

Genetic Programming to Learn an Agent's Monitoring Strategy*

Marc Atkin and Paul R. Cohen

CMPSCI Technical Report 93-26

Experimental Knowledge Systems Laboratory
Department of Computer Science
University of Massachusetts
Amherst, MA 01003

Abstract

Many tasks require an agent to monitor its environment, but little is known about appropriate monitoring strategies to use in particular situations. Our approach is to learn good monitoring strategies with a genetic programming algorithm. To this end, we have developed a simple agent programming language in which we represent monitoring strategies as programs that control a simulated robot, and a simulator in which the programs can be evaluated. The effects of different environments and tasks is determined experimentally; changing features of the environment will change which strategies are learned. The correspondence can then be analyzed.

We present the simulated robot with the task of getting as close to an obstacle as possible without touching it. We know of two strategies that might be used to solve this problem: periodic monitoring, and proportional reduction, in which the robot moves a fixed fraction of the remaining distance to the obstacle between monitoring events. Our hypotheses concerned the environmental conditions in which each strategy might be learned. An experiment confirmed some of our hypotheses and left the status of others ambiguous. At the same time, the experiment revealed new monitoring strategies and set the stage for the further exploration of this field.

* This research was supported by DARPA/AFOSR contract F30602-91-C-0076.

forward, turn left and right, turn itself towards the goal, and use the sensor. These are the "sensor" that detects the presence of obstacles along the present heading. A robot can move determine the direction of the goal point, one that reports contact with an obstacle, and a while minimizing their energy expenditure. Our robots have three sensors, one to

Our simulated robots all perform a single task: to reach a goal point in the environment strategy can be evaluated.

advantageous. The environment also serves as a testbed in which an agent's monitoring create situations in which we expect different monitoring strategies will be most be dealt with effectively. By varying aspects of the environment and the task, we can dimensions and distance—it allows us to design tasks that require monitoring behavior to operate. Although the simulated world is simple—its only feature is an obstacle, plus

We have built a simulator in which populations of simulated robots individually

Testbed Description

that tells us when to use which strategy.

periodic monitoring. This work is a first step in setting up a general classification scheme We will identify a strategy, proportional reduction, that is in many cases superior to which different monitoring strategies are learned in environments with different features, testbed and the way the robot's behavior is represented, we will discuss an experiment in environment; so in most situations, it learns a monitoring strategy. After we describe the given noisy sensor data—usually can be accomplished more efficiently if it monitors its grid world. The agent's task—to get as close as possible to an obstacle without touching it,

Specifically, our algorithm generates programs that move a simulated robot around a

or impossible to derive analytically.

overview). We use genetic programming to discover monitoring strategies that are difficult to learn programs; in particular, programs that control a robot's behavior (see [5] for an

Recent work in automated programming has shown it is possible for a genetic algorithm

strategies appropriate to environments and tasks.

monitoring strategy, etc.). We therefore use a learning algorithm to find monitoring best time to monitor, for what, how should the current state of the environment affect my factors that influence a strategy (e.g., is it better to monitor or take an action, when is the analytically what strategy is optimal or even close to optimal, because of the multitude of the same time, while periodic monitoring is often not optimal, it is difficult to say monitored, such as approximate probability distributions in time or space for the events. At not exploit knowledge the agent might have about its environment and the events being and the acquired information is noisy, periodic monitoring is probably not optimal. It does dedicated processor) or periodic (see [1] for a survey). However, when monitoring has a cost The most common monitoring strategies are continuous (i.e., offoad monitoring to a environment, yet the literature on strategies for monitoring the environment is very sparse.

It is hard to imagine an agent that does not need to acquire information about its

Introduction

was sufficient to solve the problem. This restriction also forces the robot to come up with a length restriction. We therefore chose a program length (30 commands) which we knew while in some ways simpler and more readable, our programs do have the disadvantage of commands. This distinguishes them from the ГИЗ-style 2-expressions used by Kozs [3].

