Representation of behavioral history for learning in nonstationary conditions

François Michaud [a] Maja J Matarić [b]

[a] *Department of Electrical and Computer Engineering, Université de Sherbrooke, Sherbrooke (Québec Canada) J1K 2R1*

[b] *Computer Science Department, University of Southern California, Los Angeles, CA 90089-0781*

Abstract

A robot having to operate in nonstationary conditions needs to learn how to modify its control policy to adapt to the changing dynamics of the environment. Using the behavior-based approach to manage the interactions between the robot and its environment, we propose a method that models these interactions and adapts the selection of behaviors according to the history of behavior use. The learning and the use of this "Interaction Model" are validated using a vision- and sonar-based Pioneer I robot in the context of a multi-robot foraging task. Results show the effectiveness of the approach in taking advantage of any regularities experienced in the world, leading to fast and adaptable specialization for the learning robot.

1 Introduction

A robot having to operate in real-life situations requires the ability to adapt to changing dynamics and nonstationary conditions, i.e., the control policy must change over time [27]. Mobile robots designed using the behavior-based paradigm [3] have shown good performance in adapting and operating in open environments and the approach has been praised for its robustness and simplicity of construction. In a typical behavior-based system, the constituent behaviors are designed parsimoniously, executed in parallel, and prioritized using some fixed or flexible arbitration mechanism [16]. A fixed and parsimonious behavior set, however, does not allow a robot to easily adapt to changes in the environment, thus resulting in diminished flexibility [12]. One possible improvement is to dynamically select behaviors according to external and internal states [20]. The control policy is then composed of behavior-producing modules (or behaviors) for low-level control and a selection mechanism working at a more abstract level, and knowing when to use a behavior becomes as important as the behavior-producing modules themselves. Doing so in nonstationary conditions may require learning a model of the environment that allows the robot to characterize the changing dynamics of the environment and to modify its policy accordingly at the behavior selection level.

However, deriving such model is a very challenging problem. Careful considerations must be made of what to model, what representation to use, and how to use it to select behaviors. Because of nonstationary conditions, these considerations must be as independent as possible of the robot's operating environment and of its own limitations in perception, action and processing. Since performance of situated agents is based on their ability to cope with and exploit the dynamics of interaction with their environment [1,4], it is appropriate to model these interactions using a representation consistent with the control mechanism used, in our case the behavior-based framework. Doing so also preserves two key strengths of the behavior-based approach, situatedness and emergence [4]. Modeling the interaction dynamics with the environment can be especially useful in the case of dynamic multi-agent/robot environments [9] where topological and metric models do not directly apply. These dynamics are difficult for the designer to characterize and fully predict *a priori*, and are thus best captured by the robot itself.

The idea explored in this paper is that interactions between a robot and its environment can be efficiently captured through observation of behavior use over time. The robot is initially given a set of behaviors, whose subset is used to safely accomplish the assigned tasks. The goal of the algorithm is to learn to select alternative behaviors that are not normally exploited in order to adapt the robot's behavior to the experienced dynamics. By knowing the purpose of each behavior, the robot can autonomously evaluate and model the type of interactions it is having with its environment. The representation of the "Interaction Model" is based on the robot's history of behavior use, i.e., on any regularities in the sequences of behavior use that reflect the interactions between the robot and the world [21]. The robot derives this model in a reinforcement learning fashion using a performance criterion also derived from behavior use. The proposed algorithm focuses on continuous life-time adaptation rather than on learning a one-time, static policy. Our approach is validated with a vision- and sonar-based Pioneer I robot in the context of foraging in a multi-robot environment.

This paper is organized as follows. Section 2 describes the experimental setup and the mobile robots used in the experiments, showing how behaviors can be organized to derive an Interaction Model and to change the robot's way of behaving. Section 3 gives a detailed explanation of the learning algorithm. Section 4 presents the experimental results. Section 5 explains the properties of the learning algorithm in relation to nonstationary conditions. Section 6 presents related work, especially from the field of reinforcement learning, and

Section 7 concludes the paper.

2 Experimental setup and task description

Our experiments were performed on a Real World Interface Pioneer I mobile robot (shown on the right of Figure 1) equipped with seven sonars and a Fast Track Vision System. The robot is programmed using MARS (Multiple Agency Reactivity System), a language for programming multiple concurrent processes and behaviors [5]. The experiments were conducted in an enclosed 12'×11' rectangular pen containing pink blocks and a home region, marked with a green cylinder. These colored objects were perceived using the robot's vision system, which suffered from variations in light intensity at different locations in the pen and from different angles of approach. Other obstacles and walls were detected using the sonar system; the dark stripe on the wall of the pen is a strip of corrugated cardboard used to improve sonar reflection. IS Robotics R1 robots, equipped with infra-red and contact sensors, were also used in the experiments; they were perceived as obstacles by the Pioneer's sonars.

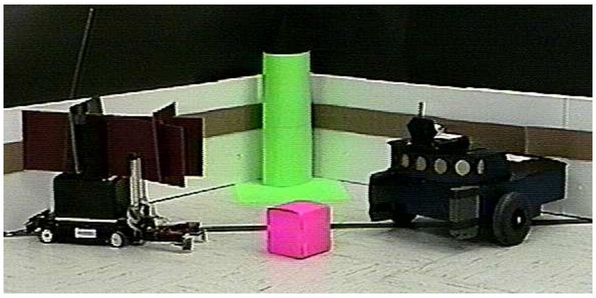

Fig. 1. A Pioneer I robot (on the right) near a block, the home region and a R1 robot.

Foraging was done in two steps: searching for a block (Searching Task), and then bringing it to the home region (Homing Task). To do this, the Pioneer was given a set of behaviors and an initial controller that governed the use of a subset of those behaviors. The behaviors controlled the velocity and rotation of the robot based on sonar readings and visual inputs. To analyze the interactions derived from observing behavior use, behaviors were classified according to their respective purposes, as follows.

First, the robot needed behaviors to accomplish its tasks: one behavior for the Searching Task, called *Searching-block*, and two for the Homing Task, *Homing* and *Drop-block*. A behavior called *Velocity-control* was also used in both of these tasks to make the robot move. All these behaviors are referred to as "task-behaviors". Conditions for activating task-behaviors were pre-programmed based on the presence or absence of a block in front of the robot (using an internal variable indicating if the robot saw a block go underneath its field of view and is pushing it) and the proximity of the robot to home.

