A MULTISTRATEGY LEARNING SCHEME FOR AGENT KNOWLEDGE ACQUISITION

Diana Gordon
Naval Research Laboratory, Code 5514
Washington D.C. 20375
gordon@aic.nrl.navy.mil

Devika Subramanian
Department of Computer Science
Cornell University
Ithaca, NY 14853
devika@cs.cornell.edu

Keywords: multistrategy learning, advice taking, compilation, operationalization, genetic algorithms

Edited by: Gheorghe Tecuci

Received: May 26, 1993 **Revised:** October 8, 1993 **Accepted:** October 15, 1993

The problem of designing and refining task-level strategies in an embedded multiagent setting is an important unsolved question. To address this problem, we have developed a multistrategy system that combines two learning methods: operationalization of high-level advice provided by a human and incremental refinement by a genetic algorithm. The first method generates seed rules for finer-grained refinements by the genetic algorithm. Our multistrategy learning system is evaluated on two complex simulated domains as well as with a Nomad 200 robot.

1 Introduction

The problem of designing and refining task-level strategies in an embedded multi-agent setting is an important unsolved question. To address this problem, we have developed a multistrategy learning system that combines two learning methods: operationalization of high-level advice provided by a human, and incremental refinement by a genetic algorithm (GA). We define advice as a recommendation to achieve a goal under certain conditions. Advice is considered to be operationalized when it is translated into stimulus-response rules directly usable by the agent. Operationalization generates seed rules for finer-grained refinements by a GA.

The long term goal of the work proposed here is to develop task-directed agents capable of acting, planning, and learning in worlds about which they do not possess complete information. These agents refine factual knowledge of the world they inhabit, as well as strategic knowledge for achieving their tasks, by interacting with the environment. *Agent knowledge acquisition* is desirable for the same reasons that knowledge acquisition for expert systems is. It is easier for a user to provide high-level knowledge about the world and the task than to provide knowledge at a lower level of detail. The latter is well-known to be a costly, tedious, and error-prone process. Although agent knowledge acquisition is desirable, it is very difficult (for the agent), as is knowledge acquisition for expert systems. The additional challenge for agent knowledge acquisition comes from the fact that the knowledge must be dynamically updated by the agent through its interactions with the environment.

There are two basic approaches to constructing agents for dynamic environments. The first decomposes the design into stages: a parametric design followed by refinement of the parameter values using feedback from the world in the context of the task. Several refinement strategies have been studied in the literature: GAs [21], neural-net learning [3], statistical learning [13], and reinforcement learning [14]. The second, more ambitious, approach [7, 30] is to acquire the agent knowledge directly from example interactions with the environment. The success of this approach is tied to the efficacy of the credit assignment procedures, and whether or not it is possible to obtain good training runs with a knowledge-impoverished agent.

We have adopted the first approach. The direction we pursue is to compile an initial parametric agent using high-level strategic knowledge (e.g., advice) input by the user, as well as a body of general (not domain-specific) spatial knowledge in the form of a *Spatial Knowledge Base* (SKB). The SKB contains qualitative rules about movement in space. Example rules in our SKB are `If something is on my side, and I turn to the other side, I will not be facing it` and `If I move toward something it will get`

`closer`. This SKB is portable because it is applicable to a variety of domains where qualitative spatial knowledge is important. A similar qualitative knowledge base was constructed by [18] for the task of pushing objects in a plane. Since the knowledge provided to our agent is imperfect (incomplete and incorrect), our agent refines the knowledge using a GA by directly interacting with the world.

First we describe our deductive advice operationalization process and the nature of the parameterization adopted for our agent. Then we describe the inductive (GA) refinement stage and compare our multistrategy approach with one that is purely inductive. Before we present the details of the method, we characterize the class of environments and tasks for which we have found this decomposition of an agent design into an initial parametric stage and subsequent refinement stage to be effective.

- *Environment characteristics:* Complete models of the dynamics of the environment in the form of differential equations or difference equations, or discrete models like STRIPS operators, are unavailable. An analytical design that maps the percepts of an agent to its actions (e.g., using differential game theory or control theory) in these domains is thus not possible. Even if a model were available, standard analytical methods for deriving agents are extensional and involve exploration of the entire state space. They are inapplicable here because the domains we consider have of the order of a hundred million states.