As it is currently implemented, programs are simply fixed-length linear strings of

in developing its programs:

the only architectural feature provided, so as to leave the robot as much freedom as possible

Apart from the choice of actuator and control commands, the interrupt structure was

operator- laden environment:

indeed found when we ran the algorithm on a scenario in which robots traveled an an obstacle is detected ahead—if surmises the move-to-goal behavior. Such a program was (EVALS) would be responsible for turning away from obstacles. It would be executed when and monitoring periodically for obstacles, whereas the goal-associated interrupt handler example, the main program could be responsible for always moving towards the goal point easily give some behaviors priority over others, in effect realizing a behavior hierarchy. For

By disabling and enabling interrupt handlers with explicit commands, a program can

the point of interruption:

program flow is interrupted, and the appropriate handler is executed before returning to an EVA value changes and the corresponding interrupt is enabled, then the normal designated to be interrupt handlers: one interrupt handler is associated with each sensor. If interrupted when a sensor value (an EVA value) changes. Parts of programs are events immediately. To this end, the normal execution of a robot's program can be

An agent operating in a real-time environment should be able to respond to external

explicit monitoring: EVA S—is there an obstacle ahead, and if so, how far away is it?

reached the goal point, and EVA I—have I hit an obstacle? Another EVA is required for sensors available to the robot. Two EVAs are updated automatically: EVA0—have I about the environment is held in environmental variables (EVAs). They correspond to the MONITOR and a few control structures (ГООБ, IE, etc.). The robot's current knowledge consists of actuator commands (MOVE, TURNLEFT, TURNRIGHT, SETCOURSDIR,

The simulated robot is controlled by a program constructed from a language that

Behavior Representation

same time limiting the monitoring cost.

therefore monitor repeatedly to keep the probability of hitting the obstacle low, while at the measured. A proportionally constant w expresses the accuracy of the sensor. A robot must a noise component, normally distributed with variance proportional to the distance and move as far as the sensor indicated. In our scenario, however, the sensor includes without touching it. If the sensor were 100% accurate, the robot would monitor only once facilitate the robot's task, that is, enable the robot to get as close to the obstacle as possible experiment described in this paper, we were interested in monitoring strategies that require energy, and an energy deduction is also charged for hitting the obstacle. In the basic actuator functions. Monitoring means using the sensor sensor. All these actions

exist only to make later insertions of new commands easier:

CELESTODIB points the robot towards the goal. The NOB commands do nothing: they terminate the program when the obstacle has reached a critical distance, the remaining before monitoring again. The command „IF (EVAL) <= \cdot THEN STOP„ proportionally constant is \cdot : the robot will move \cdot times its estimate of the distance value. Four MOVEs are executed in this loop, each moving \cdot distance units. Therefore, the takes the current distance to the goal (stored in EVAL after monitoring) and loops over its executed repeatedly. Nested within this loop is a second one, „LOOP (EVAL) times„, which

Within the program's outer loop, „LOOP times„, the „MONITOR: EVAL \cdot „ command is

monitoring event.

the proportional reduction strategy that this probability remains constant after each

The proportion determines the probability of overshooting the goal. It is a property of

to the goal and then move a fixed proportion of the estimated distance.

Proportional reduction: Repeatedly estimate the distance remaining

algorithm, demonstrates the proportional reduction monitoring strategy:

The following example program (Fig. 1) taken from an actual run of the learning

need not mean long programs.)

the explicit LOOP construct makes our representation concise: complex sequences of events compact solutions, which enhances their readability and, perhaps, generally. (In addition,

that reads the programs generated by our algorithm.

The data from the experiment described later were collected automatically by a program we can also write programs that read the monitoring strategies in the programs. In fact, we are able to read the monitoring strategies that our algorithm discovers. As an added bonus, this is one reason we have shied away from parameter-adjustment algorithms: we want to robot's actual behavior, but it is preferable to read the strategy from the program directly, describes its strategy. Strategies can sometimes be guessed by looking at a trace of the

The robot's program satisfies two goals at the same time: It controls the robot and statement in the main loop that gets executed independently of the distance monitored.