Second, an *Avoidance* behavior was required for safe navigation in the environment. It was active unless the robot was near home and carrying a block; then it was disabled in order to let the robot approach the home region. This type of behavior, used for handling harmful situations and interference while accomplishing a task, is referred to as a "maintenance-behavior". The designer determines the conditions in which maintenance-behaviors should be used but cannot indicate when they will occur during the task, as that is tied to the dynamics of the interaction between the robot and its environment.

Third, the robot can also use behaviors that introduce variety into its action repertoire to change the way it accomplishes its tasks. We call these *alternative-behaviors*. In contrast to other types of behaviors, no *a priori* conditions are given for activating alternative-behaviors: the objective is to let the robot learn when to activate these behaviors according to past experiences. In our foraging task, the Pioneer could decide that it was better to activate *Follow-side*, *Rest*, or *Turn-randomly* for some pre-set periods of time when experiencing interference in accomplishing its task.

Figure 2 shows our architecture and the specific components used in our experiments. In the Behavior Level, the behaviors are organized according to the Subsumption Architecture [3] with the difference that the activated behaviors, i.e., those allowed to issue outputs, change dynamically. Behaviors that are not activated cannot participate in the control of the robot. The actual use of a behavior depends on the rules it contains and on commands issued by other behaviors subsuming it. These behaviors were not designed to be mutually exclusive: functionality can emerge by activating multiple behaviors at the same time (e.g., going home by following a wall).

3 History-based learning for adaptive behavior selection

Figure 2 also shows how behavior observation is used to learn the Interaction Model. Using Subsumption, each behavior is assigned a symbol (e.g., F for Follow-side) sent through a wire (for sending messages between between processes) whenever that behavior is executed. This generates a symbol sequence that represents behavior use at each processing cycle, and the sequence is used to construct a tree representation of the robot's interactions with its environment and to activate alternative-behaviors. Separate learning trees are used for each task; determining which one to use is done based on the activated task-behavior. Overall, the learning algorithm is composed of four parts,

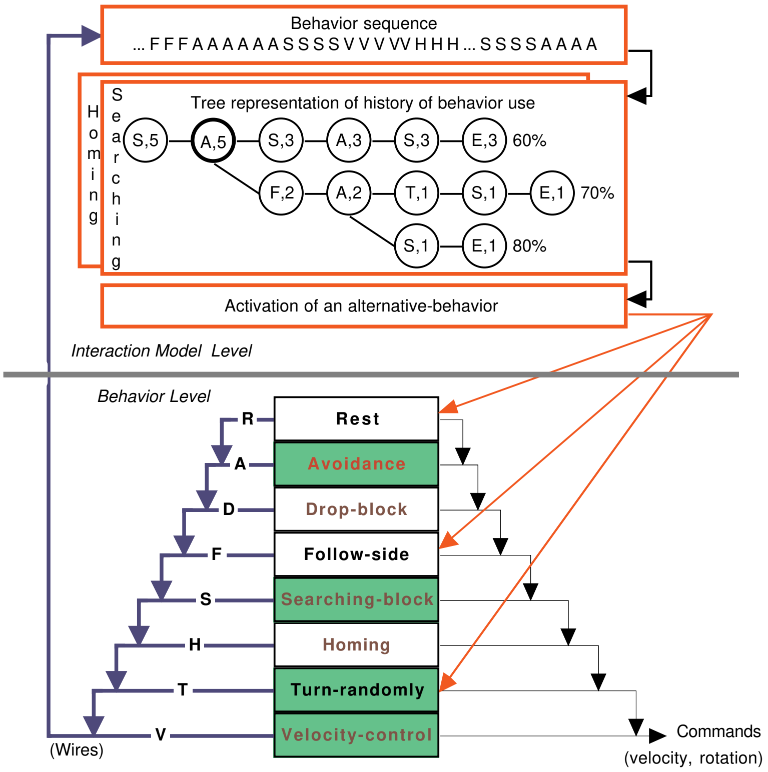

Fig. 2. Organization of the Behavior Level and the Interaction Model Level. The colored background represent an example of behaviors activated for the Searching Task with *Turn-randomly* as a chosen alternative-behavior. For clarity, sensory inputs to behaviors are not shown. Also note that behavior classification given above is independent of the subsuming organization used in our experiments.

described next.

3.1 Tree representation mechanism

The algorithm uses a tree structure to store the history of behavior use. The upper part of Figure 2 shows a typical tree in which nodes store the type of behaviors used for controlling the robot while accomplishing its task, and n, the number of times a transition between the node itself and its successor has been made (as observed from behavior use). Because there can be more than one task-behavior associated with each task, the node type for task-behaviors corresponds to the name of the predominant behavior for that task: a type S node represents the use of Searching Task behaviors, and a type H node is created when Homing Task behaviors are exploited consecutively. For synchronization reasons, the root node of a tree is always a task-behavior node, whatever is observed at the start of the task. For the maintenance-behavior and the alternative-behaviors, the letter associated with the name of the behavior is used as the node type. For example, according to the behavior sequence shown in Figure 2, the last path created in the tree was the one with the sequence *S-A-F-A-S*. The Homing task tree was then selected, since the *Homing* behavior (*H*) was used. After a while, the Searching task was initiated again and the observed behavior sequence was *S-A*. The current node in the tree is then the *A* node depicted in bold.

Initially, the tree for a particular task is empty and is incrementally constructed as the robot goes about its task. Leaf nodes are characterized by the letter *E* (for end-node) and store the total *performance* of the particular tree path. Whenever a path is completely reused, i.e., when the same sequence of behaviors is executed, the average of the stored and current performances for the last 10 trials (this number was determined empirically from various experiments) is computed and stored in *E*. The history of behavior use and the resulting performance values provide the ability to learn to trigger alternative-behaviors according to the dynamics observed by the robot in the environment.