- *Task characteristics:* Tasks are sequential decision problems: payoff is obtained at the end of a sequence of actions and not after individual actions. Examples are pursuit-evasion in a single or multi-pursuer setting and navigating in a world with moving obstacles. The tasks are typically multi-objective in nature: for instance for pursuit-evasion, the agent needs to minimize energy consumption while maximizing the time till capture by the pursuers.

- *Agent characteristics:* The agent has imperfect sensors. Imperfections occur in the form of noise, as well as incompleteness (all aspects of the state of the world cannot be sensed by our agent, a problem called *perceptual aliasing* in [32]). Stochastic differential game theory has methods for deriving agents with noisy sensors, but it requires detailed models of the noise as well as a detailed model of the environment and agent dynamics.

The action set of the agents and the values taken on by sensors are discrete and can be grouped into equivalence classes. This is the basis for the design of the parametric agent. A similar intuition underlies the design of fuzzy controllers that divide the space of sensor values into a small set of classes described by linguistic variables.

In domains with characteristics such as those just described, human designers typically derive an initial solution by hand and use numerical methods (usually very dependent on the initial solution) to refine their solution. Our ultimate objective is to automate the derivation of good initial solutions by using general knowledge about the environment, task, and agent characteristics and thus provide a better starting point for the refinement process. We begin with the SKB and advice.

2 Compiling Advice

Our operationalization method compiles high-level domain-specific knowledge (e.g., advice) and spatial knowledge (SKB) into low-level reactive rules directly usable by the agent. The compilation performs *deductive concretion* [16] because it deductively converts abstract goals and other knowledge into concrete actions. An important question is why we adopt a deductive procedure for the operationalization of advice. At this time, we are able to generate good parametric designs with deductive inference alone. We expect that as we expand our experimental studies to cover more domains, the incompleteness of the SKB will force us to adopt more powerful operationalization methods.

The advice compilation problem can be stated as follows:

Given:

- Strategic advice of the form: *Set of States* $\rightarrow$ *achieve(Goal)*, which recommends that *Goal* be achieved (be made to hold in the world) whenever the agent finds itself in a state that matches the antecedent of the advice. Note that there is no requirement that the set of states in the antecedent be directly perceivable. The agent may have to plan to acquire the information needed to determine whether or not it is in the specified set of states.

- Qualitative factual knowledge about the domain in the form of a description of a *Set of States*.

- A qualitative theory with information of the following types:

  (1) Domain-specific rules. *a.* A set of domain-specific terminological mappings that define terms needed for interpreting the advice: these mappings are of the form *Set of states* $\rightarrow$ *Set of states* and are vocabulary conversion rules. *b.* A set of state transition rules that represent the dynamics of the world: these mappings are of the form *Set of states* $\times$ A $\rightarrow$ *Set of states*, where A is the set

of actions. *c*. A set of rules for predicting other agent's actions: these mappings are of the form *Set of States* $\times A \rightarrow A$. These rules are special to multi-agent domains.

(2) A general qualitative theory of movement (SKB). This is a common subpart of the domains that we consider here. Therefore it need not be reintroduced by the user for each new domain. The SKB is of the form *Set of states* $\times A_M \rightarrow$ *Set of states*, where A_M is the subset of actions A that involve movement.

– The operational sensors P and actions A of the agent. A detailed agent design involves finding a mapping from the sense history P^* of the agent to A. This is usually simplified to be of the form $2^P \rightarrow A$ so that only current percepts determine the next action. The map need not be deterministic. Indeed, in the GA architecture, which is our target architecture, the agent map is non-deterministic.

Find: An operational mapping $2^P \rightarrow A$ from the sensors to the effectors of the agent such that the map implements the provided strategic advice. By this we mean that the map tests the conditions under which the recommended goal is to be achieved, and the action sequence it generates in the world places the agent in a state in which the goal is achieved.

In summary, advice tells the agent to achieve a goal under specified conditions. Facts, domain-specific rules, a general theory of movement, and a list of what is operational are all required for interpreting this advice and compiling it into an operational form.

As an example, consider the problem of an agent that tries to evade a pursuing adversary. (For a more detailed description of this problem, see Section 4.3.) We assume the adversary is initially faster than the agent, though it is less maneuverable. Both the adversary and the agent lose speed as they turn, and the adversary's speed loss is permanent. The agent and the adversary can sense each other's instantaneous position and velocity to within some specified accuracy. Although these are facts of the problem, they are initially unknown to the agent.