Note that this program is not pure proportional reduction, because there is one MOVE in the distance monitored would cause the interrupt handler to be called. This would mean that the distance is checked immediately after monitoring, since a change in the main loop. Occasionally, programs do actually use the EVAS handler in this way: the EVAS handler could be to test if the robot is close enough to the goal instead of doing so never called. Only the EVAS handler has any real use in this scenario. One possible use of interrupts are not explicitly enabled in this program, so the interrupt handlers are

Figure 1: An example program for proportional reduction.

```

MVAL
IOB
MOLE
DISABTE: EVA 1
EVA 3 interrupt handler:
IOB
IOB
UNKNIFEEL
IF (EVA 3) <= 4.0 THEN GLOB
EVA 1 interrupt handler:
EVA 0 interrupt handler:
MOLE
IOB
GOOB 4 time(a):
GOOB (EVA 3) timea:
IOB
IF (EVA 3) <= .2 THEN GLOB
MOLE
IOB
IOB
DISABTE: EVA 0
IOB
MONITOR: EVA 3
IOB
IOB
IOB
GETGOALDIR
GOOB 1 time(a):
GOOB 4 time(a):
IOB
UNKNIFEEL
IOB
IOB
MAIN program:

```

grows, as does the crossing over rate, during long periods in which the best individual does not program. The mutation rate is initially set to affect every 520th command on average, but if command parameters, or pointers to interrupt handlers in the selected individual, are affected by mutation or crossing over. Mutation operator randomly changes commands, raising or an individual within the population (see [2] for discussion). These 30% can be with the higher fitness value is chosen. This corresponds to a selection based on the linear with a tournament size of two: two individuals are selected with replacement and the one generation. The other 30% are selected based on fitness. We used tournament selection programs from the best 10% of the population are copied unchanged into the new

just before hitting it:

the robot with the highest total fitness is one that goes directly to the goal point, but stops obstacle (there are no other obstacles in the maze), and the penalty for hitting it is very high, minus the robot's energy expenditure on that trial. As the goal is always biased on an directionally as the distance between the robot's final position and the goal decreases, reproduction. Fitness consists of two terms, a closeness-to-goal reward, which grows learning process. The fitness of an individual determines its chances of survival and training pairs (start and goal points in the environment) that remain fixed during the

The individual is then assigned a fitness value based on its average performance on 15

by the execution of 1000 program steps:

simulated environment. Because programs need not terminate, the simulation is rounded each generation, every individual is evaluated by running the program it represents in the randomly initialized; that is, each individual is a random sequence of 30 commands. In

At the beginning of each run of the CA, a population of 800 individual programs is

Learning Procedure:

during a mutation operation, cannot affect syntactic correctness.

changing the start of an interrupt handler by changing the pointer, which can happen end when another one starts or when the end of the program is reached. Therefore, simply random pointers into the linear array of program commands; a handler is defined to length of 1000 bytes, which are taken modulo the program length. Interrupt handlers are modulo 3 with regard to any value). The same trick applies to other parameters, such as the number in particular commands has the legal range of 0 to 5, so taking the EVA number parameters, are taken modulo their maximum value, forcing them to be legal. (e.g. the EVA legal programs. Range restricted values, such as command numbers and some command program interpreter in such a way that all individuals that could possibly be generated are maintain the program's syntactic correctness. We solved this problem by designing the

The main difficulty is to ensure that the genetic operators, mutation and crossing over,

introduction to CAs).

(CAs) except that individuals in the population represent programs (see [4] for an excellent programming). This technique is completely analogous to conventional genetic algorithms

Programs are learned by a process of simulated evolution known as genetic

The Genetic Algorithm

noise should result in a more cautious program (with a low constant) that monitors after the proportional reduction constant should depend directly on the noise. A lot of influence monitoring strategy. The error in the "noise" data was of particular interest, as robots. We suspected that several features of the environment and the robot would

