AMM, we can do similarly for the variance of $e_{\alpha\beta}$:

$$\mathrm{var}(e_{\alpha\beta}) = \sum_{i=0}^{N} c_i^2 a_i^{\sigma^2}$$

with the c_i's calculated as above.

4.3 Degrees of Freedom

The number of degrees of freedom associated with the linear combination $Y = \sum_{i=1}^{N} c_i X_i$ may be expressed as

$$\frac{\left[\displaystyle\sum_{i=1}^{N} \frac{c_i^2\ \mathrm{var}(X_i)}{V_{X_i}}\right]^2}{\displaystyle\sum_{i=1}^{N} \frac{\left[c_i^2\ \mathrm{var}(X_i)/V_{X_i}\right]^2}{V_{X_i}-1}},$$

where V_{X_i} is the number of values used to calculate $\mathrm{var}(X_i)$. This is an expanded version of the formula found in [14]. To apply this to our calculation of the variance of $e_{\alpha\beta}$ we simply use $c_i^2 a_i^{\sigma^2}$ in place of $c_i^2\ \mathrm{var}(X_i)$, with the c_i's calculated as above, and with

$$V_{X_i} = \sum_{\{i \mid l_i^t = a_i\}} l_i^\delta .$$

4.4 The t-Test and F-Test

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 [4]. 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 [14] for more details on the computation of these tests.

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

5. 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 ϖ_t and ϖ_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 ϖ_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.

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

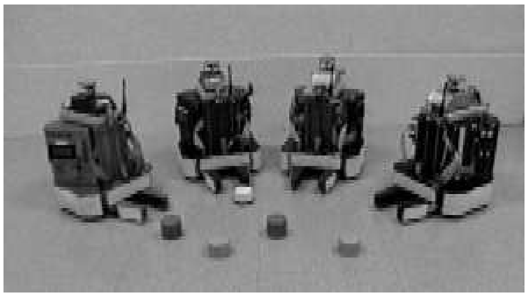

Figure 1: Four IS Robotics R2e robots used individually in the experiments.

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. In the second mode ($mode$=2), the algorithm essentially "jumps over" more abrupt non-stationarities that would otherwise cause it to stall on old data. We present experimental results for this more abrupt form of non-stationarity later in the paper.

One final item of note is the importance of the ϖ_t and ϖ_F thresholds for significance level. Both values must in the interval $[0, 1.0]$, with a significance level of 0.05, or less, generally considered significant. The effect of ϖ_t and ϖ_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 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 experiments 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 [1]. Experiments are performed in an 11×14 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.

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 reflect 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

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

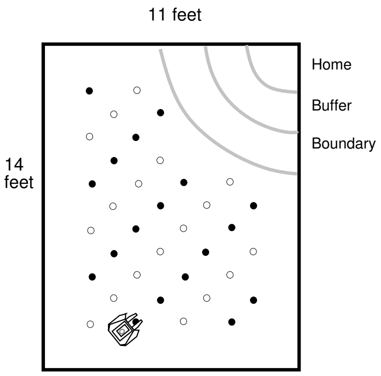

Figure 2: The mine collection task: setup for validation of the reward maximization criterion.

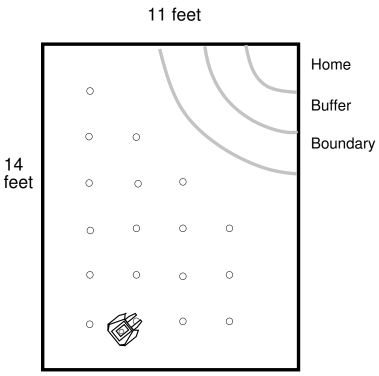

Figure 3: The mine collection task: setup for reward maximization in a non-stationary environment.

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 τ_s^f be the expected time required for the robot to find a small mine, and let τ_s^d be the expected time to deliver the mine to the goal location once it has been found. Similarly, τ_l^f and τ_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 τ^d and τ^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 τ^f and τ^d. τ^f is simply the entry in E associated with the minimum mean time from a **wandering** state to a **puck color** state, and τ^d is the minimum mean time from a **puck color** state to a **wandering** state. In other words, if s_1 is the input symbol for the **wandering** behavior and s_2 is the input symbol for the **puck color** behavior, then $\tau^d = \varepsilon_{21}$ and $\tau^f = \varepsilon_{12}$. Since these states are qualified by the color of puck the robot possesses, τ_s and τ_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 (Figure 2). 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: NON-STATIONARITY

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 τ^f and τ^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 (Figure 3). 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 white 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 white 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. In future work, we plan to explore the efficacy of the algorithm in gradually shifting non-stationary environments.

9. RELATED WORK

We confine our review to a small selection of the most relevant work. Various modeling approaches have been employed on mobile robot platforms to date. [2] uses a partially observable Markov decision process (POMDP) to model uncertainty of location in a robot navigation task for an office environment. [9] employs a POMDP to model sensor, actuator, and metric uncertainties in a similar office navigation task. [12] presents a learning technique that uses a tree-like behavioral model of a robot's interaction with the environment. [10] presents 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. [18] uses a network of nonlinear gated experts to detect regime changes in a time series. [16] considers stationary regime detection in a finite Markov process as part of the larger problem of availability analysis in a computer system. [17] discusses the importance of representing knowledge at multiple time scales in planning and decision making, and use an approach based on Markov decision processes.

10. CONCLUSIONS

In this paper we have presented an algorithm for moving average state estimation in a non-stationary environment. The algorithm uses multiple augmented Markov models (AMMs) 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 of the system and the manifested types of non-stationarity (from abrupt to gradual shifts). We have applied this algorithm to the problem of reward maximization in a mobile robot mine collection task and shown that it facilitates quicker adaptation in an environment with an abrupt change.

11. 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.

12. REFERENCES

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

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

[3] Cormen, T. H., Leiserson, C. E., and Rivest, R. L. *Introduction to Algorithms*. McGraw-Hill Book Company, 1990.

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

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

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

[7] Goldberg, D., and Matarić, M. J. Coordinating mobile robot group behavior using a model of interaction dynamics. In *Proceedings, The Third International Conference on Autonomous Agents (Agents '99)* (Seattle, Washington, May 1999).

[8] Kemeny, J. G., Snell, J. L., and Knapp, A. W. *Denumerable Markov Chains*. D. Van Nostrand Company, Inc., 1966.

[9] 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).

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

[11] 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.

[12] 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).

[13] Mitchell, T. M. *Machine Learning*. The McGraw-Hill Companies, Inc., 1997.

[14] Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, 1992.

[15] Roberts, F. S. *Discrete Mathematical Models: With Applications to Social, Biological, and Environmental Problems*. Prentice-Hall, Inc., 1976.

[16] Sericola, B. Availability analysis and stationary regime detection of markov processes. Tech. Rep. 2886, INRIA, France, May 1996.

[17] Sutton, R. S., Precup, D., and Singh, S. Between mdps and semi-mdps: Learning, planning, and representing knowledge at multiple temporal scales. Tech. Rep. 98-74, Department of Computer Science, Univeristy of Massachusetts, Amherst, MA, April 1998.

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