The strategic advice offered by a human expert for this problem is the following: if the adversary's speed is high, then try to slow down the adversary; if the adversary's speed is low, then avoid the adversary. In this example the preconditions of the advice involve predicates that the agent can sense; however it needs to pin down the semantics of *high* and *low*. The goals to be achieved are: slow down the adversary in the one case, and avoid the adversary in the other. These are not directly implementable using a single action from A. We need to devise a plan to realize these goals using knowledge about the domain, including knowledge about movements. These plans form the initial parametric design with *high* and *low* being two of the parameters whose ranges need to be learned by interaction with the environment. An example SKB rule that is used to compile our advice is:

```
IF bearing(agent,adversary) = right AND
   turn(agent) = left
THEN heading(adversary,agent) ≠ headon.
```

`Heading(`X`,`Y`)` refers to the direction of motion of Y relative to X, and `bearing(`X`,`Y`)` refers to the direction of Y relative to X. This SKB rule states that an agent no longer faces an adversary when it turns away from it. `Turn` is an operational action.

To compile the advice, we also need terminological mappings that define terms such as *avoid*. An example definition needed to compile our advice is:

```
IF range(X,Y) ≠ close AND
heading(Y,X) ≠ headon
THEN avoids(X, Y).
```

Avoids(X, Y) means X avoids Y. Factual knowledge is also required for compilation, e.g.,

```
Moving(agent).
Moving(adversary).
```

From our advice, the set of operational sensors and effectors, and our qualitative theory, compilation results in a set of operational sensor-effector mappings. The set of all these operational sensor-effector mappings is what we consider to be a *reactive plan* (agent map). A portion of the reactive plan that implements our advice, and is the product of our compilation procedure, is shown below.

```
IF speed(adversary) = high AND
   range(agent,adversary) = close
THEN turn(agent) = hard
IF speed(adversary) = high AND
   range(agent,adversary) ≠ close AND
   bearing(agent,adversary) = left
THEN turn(agent) = left
IF speed(adversary) = high AND
   range(agent,adversary) ≠ close AND
   bearing(agent,adversary) = right
THEN turn(agent) = right
IF speed(adversary) = low AND
   bearing(agent,adversary) = right
THEN turn(agent) = left
IF speed(adversary) = low AND
   bearing(agent,adversary) = left
THEN turn(agent) = right
```

This is a fairly complex plan. It implements the strategic advice in the following manner: when the adversary's speed is high, the agent moves toward the adversary and makes a hard turn when it is close enough. Since the adversary is chasing the agent, this plan will

cause the adversary to turn hard, which then causes the adversary to permanently lose speed. If the adversary is moving slowly, the agent simply stays out of its way. In this section we present an algorithm that generates plans of this form starting from high level advice and a qualitative domain theory.

Our compilation algorithm is tailored for situations in which the execution of plans, but not the learning of plans, is highly time-critical. We assume the advice provided is operationalized immediately, though it need not be applied immediately. We precompile all high-level knowledge into an operational form because the agent will apply it in a time-critical situation. In our approach, advice and SKB rules have nonoperational elements, and the compilation process results in rules that are fully operational.

The compilation algorithm, shown in Figure 1, uses two stacks: a GoalStack and an (operational condition) OpCondStack. Four types of knowledge are initially given to the compiler: advice, facts, nonoperational rules (abbreviated *nonop rules*), which include domain-specific rules and the SKB, and the set of operational sensors and actions. The output from compilation is a set of *op rules* directly usable by the agent, i.e., operational. Advice has the form:

```
<IF cond AND ... AND cond THEN> ACHIEVE goal.
```

Facts have the form:

`Predicate(`X_1`, ...,` X_n`)`

Nonop rules have the form:

```
IF cond AND ... AND cond <AND action>
THEN goal.
```

The set of operational sensors and actions are presented to the agent as a list.

Anything in angle brackets is optional. The portion preceding the `THEN` is the rule *antecedent*, and the portion following the `THEN` is the rule *consequent*. A nonop rule consequent is a single goal. The syntax for a goal is `function(`$X_1,\ldots,X_n$`)` $=/\neq$ `value` or `predicate(`$X_1,\ldots,X_n$`)`. Each X_i is an object (e.g., an agent). The syntax for a `cond` (condition) or `action` in the rule antecedent is the same as for goals.