We wanted to see if and when proportional reduction would emerge in our simulated monitoring strategy [Δ].
 broblem show that proportional reduction is only a close approximation to the optimal now really the decrease is missed, but dynamic programming solutions to a more general that proportional reduction is optimal if the cost of missing a decrease does not depend on simply make not developed the proportional reduction strategy [θ]. Cohen demonstrated monitoring for the same task, because their internal clock is less accurate or because they rely on an inaccurate internal clock. Ten-year-olds, however, use periodic solving the "Субсакс Брoлeм", which requires them to take subtasks out of an oven before shown that 14-year-old children use a form of proportional reduction monitoring when for that matter, any other strategies besides periodic monitoring). Psychologyists make little is known about optimal monitoring strategies such as proportional reduction (or

Experiments Hypotheses

commands) was handled similarly.
 grew, the effect became weaker after about 200. The determination of code length (30 performance (in terms of fitness values achieved) generally increased as the population The population size of 800 was determined empirically: Although the algorithm's fitness so as to achieve a better approximation to the best strategy.
 dominated the rest of the population. We finally chose to run the genetic algorithm several appropriate to this problem, once an individual had evolved a suitable strategy, it often strategies make common modules, which is one of the reasons a genetic algorithm should be variation of up to 12% in terms of fitness improvement. Although many monitoring passage fitness for a program doing nothing was around 1000 for this case, this means a standard deviation 3σ), corresponding to a set of different monitoring strategies. As the world produce individuals whose fitness values ranged from 1205 to 1288 (mean 1248, fitness when the genetic algorithm was run repeatedly. For example, one typical test case algorithm's learning rate. Still, there was a considerable variance on the best individual's tournament selection and dynamically adjusting mutation and crossing over rate to the premature convergence was often a problem. We hoped to counteract it by using fitness after 120 generations, the algorithm was stopped.

usually converged within 300 to 800 generations. When there was no improvement in Although the maximum number of generations was set to 1000, the learning algorithm Performance:

initially set to affect every 10th individual on average.
 edual length are exchanged between two selected individuals. The crossing over rate is not improve. Crossing over simulates sexual reproduction: two random code segments of

times the agent monitored, as monitoring has a fixed cost, but it should not be hard to automate also. Note that measure 3 basically tells us how many programs execution trace. The fourth was determined manually by looking at the program. The first three measures were calculated by an evaluation function directly from the

- 4: proportional reduction constant
- 3: percentage of distance travelled while using proportional reduction
- 2: total cost of monitoring
- 1: fitness

strategy:

Four measures were averaged over the 30 trials to evaluate the program's performance and this program was then re-tested on 30 start-goal pairs in each specified length range.

of fitness) from these runs as our final result.

random number seeds for the genetic algorithm, taking the overall best program (in terms a good strategy in each of these cases, we repeated the process 10 times with different 30) and three ranges of pair length (1-4, 2-10, 15-18). To increase the chances of getting

We ran the genetic algorithm on five levels of λ , the monitoring error (0, 0.2, 1.0, 5.0,

Experiment Results

decreases.

- 3: Periodic monitoring becomes more likely as λ increases and pair length
- 2: The FB constant becomes smaller as λ increases.
- 1: If the monitoring error, λ is 0, a robot should monitor once only.

To summarize our hypotheses:

completely circumventing any need to monitor.)

otherwise the genetic algorithm will over adapt to move exactly the required distance, separated. (Note that this distance cannot be a specific length, but must be a range, proportional reduction should be more apparent when the start and goal states are more time, while the periodic strategy will monitor twice as much. This is why the advantage of and goal states is doubled, the proportional reduction strategy will monitor one additional travelling 10 units, then 8, then 4, then 2, then one. But if the distance between the start strategy with .2 as its constant. On average it will monitor once to start and then after distance between the start state and the goal is 35. Now imagine a proportional reduction strategy that monitors every five distance units; this strategy will monitor six times if the between the robot's starting position and the goal grows. Imagine a periodic monitoring

Theoretically, proportional reduction should show more of an advantage as the distance

event.)

periodic monitoring is optimal when no information is given about the occurrence of an data no longer contain any useful information. (It is easy to show mathematically that proportional reduction strategy would show no advantage over periodic monitoring, as the moving shorter distances. We hypothesized that when some data are extremely noisy, the