3.2 Metric for self-evaluation of performance

Depending on the domains and tasks, a variety of factors can be used to evaluate performance. Different metrics can be used with our learning algorithm. The one that we chose to study in our experiments is based on time of behavior use. Our objective was to see how far we could take the idea of self-evaluation and learning by observing behavior use, and see how it can be useful to a robot with limited capabilities operating in a dynamic environment. Optimality is then difficult to define and may also be nonstationary. A more detailed justification of this choice is given in Section 5. Our evaluation function, expressed by Relation 1 below, is characterized by the amount of time required to accomplish the task and the interference experienced during that process:

$$eval(t) = \max\left(0, \frac{(t_{tb} - t_{mb} + t_c)}{t_{tb}} - \max\left(0, \frac{(t - TT)}{TT}\right)\right) \qquad (1)$$

where tb refers to task-behaviors, mb to maintenance-behaviors, and t to the current time-step (corresponding to one cycle of evaluation of all currently active behaviors). TT represents the average total number of time-steps taken to accomplish the task over the last 10 trials, updated upon completion of a task. Again, the number 10 was determined empirically by analyzing typical results; it establishes a compromise between adaptability to changes in the environment and tolerance for poor performance resulting from necessary exploration of different options. TT captures changes in the dynamics of the environment that affect the time to complete the task. t_c is a correction factor that increases when $eval(t)$ decreases below 0, to boost $eval(t)$ as soon as the robot uses its task-behaviors. The first term of Relation 1 reflects the amount of control the robot has over the task compared to the effect of other events that maintenance-behaviors must handle. The second term penalizes the robot's performance if it takes more time to accomplish the task than it has in the past. Also note that, as shown in Figure 2, the sequence of behavior use situates the robot in regard to past experiences, and $eval(t)$ takes into consideration the amount of time a single behavior was used. In this way, temporal separation of nodes is considered in the representation and has an effect on what is decided for behavior activation and what is memorized in the end node of a trial.

3.3 Activation criteria for behavior selection

Using the tree and the evaluation function, our algorithm has two basic options: not to make any changes in its active behavior set (this option is referred as the *Observe* option), or to activate an alternative-behavior. The sequence of nodes in a tree path characterizes the interactions experienced by the robot in the world, and situates the robot relative to those interactions. Different selection criteria can be used by comparing the performance at the current position in the tree (i.e, $eval(t)$) with the expected performance. The expected performance of a given node is the sum of the stored end-node performances in its sub-paths, multiplied by the frequency of use of the sub-paths relative to the current position in the tree. Formally, the expected performance for sub-path i starting from a given position in the tree is given by Relation 2:

$$E_i(eval) = \sum_j eval(T)_j \cdot \frac{n_j}{n_i} \qquad (2)$$

where j refers to the end-nodes accessible from sub-path i, $eval(T)_j$ is the performance memorized in these end-nodes, and n is the number of times the node j of sub-path i was visited. This gives the expected performance for each sub-path in a tree, i.e., for each possible option at the current position in the tree. This formula can also be used to determine the expected performance at this position, $E(eval)$.

A simple example can be given by using the bold node in the tree of Figure 2 as the current node. The expected performance of the sub-path starting with S node is $60 \cdot 3/3 = 60\%$, while the expected performance of the F sub-path is $(70 \cdot 1/2) + (80 \cdot 1/2) = 75\%$. The total expected performance $E(eval)$ for this node is $(60 \cdot 3/5) + (75 \cdot 2/5) = 66\%$. Note that longer paths do not necessarily mean lower performance; the evaluation function is not influenced by the number of different nodes in the path but by the time spent in behavior use.

Using the expected performances at the current position in the tree and those at each of its sub-paths, the choice made by the algorithm is based on the following criteria:

- If ($eval(t) \geq E(eval)$), then take the option with the maximum expected performance $E_i(eval)$. This criterion is referred as the TM criterion, for *Try Maximum*.
- If ($eval(t) < E(eval)$), then two possible situations exist:
  · If there are untried options at this junction, then select an untried option. This selection criterion is referred to as UO, for *Untried Option*.
  · If there are no untried options at this junction, choose the one with overall best performance (i.e., over all the trials up to that point). The overall best option is obtained from a *Global Options Table* (GOT), an array of past performances indexed by the name of the option. The GOT is updated at the end of each trial by averaging the resulting performance with the ones stored, for all the options used for that trial.

The evaluation of the behavior activation criteria takes place only at a maintenance-behavior node (i.e., node created by the observation of a maintenance-behavior) because that is when interference occurs in the accomplishment of the robot's tasks, resulting in decreased $eval(t)$. When one of these nodes is reached, the selection can be directly set by UO or GOT if $eval(t)$ is already lower than $E(eval)$, or by TM if $eval(t)$ is greater or equal to $E(eval)$. If this condition changes at the given position in the tree, UO or GOT can override the selection made using TM. Note also that when a new path is constructed, its expected performance is 0 and the default option is to behave without using any alternative-behaviors. This allows the robot to explore what can be accomplished by using task- and maintenance-behaviors only.

3.4 Path deletion mechanism

Since this algorithm is used in noisy and nonstationary conditions, deleting paths is necessary to keep the Interaction Model that the tree represents up to date. More specifically, because the expected performance is a weighted average based on the frequency of occurrences of the sub-paths, the most tried option at a given maintenance-node will have a greater influence on $E(eval)$. This decreases the variation of $E(eval)$ as an option is tried frequently, making it harder to modify it and to change strategies if the option becomes less useful. Forgetting paths may lead to exploring some of the choices again and discovering changes in the world. Node deletion also serves to regulate memory use. To enable the robot to respond to recent changes in the environment, we opted for deleting the oldest path in the tree when more than 25 different paths were stored. Other criteria can be easily substituted for different applications.

4 Experiments

Although the goal of the proposed algorithm is to learn in dynamic environments, it would have been inappropriate to start directly with a multi-robot domain experiment where nonstationary conditions make the analysis and the validation of the properties of the algorithm extremely difficult. To gradually address this problem, we first used static environment conditions with increasing complexity to verify that the algorithm could learn in stationary conditions. Next, experiments with multiple robots were performed. The same behavior repertoire was used in both sets of experiments, without optimization or retraining of the vision system parameters for color detection. The objective was to create situations where the designer could not know or adjust the behaviors according to *a priori* knowledge of the environment. Learning occurred over several trials of searching for the block (Searching Task) and bringing it home (Homing Task). Also note that during the trials, the learning robot did not always start its task in the same position: the searching task started 10 seconds after having dropped the block at home, and the direction taken from home depended on the proximity of the robot to the wall and the position of the dropped block relative to home; the homing task started right after a block was found, in the same direction the robot faced when it got the block. All the computations were done on board, with a limited amount of memory (the entire robot program used 66 Kbytes of memory, with 27 Kbytes used by the learning algorithm, leaving 43 Kbytes for storing the trees; a tree of 200 nodes takes approximately 7 Kbytes).