Although advice has a similar syntax to nonop rules, its interpretation differs. Advice recommends achieving the given `goal` under the given `conds`. A nonop rule, which is needed to compile this advice, states that the given `goal` will be achieved if the given `conds` (and `action`) occur. Compilation results in stimulus-response op rules of the form:

```
IF cond AND ... AND cond THEN action.
```

Push advice on GoalStack: goal followed by conditions.
Initialize OpCondStack to be empty and invoke
  Compile(GoalStack,OpCondStack).
Procedure Compile (GoalStack, OpCondStack)
  if GoalStack is not empty **then**
    g ← pop(GoalStack);
    case g:
      1. g is an operational condition:
          Push(g, OpCondStack);
          Compile(GoalStack, OpCondStack);
      2. g matches a fact:
          Compile(GoalStack,OpCondStack);
      3. g is nonoperational:
          foreach nonop rule R_i whose consequent matches g **do**
            Push(antecedent(R_i), GoalStack);
            Compile(GoalStack,OpCondStack);
      4. g is an operational action
          Form a new op rule from the contents of OpCondStack and g;
          Clear OpCondStack;
  else Clear OpCondStack;

Figure 1: Algorithm for operationalizing advice

The conditions of an op rule are sensor values detectable by the agent. The action can be performed by the agent.

Our compilation algorithm in Figure 1 takes advice and backchains through the SKB and user-provided nonop rules until an operational action is found. Once an operational action is found, it pops back up the levels of recursion, attaching operational conditions along the way, to form a new reactive agent op rule. To prevent cycles, the last nonop rule used in step 3 is marked as "used" so that it will not be used again.

Let us examine a simple example of how this algorithm operates, as shown in Figures 2 and 3. Consider the advice which advocates avoiding the adversary when the adversary's speed is low.

```
IF speed(adversary) = low THEN
ACHIEVE heading(adversary,agent) ≠ headon
```

Figure 2 shows how SKB nonop rules match this advice for backchaining, thereby creating an "and" tree. Anything preceded by a "*" is operational.

Figure 3 shows this algorithm in operation. Note that stacks grow downward. The algorithm begins by pushing the advice goal, followed by the advice condition, on the GoalStack. It then calls procedure Compile, which moves the advice condition to the OpCondStack because it is operational. The advice goal is not operational. In our example, the advice goal can be unified with the goal of SKB RULE1, which states

```
IF bearing(agent,adversary) = right AND
```

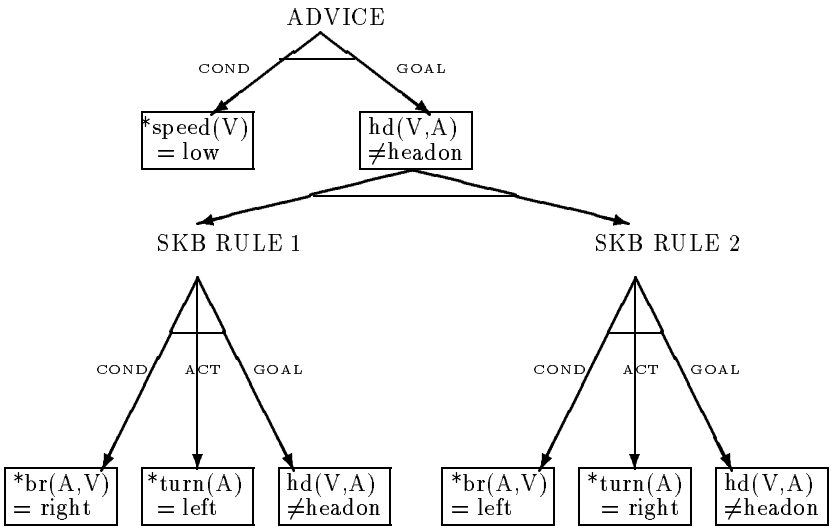

Figure 2: Graph of example. A is short for agent, V for adversary, br for bearing and hd for heading.

```
   turn(agent) = left
THEN heading(adversary,agent) ≠ headon.
```

The condition and action of RULE1 are pushed on the GoalStack. Because the condition of RULE1 is operational, it is moved to the OpCondStack. Note that if this condition had not been operational, our compilation algorithm would have searched for all nonop rules whose consequents unify with this nonoperational condition. Further backchaining would have continued until operational conditions were found.

At this point, the action of RULE1 is at the top of the GoalStack, and it is operational, so we can create an op rule. The conditions from the OpCondStack are added to the action. This creates an op rule that states

```
IF speed(adversary) = low AND
   bearing(agent,adversary) = right
THEN turn(agent) = left
```

Both stacks are cleared. The algorithm (see Figure 2) continues similarly to generate a second op rule from SKB RULE2 that states

```
IF speed(adversary) = low AND
   bearing(agent, adversary) = left
THEN turn(agent) = right
```

Next, we apply a conversion from qualitative to quantitative op rules. The rules are given default quantitative ranges. For example, if `speed` has two values, `slow` and `fast`, we bisect the range of all possible values into two subranges. Then, we allow the system to improve this initial choice of quantitative ranges by using a GA to refine the initial ranges while interacting with the environment.