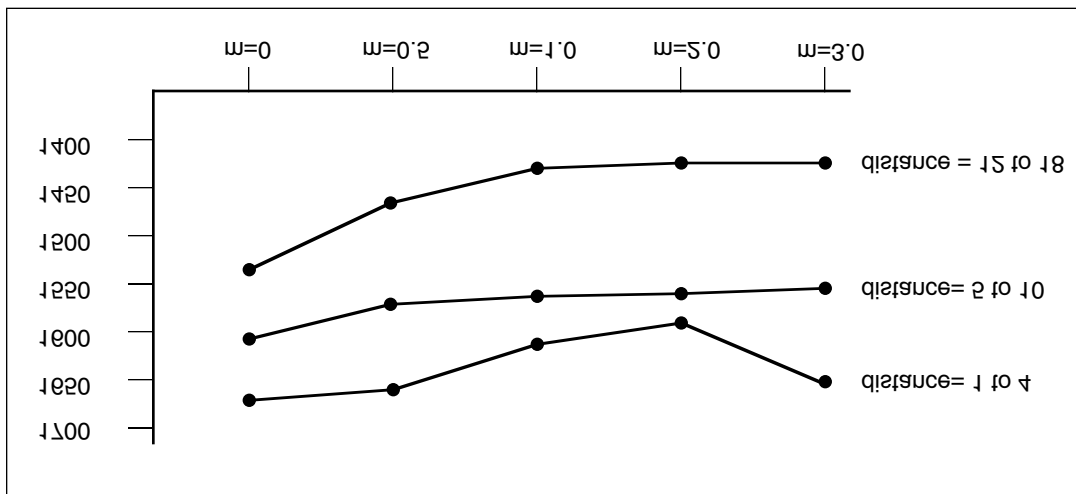
time of the algorithm.

to this problem is to train on more pairs, but it is a trade-off between better results and run time on the 30 test pairs, whereas it didn't on the 15 training pairs. Obviously the solution $w=3.0$ (range 1-4) cases being too low. In these cases, the robot hit the obstacle once or

The exception mentioned above is in fact actually due to the values for the $w=1.0$ and $w=3.0$ trials.

dependent measure is the fitness of the best individual to evolve in 15 trials, re-tested on 30

Figure 5: A plot of cell means for a two-way ANOVA, monitoring error (m) by distance. The



$p < .0001$), but no interaction effect.

shows a clear effect of path length and monitoring error on fitness (for both main effects genetic algorithm copes quite well with finding good monitoring strategies. An ANOVA however, that path length has a much greater effect than monitoring error. Apparently the values of w mean more elaborate and energy-consuming monitoring strategies. Note longer path lengths mean higher energy consumption, which lowers fitness, and higher the exception of the $w=3.0$, path length=1-4 situation. Both these effects were predicted: the most part, fitness values decrease monotonically as w and path length increase, with

Figure 5 shows the average fitness values for the final program on the 30 test pairs. For

Table 1: The best individual's monitoring strategy.

- (**) strategy is only capable of monitoring once
- (*) deliberately moves base goal to one side
- : no monitoring strategy
- check: stops if obstacle is visible after monitoring
hits a constant distance.
- DBB: proportional reduction: moves a proportion of distance remaining
- BB₂: moves a proportion of the squared distance remaining
- BB (c): proportional reduction strategy with proportional reduction constant

Key:


15-18 length	DBB	BB ('Δ)	(*) (**) DBB	DBB	DBB
2-10 length	BB ₂	DBB	BB ('∇)	-- (*)	-- (*)
1-4 length	BB ₂ (**)	BB ('Δ)	(*) (**) BB ₂	(**) BB ('∇)	(*) (**) check
	w = 0.0	w = 0.2	w = 1.0	w = 5.0	w = 3.0

ensuring that if proportional reduction fails, it will still not touch it.
15-18), it will combine the strategy of moving base the obstacle with proportional reduction,
monitoring strategy was found. Sometimes (case w=1.0, both length 1-4 and both length
test case. As you can see, in cases w=3.0, both length 2-10; and w=5.0, both length 2-10, no
training set. Table 1 gives an overview of the best individual's monitoring strategy for each
the left or right) for a distance that is approximately the average of the both lengths of its
instead tries to find a different solution such as deliberately moving base the goal point (to
too high, the genetic algorithm does not come up with a monitoring strategy at all, but
basically verified. The expected pattern is distorted because when the monitoring error gets
increase as both length and especially monitoring error increases. This prediction is only