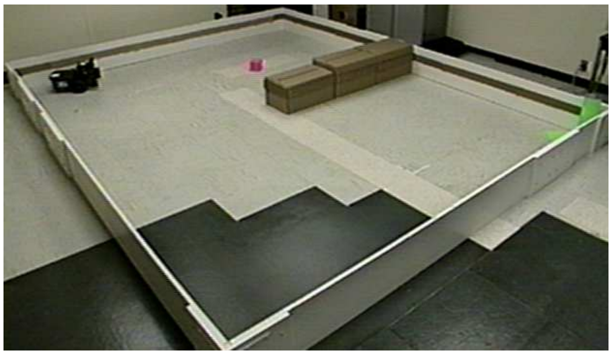

Fig. 3. Environment configuration with the pen divided into two regions.

Even in simpler static conditions, unexpected behavior selection was learned by the robot. For example, in the configuration with a single block placed at the center of the pen, following a wall to a corner was a good strategy to easily find the block. But the robot also learned to use *Rest* in front of a wall in order to rotate more than when the *Avoidance* behavior was exploited right away, helping it see the block directly from that location. Without the learning algorithm, the average total number of time-steps to accomplish the Searching Task was 492 cycles, compared to 317 cycles when the learning algorithm was used. We were surprised to find that *Turn-randomly* was not frequently useful: at first, we thought that this behavior would help the robot move toward the center of the pen to see the block. Instead, because this behavior made the robot rotate by a random angle at a random interval, it was not very reliable in helping the robot see the block. In fact, it frequently made the robot loop around the block. It is important to note that all of these discoveries were made by the robot from its own point of view, instead of by the designer.

We also performed experiments by starting with one block placed in the center (for the first 60 trials), and then dividing the pen into two regions (shown in Figure 3). Figure 4 shows the plot of $E(eval)$ at the first junction of the tree and the number of nodes in the tree for the learning trials of one such experiment. In the first 60 trials, the robot learned to use the *Follow-side* behavior to follow the wall to a corner, and then go to the block. Then, a two-phase change occurred: the strategy changed to using the *Observe* option for a while, and then *Rest* from trial 87 on. This can be seen from the changes in $E(eval)$. In fact, *Observe* and *Rest* were the second and third best options at the first junction of tree before the change of environment configuration occurred. The change of strategies and the stabilization starting around the 87th trial is also corroborated in Figure 4, demonstrating that the algorithm was able to quickly adapt to changes in the environment. In effect, the approach shows transfer of what was learned in the first configuration to what was useful in the new, second configuration: an initial strategy can thus serve as a bias toward the elaboration of a specialized way of behaving, but does not prevent it from changing dynamically. Also, at the end of these experiments, the average total number of time-steps to accomplish the Searching task is similar to the results obtained when the robot learns only in the pen divided in two regions configuration. This shows that the algorithm is able to improve rapidly and adapt the behavior selection strategy according to the situations experienced in the environment. Overall, the algorithm was able to converge to stable policies that revealed to be optimal in regard to the robot's perspective (which includes limitations in its perceptual field and its actions). A more detailed description of the results obtained in static conditions is presented in [21].

In the multi-robot experiments, the block to be brought home was put in the center of the pen, as shown in Figure 5. Two or three R1 robots were added to the scenario, and generated interference in the environment. The R1s were programmed to move about, avoid all obstacles and stay out of forbidden regions, using their contact and infra-red sensors, mounted in the front and below the grippers. These forbidden regions were the areas of the pen that were marked with dark tiles or tape, i.e., the area surrounding the block, the home region, and one dark-tiled corner of the pen. The R1 controller was simple: when detecting an obstacle or a forbidden area, the robots moved backward and turned away from the obstacle, then moved forward again; otherwise,

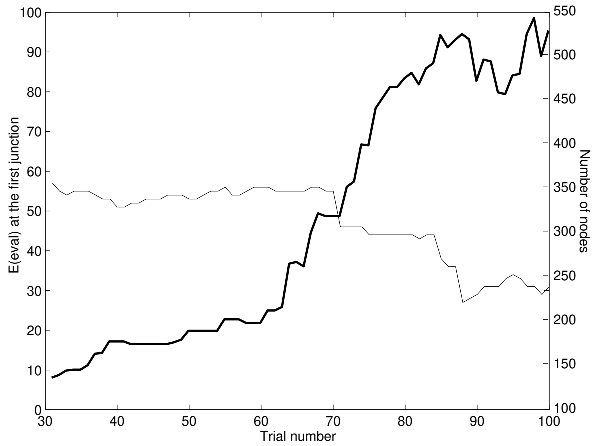

Fig. 4. Graph of $E(eval)$ and of the number of nodes (in bold) for the Searching Task tree of an experiment starting with one block in the center, and then changing by dividing the pen into two regions.

it moved forward freely. These constraints on the R1s' motion resulted in a great deal of dynamic interference. Additionally, although dark areas were inaccessible to the R1s, sensor noise enabled them to occasionally enter those areas nonetheless.

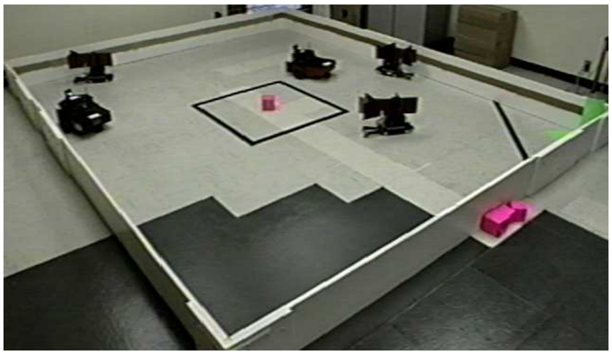

Fig. 5. Environment configuration for the multi-robot experiments, with two Pioneers and three R1s. The R1s are programmed to avoid obstacles and areas marked with dark tiles or tape.