3 Executing and Refining Advice

The system we use to refine and apply the op rules derived from our compiled advice is the SAMUEL reactive planner [7]. We have chosen SAMUEL because this system has already proven to be highly effective for refining rules on complex domains [7, 25]. SAMUEL adopts the role of an agent in a multiagent environment in which it senses and acts. This system has two major components: a performance module and a learning module. Section 4.2 explains how performance interleaves with learning in our experiments.

The performance module, called the Competitive Production System (CPS), interacts with a simulated or real world by reading sensors, setting effector values, and receiving payoff from a critic. CPS performs matching and conflict resolution on the set of op rules. This performance module follows the match/conflict-resolution/act cycle of traditional production systems. Time is divided into *episodes*: the choice of what constitutes an episode is domain-specific. Episodes begin with random initialization and end when a critic provides payoff. At each time step within an episode, CPS selects an action using a probabilistic voting scheme based on rule strengths. All rules that match (or partially match - see [7]) the current state bid to have their actions fire. The actions of rules with higher strengths are more likely to fire. If the world is being simulated, then after an action fires, the world model is advanced one simulation step and sensor readings are updated.

CPS assigns credit to individual rules based on feedback from the critic. At the end of each episode, all rules that suggested actions taken during this episode have their strengths incrementally adjusted to reflect the current payoff. Over time, rule strengths reflect

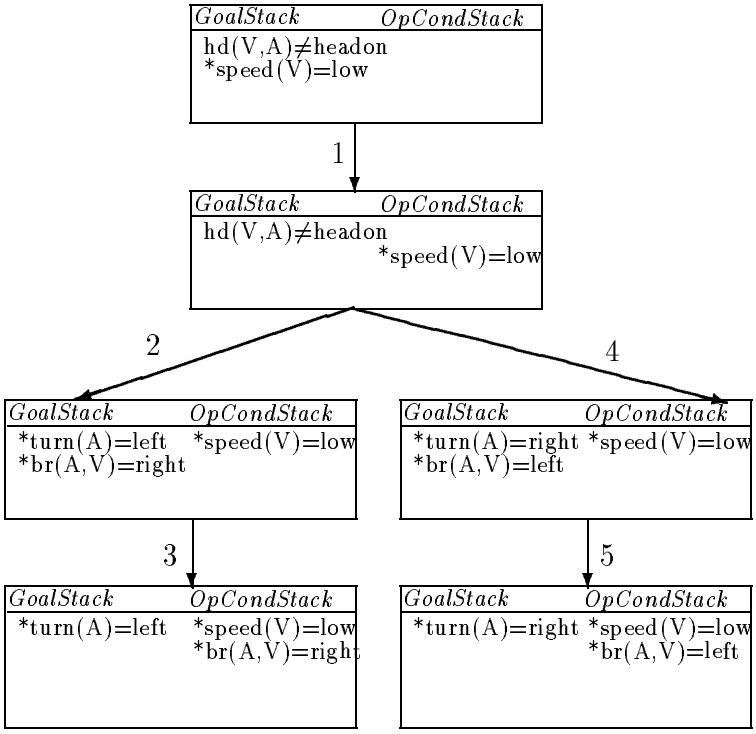

Figure 3: Example of compilation algorithm

the degree of usefulness of the rules.

SAMUEL's learning module is a genetic algorithm. GAs are motivated by simplified models of heredity and evolution in the field of population genetics [8]. GAs evolve a population of individuals over a sequence of *generations*. Each individual acts as an alternative solution to the problem at hand, and its *fitness* (i.e., potential worth as a solution) is regularly evaluated. During a generation, individuals create *offspring* (new individuals). The fitness of an individual probabilistically determines how many offspring it can have. Genetic operators, such as *crossover* and *mutation*, are applied to produce offspring. Crossover combines elements of two individuals to form two new individuals; mutation randomly alters elements of a single individual. In SAMUEL, an individual is a set of op rules, i.e., a reactive plan. In addition to genetic operators, this system also applies non-genetic knowledge refinement operators, such as *generalize* and *specialize*, to op rules within a rule set.