Figure 3 shows the mean monitoring costs. We had expected to see monitoring cost

and four levels of monitoring error:


Figure 3: Average monitoring cost for 36 test cases for each of three levels of distance



respect to the distance moved vs proportional reduction. It is therefore very possible that implemented this behavior. Note that the fixed additional distance was usually small with constant throughout the strategy. The example program presented in a previous section the goal as if gets closer to the goal. In pure proportional reduction, this proportion remains reduction, as the robot is in effect moving a larger proportion of the distance remaining to remaining distance, the robot moves a fixed distance as well. We call this "disproportionate length. Sometimes after monitoring, in addition to moving a distance proportional to the robot was guaranteed a minimum distance from the goal by the lower bound on the path proportional reduction strategy. This is not surprising because in many of the trials, the frequently, the robot would move a short distance before beginning to monitor with the. The proportional reduction strategy did not usually manifest itself in its pure form.

robot uses the proportional reduction strategy.

Figure 4: Percentage of the distance between the starting state and the goal in which the



inaccurate: the sensor is used only to detect the presence or absence of the obstacle. This is notable because it is not a proportional reduction strategy, the sensor data are too spread, it stops; otherwise it moves a fixed distance further, curving around the obstacle. robot moves a certain distance, then monitors for the obstacle. If the obstacle is visible proportional reduction. The one clear exception is the w=3.0, range=1-4 case. Here, the several high-monitoring-error situations, but its fitness was slightly lower than mention, however, that periodic monitoring was indeed discovered by the algorithm in dominant for high values of monitoring error was not confirmed. It is important to proportional reduction was used. Our hypothesis that periodic monitoring would become shows. In fact, in all but one of the cases where a monitoring strategy was found,

Some form of proportional reduction was found in nearly all situations, as Figure 4

imagine how this could be advantageous:

path lengths and no sensor error, programs still monitor more than once. It is hard to the goal is not worth the energy necessary to achieve it. Another surprise is that for longer apparently, over such short distances, the extra effort involved in coming extremely close to We had expected the "monitor only once" strategy only in the absence of sensor noise, but

Interestingly, no program monitors more than once when the path length is very small.

strategy as beam length increases might be worth looking into: proportional reduction can do. But the higher rate of occurrence for the disproportional one opens question if either of these strategies is really doing some useful, beyond what pure approaches the goal, the second moves more quickly while it is still further away. It is an unmodified proportional reduction. The first moves relatively greater distances as if proportional reduction. They seem to be quite complementary in their relationship to strategies emerged - ones we hadn't foreseen: disproportional reduction and squared reduction, although it did come close in the $w=5.0$ cases. Two interesting new monitoring emerge, but predominantly not in its pure form. Periodic monitoring never beat proportional w and beam length increase, is not so clear-cut. The proportional reduction strategy did is no monitoring error, and that periodic monitoring will surpass proportional reduction as

The evidence for our other hypotheses, that the robot will monitor only once when there

alone.

strongly as the robot only monitors once here. Even so, it is still close to the previous $.4$ case. The constant of $.2$ in the $w=5.0$ beam length = 1.4 case should not be weighed too far of the constant decreases from $.4$ in two of the $w=0.2$ cases to $.4$ in the second $w=1.0$ data points. But the few applicable cases do show the predicted effect (see Table 1). The only really defined in cases of unmodified proportional reduction, which limits our available the goal that the robot actually moves before monitoring again. Of course this constant is grows. The proportionality constant is the proportion of the robot's estimated distance to

Our second hypothesis was that the proportionality constant should decrease as w

with proportional reduction alone.