Using this configuration, fifteen experiments of 60 to 80 trials were run, lasting between 1.5 and 2 hours and using two or three R1s. In some of these experiments, two learning Pioneer robots were used to compare their learned strategies. However, when placed together, additional sources of noise affected the behaviors of the robots: 1) when facing each other at particular angles, crosstalk between the sonars sometimes occurred; 2) one Pioneer robot searching for a block could perceive and pursue the block pushed by the other; 3) the blue Pioneer robot sometimes confused the red Pioneer robot with a pink block. During these experiments, the nonstationary conditions arose from the presence and movement of the R1s, and varied greatly during and between experiments. In some cases, one or two R1s stayed close to the home region, while in others all moved around the region near the block. As the learning is primarily dependent on the sequence of past experiences, which changes in each experiment run, optimality cannot be determined from the task alone (unless a very large amount of data is available for statistical analysis). To counter this, we used the run-time data and studied the experiment video tapes in order to analyze what, why, and how strategies were learned, if regularities were found from the interactions dynamics, and how the algorithm adapted to changes detected from these interactions.

For conciseness, we describe only the results obtained for the Searching Task because they best illustrate the properties of the learning algorithm. In this task, the learned strategies largely involved the use of the *Observe*, *Rest* and *Turn-randomly* options. The sequences of options learned were: *Observe – Rest*; *Observe – Rest – Observe*; *Rest*; *Rest – Turn-randomly*; *Rest – Observe*; and *Turn-randomly.* The strategies changed over time according to past experience: when performance for one strategy decreased, another was used based on what was learned previously. One or two stable strategies "won out" in most experiments.

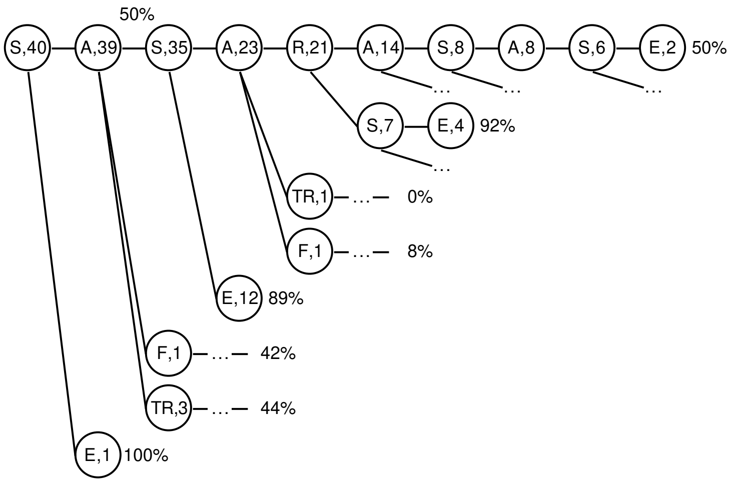

Fig. 6. Representation of a Searching Task tree for a multi-robot experiment.

Figure 6 represents part of a tree learned by the robot after 68 trials. The complete tree is made of approximately 200 nodes. The *Observe* option at node $(A, 39)$ (meaning that a maintenance node was observed in 39 trials at this position in the tree) has $E(eval)_i$ equal to 50% and is the one with maximum $E(eval)_i$ at this position in the tree. Note that this path may or may not be followed depending on $eval(t)$ and the use of the selected option. However, this path was taken 35 times over 39 trials.

Figure 7 shows the trace of the choices made for the first 16 trials and the last 10 trials of a test. Data associated with a choice are delimited by parentheses and consist of the trial number, the name of the option chosen, the name of the selection criterion used, $eval(t)$ and $E(eval)$ at the position of the tree where

```
1:(O TM 0 0)
2:(S TM 98 17)
3:(F UO 0 40)
4:(R UO 0 60) (O TM 98 0)
5:(TR UO 0 58) (O TM 96 0)
6:(F TM 75 57)
7:(F GOT 0 62) (S TM 96 91) (F UO 89 91)
8:(F TM 90 63)
9:(F TM 75 65)
10:(F TM 95 66) (TR UO 70 76) (O TM 42 0)
11:(F GOT 0 61) (F UO 87 91)
12:(F TM 96 62) (O TM 64 0)
13:(F TM 96 57) (R UO 45 64) (O TM 43 0)
14:(F TM 92 53) (F UO 45 82) (O TM 43 0)
15:(S GOT 0 50) (S TM 98 17)
16:(TR TM 96 49)
...
58:(S GOT 0 49) (R TM 92 37) (F TM 96 32) (O TM 53 0)
59:(TR TM 92 49)
60:(S GOT 0 50)
61:(TR TM 94 51) (S TM 77 55) (F UO 54 55)
62:(S TM 98 52) (S GOT 51 52) (R TM 58 37) (S TM 57 29) (R UO 26 32) (O TM 8 0)
63:(S TM 97 50) (R TM 89 35) (F TM 88 27) (O TM 13 0) (O TM 13 0)
64:(S TM 87 50) (64 S GOT 37 50)
65:(S GOT 0 50) (R TM 88 34) (E TM 96 27) (O TM 92 0) (O TM 0 0)
66:(S TM 96 50) (S GOT 48 50) (R TM 56 32)
67:(S GOT 0 51) (R TM 99 34) (S TM 97 25) (S TM 83 27) (R TM 57 31) (S TM 65 35)
   (F UO 34 35) (O TM 3 0)
68:(S TM 98 51) (R TM 75 33) (TR TM 74 23) (E TM 66 3) (R UO 2 3) (O TM 0 0)
```

Fig. 7. Trace of choices made for the Searching Task tree in a multi-robot experiment.

the choice was made. In the initial trials, *Follow-side* was the first favored option. However, this strategy did not remain stable for long; from trials 28 to the end of the experiment, the *Observe – Rest* strategy was preferred (e.g., in trials 58, 62, 63, 65 to 68). Between these trials, different options were explored. These choices were made using the TM and GOT criteria. Consecutive choices with the same $E(eval)$ values indicate that they were made from the same position in the tree. Notice that for trials 62, 64 and 66, a choice using the TM criterion is overridden by one made using the UO or the GOT criterion because $eval(t)$ fell below $E(eval)$. Shorter paths are a sign that the robot was able to find the block without experiencing much interference from the R1s. Otherwise, longer sequences of choices were generated because longer paths were reused. From Figure 6 at node $(R, 21)$, it is also possible to see that the algorithm allowed the robot to discriminate two situations using the *Observe – Rest* strategy: using *Rest* in front of a static obstacle and using *Rest* in front of a moving obstacle. This distinction is made according to the behavior used after resting, without using any special processing to distinguish a moving from a static obstacle, and lead to distinct strategies for behavior selection. Figure 8 shows that the number of nodes for this tree stopped increasing and started to decrease around the 50th trial, demonstrating reuse of stored paths in the tree even in nonstationary conditions.