The interface between our compilation algorithm and the SAMUEL system is straightforward. The output of our compilation algorithm is a set of op rules for the SAMUEL agent. Because the op rules may be incomplete, a *random rule* is added to this rule set. The random rule recommends performing a random action under any conditions. The final rule set, along with CPS and the GA learning component for improving the rules, is our initial agent.

4 Evaluation

We have not yet analyzed the cost of our compilation algorithm. The worst case cost appears to be exponential in the sizes of the inputs to the algorithm because the STRIPS planning problem can be reduced to it. In the future, we plan to investigate methods to reduce this cost for complex realistic problems. Potential methods include: (1) attaching a likelihood of occurrence onto advice, which enables the agent to prioritize which advice to compile first if time is limited, (2) tailoring the levels of generality and abstraction of the advice to suit the time available for compilation (e.g., less abstract advice is closer to being operational and therefore requires less compilation time), and (3) generating a parallel version of the algorithm.

We have evaluated our multistrategy approach empirically. We focus on answering the following questions:

1. Will our advice compilation method be effective for a reactive agent on complex domains?

2. Will the coordination of multiple learning techniques lead to improved performance over using any one learning method? In particular, we want the GA to increase the success rate of the compiled advice, and the advice to reduce the *convergence time* of the GA. Success and failure have domain-specific definitions (see Sections 4.3 and 4.4). The convergence time is defined as the time

taken to ascend to a given success rate. A reduced convergence time is useful when learning time is limited.

3. Can we construct a portable SKB?

4.1 Domain characteristics

To address our questions, we have run experiments on two complex problems: Evasion and Navigation. Our choice of domains is motivated by the results of Schultz and Grefenstette, who have obtained large performance improvements by initializing the GA component of SAMUEL with hand-coded op rules in these domains [25]. Their success has inspired this work: our objective is to automate their tedious manual task, and the work described here is one step toward the goal.

Both problems are two-dimensional simulations of realistic tactical problems. However, our simulations include several features that make these two problems sufficiently complex to cause difficulties for more traditional control-theoretic or game-theoretic approaches [7]:

- *A weak domain model*: The learner has no initial model of other agents or objects in the domain. Most control-theoretic and game-theoretic models make worst case assumptions about adversaries. This yields poor designs in the worlds we consider because we have statistical rather than worst case adversaries.

- *Incomplete state information*: The sensors are discrete, which causes perceptual aliasing: a many-to-one mapping from actual states to perceived states.

- *A large state space*: The discretization of the state space makes the learning problem combinatorial. In the Evasion domain, for instance, over 25 million distinct feature vectors are observed.

- *Delayed payoff*: The critic only provides payoff at the end of an episode. Therefore a credit assignment scheme is required.

- *Noisy sensors*: Gaussian noise is added to all sensor readings. Noise consists of a random draw from a normal distribution with mean 0.0 and standard deviation equal to 5% of the legal range for the corresponding sensor. The value that results is discretized according to the defined granularity of the sensor. A 5% noise level is sufficient to slightly degrade SAMUEL's performance.

4.2 Experimental design

Two sets of experiments are performed on each of the two domains. Perception is noise-free for the first set, but noisy for the second. The primary purpose of the first set is to address our question about the effectiveness of our advice compilation method alone, without GA refinement. Facts, nonop rules, advice, and the set of operational sensors and actions are given to the compiler and the output is a set of op rules. The random rule is added to the op rules and this rule set is given to SAMUEL's CPS module to be applied within the simulated world model. The baseline performance with which these rules are compared is the random rule alone. These experiments measure how the success rate of the compiled rules compares with that of the baseline as problem complexity increases, where the success rate is an average over 1000 episodes. Statistical significance of the differences between the curves with and without advice are presented. Significance is measured using the large-sample test for the differences between two means.

The primary purpose of the second set of experiments is to address our question about the effectiveness of the multistrategy approach (compilation followed by GA refinement). We used the same set of op rules (i.e., the output of the compiler) for this second set of experiments as was used for the first set of experiments. This rule set, plus the random rule, becomes every individual in SAMUEL's initial GA population, i.e., it seeds the GA with initial knowledge. The baseline performance with which these rules are compared is SAMUEL initialized with every individual equal to just the random rule. In either case, GA learning evolves this initial population. In other words, we compare the performance of advice seeding the GA with GA learning alone (i.e., random seeding). Random seeding produces an initially unbiased GA search; advice initially biases the GA search - hopefully into favorable regions of the search space.