proportional reduction and no monitoring error, the robot can reach the goal faster than adjusted to fit the particular range of beam lengths in that situation. Perhaps with squared strategy. The second case also seems to be a form of over-adaptation: the constant is again lengths. As it is expressed here, we do not see squared proportional reduction as a distinct adjusted to provide a good estimate of the distance remaining given the specific beam constant (which corresponds to the number of MOVES within the doubly nested loop) can be where it occurs, the agent only monitored once, or $w=0$. In the first case, the proportionality As shown in Table 1, this strategy was learned only occasionally. In all the situations

MOVE
 GOOP (EVA S) times:
 GOOP (EVA S) times:
 MONITOR: EVA S

loops within each other after the MONITOR command:

moves a proportion of the squared distance remaining. It is realized by nesting two EVA S- "squared proportional reduction". In this strategy, the robot monitors repeatedly, and

Another variant of the original proportional reduction scheme is something we termed

They did not cause much harm, and were therefore not removed in later generations: the extra MOVE commands that cause it are simply a relic of the evolutionary process:

across domains.

rules of when to use a specific strategy: rules that are general enough to be applicable justice, many further experiments are necessary. The ultimate goal must be to find general how the environment determines what is a good monitoring strategy. To do this topic and clearly the work presented here is only a first step in achieving an understanding to

the environment determine their success;

what form? For what other situations are the discovered strategies useful? What features of freedom to realize more sophisticated strategies, will proportional reduction still emerge? In do it they don't have the constructs to implement proportional reduction. Or when given the

As a follow-up experiment to this one, it would be interesting to see how well robots can

more neutral evaluation.

other domains on similar problems, thus compensating for over-fitting, could provide a is the task at hand, there is no absolute by which to judge strategies. Testing strategies in actually an interesting result, or merely sub-optimal. As finding the best possible solution results. In our experiment, it was often difficult to determine whether a new strategy is

Genetic algorithms are however challenging in that they do not guarantee optimal

evaluating these blocks in terms of their global usefulness.

building the idea of learning to construct programs from clearly identified sub-blocks, and modularity of programs, combining them in different ways to find novel solutions. We are part behavioral control problems such as monitoring. It takes particular advantage of the

Genetic programming might turn out to be a very robust way of finding solutions to

Discussion and Future Work

strategy, choosing instead a scheme of passing by the goal point and not monitoring at all. seems that these particular individuals had evolved beyond the proportional reduction proportional reduction, except the actual (and very expensive) monitoring command. If monitoring program in the $m=50$, $h=10$ case had all the constructs to do large, programs emerge that try to do without monitoring. It was interesting that the non-

Another phenomenon was also quite surprising: When the monitoring error becomes too

- Massachusetts, Amherst:
 Experimental Knowledge Systems Laboratory, Computer Science Dept., Univ. of
- [1] Hansen, E.A., 1985. Note on monitoring subcases. EKSL Memo #55.
 Development, Vol. 20, Pp. 125-104.
 the oven. Prospective memory, strategic time-monitoring, and context. Child
- [6] Ceci, S.L. & Bronfenbrenner, U., 1982. "Don't forget to take the subcases out of
 (Gregory J.E. Rawlings ed.). Morgan Kaufman, San Mateo, CA.
 Schemes Used in Genetic Algorithms, in Foundation of Genetic Algorithms
- [2] Goldberg, D.E. & Kalyanmool, D., 1981. A Comparative Analysis of Selection
 Learning Addison-Wesley, Reading, MA.
- [4] Goldberg, D.E., 1989. Genetic Algorithms in Search, Optimization & Machine
 Programming. AAAI-85, Pp. 104-501
- [3] Koza, J.B. & Bice, J.P., 1985. Automatic Programming of Robots using Genetic
 Means of Natural Selection and Genetics. MIT Press, Cambridge, MA.
- [5] Koza, J.B., 1985. Genetic Programming: On the Programming of Computers by
 preparation for AI Magazine.
- [1] Hansen, E.A. & Cohen, P.B., 1983. Monitoring plan execution: A survey. In

REFERENCES

governmental purposes notwithstanding any copyright notation hereon.
 United States Government is authorized to reproduce and distribute reprints for
 This research was supported by DARPA/AFOSR contract F30602-81-C-0030. The

ACKNOWLEDGEMENTS