Finally, when two Pioneer robots were put in the pen at the same time, they specialized differently in their ways of accomplishing the tasks, based on their

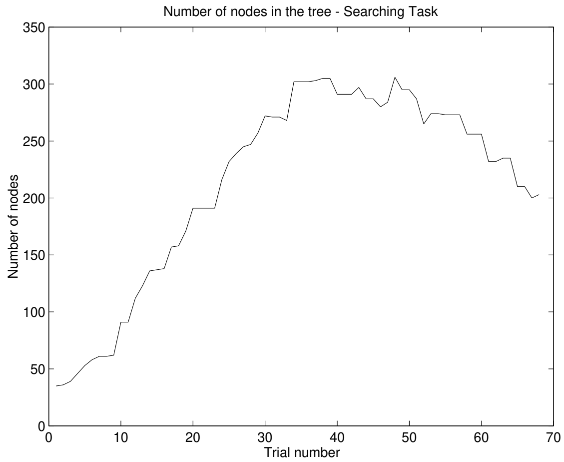

Fig. 8. Graph of the number of nodes in a Searching Task tree in a multi-robot experiment.

specific experiences. For example, in one test the strategies learned by one Pioneer for the Searching Task were *Observe – Rest – Observe* and *Rest – Turn-randomly*, while the other learned to mostly use *Turn-randomly*. Their trees were also quite different: the first one had 468 nodes and the other 256 nodes with frequently reused paths. In this case, what actually happened is that the first Pioneer got to interact greatly with the R1s, while the second one had easier opportunities to find the block. Having two Pioneers or more R1s in the pen did not influence the Pioneers' ability to find stable strategies; the dynamics of interaction with the environment had a more significant effect. Thus, dynamic specialization can occur without requiring complex communication and arbitration schemes for managing the group.

5 Discussion

The experiments done to validate the properties of our algorithm involved four important properties of the environment: dynamic, inaccessible, non-deterministic [25] and nonstationary. An environment is dynamic if it can change while the agent is deliberating, and so our multi-robot experiments are clearly done in dynamic conditions. Determinism is present when the next state of the environment is completely determined by the current state and the actions selected by the agent (according to what is accessible by perception). Perception of our learning robot was limited to sonars and color detection, and it did not have any information on its position in the pen, the position of other robots or on the configuration of the environment. By also considering the effects of noise and imprecision of the actuators, taking the same action in the same state in two different occasions could result in different next states [10] in our experiments, making them non-deterministic. Stationarity is more difficult to explain since definitions are usually given in relation to the learning algorithm. Kaelbing et al. [10] indicate that an environment is stationary when the probabilities of making state transitions or receiving specific reinforcement signals do not change over time. Sutton and Barto [27] explain that nonstationary means that the true values of actions (i.e., the actual Q-values/the expected or mean reward given that that action is selected) change over time. Changing the reinforcement signals during an experiment is one way of creating nonstationary conditions. Such changes can also occur when the environment is modified so much that the 'optimal' policy also has to change. This is what we wanted to validate by first doing an experiment in static conditions starting with one block in the center and continued with the pen divided in two regions: the policy learned in the first configuration was not a valid one for the second. Even more changes occurred during the multi-robot experiments because of the distribution of the robots in the pen. We would have had stationary conditions with moving objects in the pen if they could have remained uniformly distributed in the environment and had the same effect on the learning robot. However, that was not the case here: R1s started by being evenly distributed in the pen, and then changed their position to move around the center of the pen (where the block was), to move near the home region (causing a lot of difficulties for the learning robot to deposit the block), or to move in a remote area where the learning robot did not interfered with them. The behavior selection policy had to be continuously adapted based on these changes.

Our approach shares similarities with but also differs from the reinforcement learning paradigm. Learning is done in a manner similar to Monte Carlo learning algorithms [27], i.e., by averaging the observed returns to eventually converge to the expected value. In our learning algorithm, the decision is derived from stored past experiences and the current situation in the tree; an action takes the form of a behavior activation, and the reward is based on the internal measure of performance, which is also deduced from behavior use over time. This is different from reinforcement learning methods that use immediate sensations for states and commands for actions. Sutton and Barto [27] indicate that "there is no reason to restrict the state representation to immediate sensations; in typical applications we should expect the state representation to be able to inform the agent of more than that". Our approach is also very different from dynamic programming approaches which aim at finding a probability-based state/action table for maximizing utility values. In our case, behaviors serve as an abstraction for both state (i.e., behavior use) and action (i.e., behavior selection), and the decision process does not optimize a fixed utility function. As a result, having a state representation based on behavior use combines both a representation of the environment (i.e., the sensory information used by the behavior) and the control policy (since the use of a behavior depends on its activation and on behavior organization). "Behavior activation actions" made by the Interaction Model have a similar effect: they combine control signals (issued by a newly activated behavior) with sensory inputs (which influence control signals issued by the newly activated behavior and possibly by another behavior that subsume it). This way, we follow the principles of emergence and situatedness of the behavior-based approach (especially because an activated behavior may not get to control the robot right away in accordance to what is experienced in the environment), and we construct a state-action representation that is derived from behaviors.