In this second set of experiments, performance of the agent in the environment interleaves with GA refinement. SAMUEL runs for 100 generations using a population size of 100 rule sets. Every 5 generations, the *best* (in terms of success rate) 10% of the current population are evaluated over 100 episodes to choose a single plan to represent the population. This plan is evaluated on 1000 randomly chosen episodes and the success rate is calculated. This entire process is repeated 10 times and the final success rate, averaged over all 10 trials, is found. The curves in our graphs plot these averages. For this set of experiments, statistical significance is measured using the two-sample t-test, with adjustments as required whenever the F statistic indicates unequal variances.

We add sensor noise, as defined in Section 4.1, for this second set of experiments because GAs can learn robustly in the presence of noise [7]. Two performance measures are used: the success rate and the convergence time. The convergence time is defined as the number of GA generations required to achieve and

maintain an n% success rate, where n is different for each of the two domains. The value of n is set empirically.

4.3 Evaluation on the Evasion problem

Our simulation of the Evasion problem is partially inspired by [4]. This problem was introduced in Section 2. The problem consists of an agent that moves in a two-dimensional world with a single adversary pursuing the agent. The agent is controlled by SAMUEL, and the adversary is controlled by a simple set of rules. The agent's objective is to avoid contact with the adversary for a bounded length of time. Contact implies the agent is captured by the adversary. The problem is divided into episodes that begin with the adversary approaching the agent from a random direction. The adversary initially travels faster than the agent, but is less maneuverable (i.e., it has a greater turning radius). Both the agent and the adversary gradually lose speed when maneuvering, but only the adversary's loss is permanent. An episode ends when either the adversary captures the agent (failure) or the the agent evades the adversary (success). At the end of each episode, a critic provides full payoff for successful evasion and partial payoff otherwise, proportional to the amount of time before the agent is captured. The strengths of op rules that fired are updated in proportion to the payoff.

The agent has the following operational sensors: time, last agent turning rate, adversary speed, adversary range, adversary bearing, and adversary heading. The agent has one operational action: it can control its own turning rate. For further detail, see [7].

In our experiments, we provide the following domain-specific knowledge to the agent:

```
FACTS:
Chased-by(agent,adversary).
Moving(agent).
Moving(adversary).

DOMAIN-SPECIFIC NONOP RULES:
IF chased-by(X,Y) AND
   range(X,Y) = close AND
   turn(X)=Z
THEN turn(Y) = Z.
IF range(X,Y) ≠ close AND
   heading(Y,X) ≠ headon
THEN avoids(X,Y).
IF turn(adversary) = hard
THEN decelerates(adversary).

ADVICE:
IF speed(adversary) = high THEN
ACHIEVE decelerates(adversary).
IF speed(adversary) = low THEN
ACHIEVE avoids(agent,adversary).
```

We also include knowledge of the agent's operational sensors and actions as facts. Ten SKB nonop rules are used (they are instantiations of the two rules described in English in the introduction of this paper). Although room does not permit listing them all, some examples are:

```
IF bearing(X,Y) = right AND
   turn(X) = left
THEN heading(Y,X) ≠ headon.
IF bearing(X,Y) = left AND
   moving(X) AND
   turn(X) = left
THEN range(X,Y) = close.
```

From our input and our SKB rules, the compilation method of Section 2 generates op rules. The sensor values of these rules are translated from qualitative values to default quantitative ranges. For example, `bearing = left` is translated into `bearing = [6..12]`, where the numbers correspond to a clock, e.g., 6 means `6 o'clock`. Every new rule is given an initial strength of 1.0 (the maximum). The final op rule set includes rules such as[1]

```
IF speed(adversary) = [700..1000] AND
   range(agent,adversary) = [0..750]
THEN turn(agent) = hard-left.
```

The total number of op rules generated from our advice is 16.

We begin our experiments by addressing the first question, which concerns the effectiveness of our advice-taking method. We do not use the GA. Problem difficulty is varied by adjusting a *safety* envelope around the agent. The safety envelope is the distance at which the adversary can be from the agent before the agent is considered captured by the adversary.

Figure 4 shows how the performance (averaged over 1000 episodes) of these op rules compares with that of just the random rule. All of the differences between the means are statistically significant (using significance level $\alpha = 0.05$). From Figure 4 we see that from difficulty levels 80 to 120, the agent is approximately twice as successful with advice than without it. This is a 100% performance advantage. Furthermore, for levels 130 to 160, the agent is about four times more effective with advice. For levels 160 to 200, the agent is an order of magnitude more effective with advice. We conclude that as the difficulty of this problem increases, the advice becomes more helpful relative to the baseline. These results answer our first question: our advice compilation method is effective on this domain.