Another key difference from the commonly used reinforcement learning algorithms is that we did not use a discrete reinforcement signal, or one completely determined by perceived conditions. The evaluation function used in our experiments differs from approaches attempting to optimize external criteria such as time required to accomplish the task, total distance traveled, number of collisions, etc. Mahadevan and Kaelbling [15] indicate that it is crucial to approximate the value function during reinforcement learning for continuous or large finite state spaces problems. Using the non-monotonic decreasing performance function of Relation 1 enabled 'self-evaluation' by the learning algorithm. This makes the learning independent of the constraints of the environment, the tasks, the limitations of the robot and the specifics of the designed behaviors [21]. There may exist many efficient strategies for the given task; our approach aims for reasonably efficient behavior in accordance with the robot's abilities to interact with the environment. In those conditions, instead of using fixed and external criteria, it may be more appropriate to evaluate learning performance by the ability of the approach to make consistent decisions based on past experiences, with a good balance between exploration and exploitation of the options. Our results demonstrate that the algorithm has this ability. Our approach also helped generate unanticipated behavior selection strategies, decreasing the burden of behavior design by not having to fine-tune the behaviors or acquire knowledge about the specific and complex dynamics of the system. This allows the robot to learn strategies that are difficult and unlikely for the designer to anticipate.

Overall, our algorithm attempts to capture and characterize the dynamics experienced based on the tree representation and the evaluation function. In fact, the different components of the learning algorithm try to establish a compromise between *exploration* (learning to adapt to noise and changes in the environment) and *exploitation* [10] of a stable behavior selection strategy. The evaluation function characterizes the current situation and past experiences. The tree representation captures sequences of behavior use in a compact fashion to make a decision based on past experiences. The behavior activation criteria exploit both of these components to determine when and what choice to make: the TM criterion favors exploitation of the options, while the UO criteria favor exploration for the purpose of finding a different best option. The GOT plays an intermediate role between exploration of options at a given node (when the algorithm does not know which choice to make), and exploitation of the option with the best performance. It serves as a *medium-term memory* of the performances of the options. At the same time, the GOT makes the learning algorithm adapt more easily to changes in the environment that can be reflected in a more constant fashion in the performance of the trials. For a frequently used strategy, $E(eval)$ is more difficult to change because its evaluation is made according to a larger number of trials. In contrast, the GOT is not subject to this and allows for faster exploration by exploiting outcomes of recent trials. The GOT criterion is preferred to random behavior selection because it favors using knowledge derived from past experiences instead of blindly exploring all of the options. Finally, path deletion influences the choices made by affecting $E(eval)$ in the tree. The deletion factor defines the memorizing ability of the algorithm, which can also be viewed as a principled bias set up by the designer. It does not affect the type of strategy learned, but rather how it is learned. This results in continuous adaptation of what is being "optimized" in the behavior of the robot, without requiring extensive training periods to model changes in nonstationary conditions.

In effect, these influences allow the algorithm to establish *critical points* for behavior selection, i.e., decision factors for activating behaviors. These critical points are self-determined and modified in response to continuing changes in the dynamics, which could be caused by alterations of the environment, the behavior of other robots, or the variability of the robot's abilities (e.g., due to noise, battery power, etc.). Even in a nonstationary environment, these critical points provide the robot with some ability to make expectations based on previous experiences, and thus to evaluate its current performance.

6 Related work

Our approach fits in the category of hierarchical methods for dealing with large state spaces [10]. Hierachical learning usually involves a structure of gated behaviors, i.e., a collection of behavior-producing modules, and a gating function that decides, based on the state of the environment, which behavior's actions should be switched through and actually executed [10]. In some works [8,14], the gating function is fixed and learning occurs in behaviors by supplying them with individual reinforcement functions. Other works [13,17] like ours involve learning the gating function, or both [2,7]. The main difference from the above approaches is that ours uses history of behavior use instead of percepts as states for selecting behaviors.

Compared to other learning approaches like POMDPs, our approach also tackles the problem of hidden state, particularly important for mobile robots equipped with limited and noisy sensors. POMDPs try to compensate for the limitation of the Markov property [27] of the states by maintaining belief states in the form of a probability distribution. However, there is no theory of reinforcement learning for a non-Markovian model [26] and it seems that exact solutions of even medium-sized POMDPs are generally considered infeasible [24] and present a very difficult optimization task [15]. They require exponential amount of space and time to compute a policy, even if the policy itself does not require an exponential size representation (exact solutions to POMDPs with tens of states can take anywhere from minutes to days, if convergence is obtained at all [22], and it has been shown that the exact solution of these problems is intractable [15].). Cassandra et al. [6] use a POMDP model to make a mobile robot navigate in a real office environment. No learning is involved, and the POMDP model is valid only in static environment conditions. Learning POMDP models is done by Koenig and Simmons [11], and to make the approach feasible, subsets of probabilities are kept constant or made identical (from initial knowledge), and the backward probability variables are approximated using a sliding time window. Our approach is different from POMDPs: it does not consider the probability of observations inside a node, it is not based on a state-action representation, and it does not use the Baum-Welch algorithm for learning.

Taking into consideration the history of past experiences was inspired by McCallum [18,19], whose approach partitions the state space from raw sensory experiences, and uses memory of features encoded in a U-tree to augment the agent's perceptual inputs. These features that label the branches of the U-tree have a perceptual dimension (sensory data or actions) and a history index (time step backward). The leaves of the tree store utility estimates (Q-values) that are used to choose the next action. This mechanism is successfully applied to learning to drive in a virtual highway environment. Globally, the U-Tree approach learns a variable-width window that serves simultaneously as a model of the environment and a finite-memory policy [10]. In our case, the algorithm learns a finite-memory strategy of behavior use and selection, since it uses behaviors as an abstraction. While McCallum's approach is able to learn a complete policy (based on a discrete state space that gives a good representation of the environment), ours is more concerned with how to arrange behaviors to control the robot. This allows us to be independent of having discrete or continuous sensory state space, to work with systems that have less information to discriminate about the state of the world, and to focus learning on things not known *a priori* and improve learning time. In addition to applying a different performance measure that is independent of the task and more concerned with how it is accomplished, our approach differs from McCallum's in the selection criteria. While the U-Tree algorithm follows the path with maximum expected reward and random exploration, our algorithm compares the current with the expected performance from the tree and the GOT. Additionally, exploration is also affected by the path deletion mechanism, while the U-Tree uses statistic evaluation at the leaf nodes.

Learning from observed behavior activation was also explored within the case-based reasoning framework, in work by Ram and Santamaria [23]. However, the computational complexity of their approach is much greater than ours. Finally, the development of our algorithm was done in the context of a general control architecture for an intelligent agent [20] based on dynamic selection of behaviors.

7 Conclusion

Learning in dynamic, inaccessible, and nonstationary environments, as in our multi-robot domain, is a very challenging problem. Since interaction with the environment is affected by the dynamics and the robot's own limitations in perception, action, and processing, our approach advocates internal evaluation of the interaction dynamics, and taking advantage of them while they are still valid. The approach characterizes these dynamics by observing the use of behaviors and by using a representation and a selection mechanism that are consistent with the behavior-based framework. Everything is done relative to the exploitation of the controlling resources (i.e., the behaviors) of the robot, and thus defined according to the robot's own perspective. Results show the effectiveness of the approach in taking advantage of any regularities experienced in the world. This allows fast and adaptable agent specialization, and generates temporarily-stationary conditions in a multi-robot environment, when possible. Results also show that the approach leads to learning unanticipated behavior strategies from the perspective of the robot instead of the designer, thus increasing robot autonomy. The approach provides a predictive capability that is typically lacking in behavior-based systems, thus making a move toward a more abstract representation level for generating complex behavior.

Acknowledgments

Financial support to F. Michaud is provided by NSERC (Natural Sciences and Engineering Research Council of Canada). This research is also funded by the Office of Naval Research Grant N00014-95-1-0759, the National Science Foundation Infrastructure Grant CDA-9512448 and the Career Grant IRI-9624237 to M. Matarić. The authors also want to thank P. Melville, B. Werger, D. Goldberg and the rest of the Interaction Lab for their insightful comments.

References

[1] P. E. Agre, The Dynamic Structure of Everyday Life, PhD Thesis, Massachusetts Institute of Technology, 1988.

[2] M. Asada, E. Uchibe, S. Noda, S. Tawaratsumida, and K. Hosoda, Coordination of multiple behaviors acquired by a vision-based reinforcement learning, in: Proc. IEEE/RSJ/GI Int'l Conf. on Intelligent Robots and Systems (Munich, Germany, 1994).

[3] R. A. Brooks, A robust layered control system for a mobile robot, IEEE Journal of Robotics and Automation, RA-2(1) (1986)14–23.

[4] R. A. Brooks, Intelligence without representation, Artificial Intelligence, 47 (1991) 139–159.

[5] R. A. Brooks, MARS: Multiple Agency Reactivity System, Technical report, IS Robotics (1996).

[6] A. Cassandra, L. Pack Kaelbling, and J. A. Kurien, Acting under uncertainty: discrete bayesian models for mobile-robot navigation, in: Proc. IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems (1996).

[7] M. Dorigo and M. Colombetti, Robot shaping: Developing autonomous agents through learning, Artificial Intelligence, 71(4) (1994) 321–370.

[8] D. Floreano and F. Mondada, Evolution of homing navigation in a real mobile robot, IEEE Trans. on Systems, Man, and Cybernetics—Part B: Cybernetics, 26(3) (1996) 396–407.

[9] D. Goldberg and M. J. Matarić, Interference as a tool for designing and evaluating multi-robot controllers, in: Proc. National Conf. on Artificial Intelligence (AAAI), (Providence, RI, 1997) 637–642.

[10] L. Pack Kaelbling, M. L. Littman, and A. W. Moore, Reinforcement learning: A survey, Journal of Artificial Intelligence Research, 4 (1996) 237–285.

[11] S. Koenig and R. G. Simmons, Unsupervised learning of probabilistic models for robot navigation, in: Proc. IEEE Int'l Conf. on Robotics and Automation (1996).

[12] P. Maes, The dynamics of action selection, in: Proc. Int'l Joint Conf. on Artificial Intelligence (IJCAI), (Detroit, MI, 1989) 991–997.

[13] P. Maes and R. A. Brooks, Learning to coordinate behaviors, in: Proc. National Conf. on Artificial Intelligence (AAAI), Vol. 2 (1990) 796–802.

[14] S. Mahadevan and J. Connell, Automatic programming of behavior-based robots using reinforcement learning, Artificial Intelligence, 55 (1992) 311–365.

[15] S. Mahadevan and L. Pack Kaelbling, The NSF workshop on reinforcement learning: Summary and observations, AI Magazine (1996) 89–97.

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

[17] M. J. Matarić, Reinforcement learning in the multi-robot domain, Autonomous Robots, 4(1) (1997).

[18] A. K. McCallum, Learning to use selective attention and short-term memory in sequential tasks, in: From Animals to Animats: Proc. 4th Int'l Conf. on Simulation of Adaptive Behavior, (Cape Cod, The MIT Press, 1996) 315–324.

[19] A. K. McCallum, *Reinforcement learning with selective perception and hidden state*, PhD thesis, Department of Computer Science, University of Rochester, 1996.

[20] F. Michaud, G. Lachiver, and C.T. Le Dinh, A new control architecture combining reactivity, deliberation and motivation for situated autonomous agent, in: From Animals to Animats: Proc. 4th Int'l Conf. on Simulation of Adaptive Behavior, (Cape Cod, The MIT Press, 1996) 245–254.

[21] F. Michaud and M. J. Matarić, Learning from history for behavior-based mobile robots in non-stationary conditions, Machine Learning & Autonomous Robots, Joint Special Issue on Learning in Autonomous Robots, 31 & 5(1/2/3 & 3/4) (1998) 141–167 & 335–354.

[22] R. Parr and S. Russell, Approximating optimal policies for partially observable stochastic domains, in: Proc. Int'l Joint Conf. on Artificial Intelligence (IJCAI), (Montréal, Québec, 1995) 1088–1094.

[23] A. Ram and J. C. Santamaria, Multistrategy learning in reactive control systems for autonomous robotic navigation, Informatica, 17(4) (1993) 347–369.

[24] S. Russell, Machine learning, in: M. A. Boden, editor, Handbook of Perception and Cognition, Vol. 14, Chap. 4. (Academic Press, 1996).

[25] S. Russell and P. Norvig, Artificial Intelligence A Modern Approach, (Prentice-Hall, 1995).

[26] S. P. Singh, T. Jaakkola, and M. I. Jordan, Learning without state-estimation in partially observable markovian decision process, in: Proc. 13th Int'l Conf. on Machine Learning (1994).

[27] R. S. Sutton and A. G. Barto, Reinforcement Learning: An Introduction. (MIT Press/Bradford Books, Cambridge, MA, 1998).