We address the second question about multistrategy effectiveness by combining the compiled advice with

[1] To generate a few of these rules, we used a variant of our compilation algorithm. We omitted a description of this variation for the sake of clarity. See [6] for details.

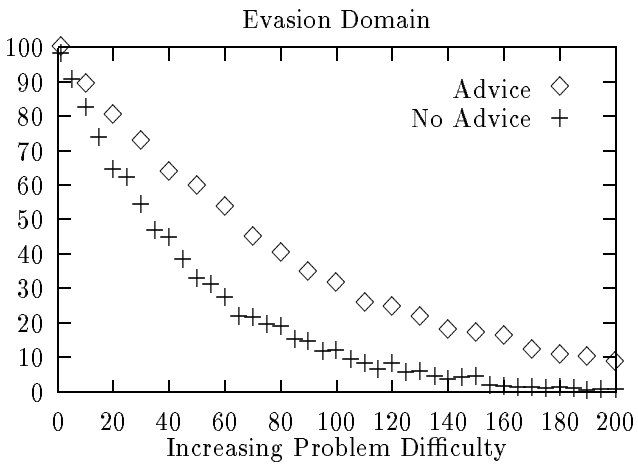

Figure 4: Results on the Evasion domain without GA refinement. Y axis denotes success rate.

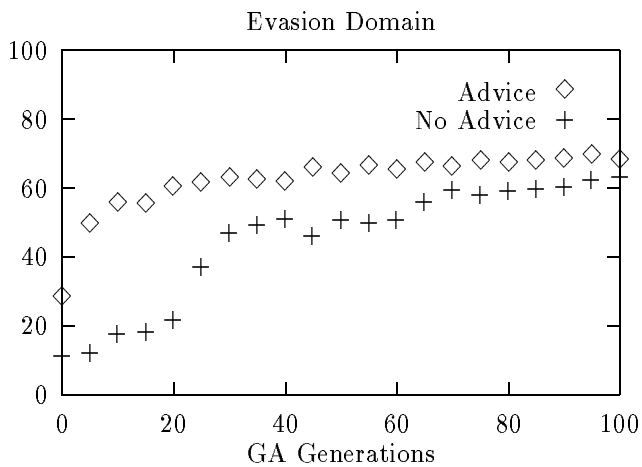

Figure 5: Results on the Evasion domain with the GA. Y axis denotes success rate.

GA refinement. Figure 5 shows the results of comparing the performance of the GA with and without advice. The safety envelope is fixed at 100 (chosen arbitrarily) and noise is added to the sensors. For this domain, the convergence time is the number of GA generations required to maintain a 60% success rate.

Figure 5 shows that in a small amount of time (approximately 10 generations), the GA doubles the success rate of the advice rules. Furthermore, the addition of advice produces a 3.5-fold reduction in the convergence time over using a GA with random initialization. The differences between the means are statistically significant for the first 65 generations, but not afterwards ($\alpha = 0.05$).[2] Apparently, the advice initially biases the GA into a favorable region of the search space, which improves the convergence time over random initialization. The convergence value, however, does not appear to be higher with advice than without it because eventually the performance of the randomly initialized GA is roughly equivalent to that of the GA seeded with advice.

4.4 Evaluation on the Navigation problem

In the Navigation domain, our agent is again controlled by SAMUEL in a two-dimensional simulated world. The agent's objective is to avoid obstacles and navigate to a stationary target with which it must rendezvous before exhausting its fuel (implemented as a bounded length of time for motion). Each episode begins with the agent centered in front of a randomly generated field of obstacles with a specified density. An episode ends with either a rendezvous at the target location (success) or the exhaustion of the agent's fuel or a collision with an obstacle (failure). At the end of an episode, a critic provides full payoff if the agent reaches the target, and partial payoff otherwise, depending on the agent's distance to the goal.

The agent has the following operational sensors: time, the bearing of the target, the bearing and range of an obstacle, and the range of the target. The agent has two operational actions: it can control its own turning rate and its speed. For further detail, see [24].

We provide the following domain-specific knowledge (in addition to a list of operational sensors and actions):

```
FACTS:
Moving(agent).

DOMAIN-SPECIFIC NONOP RULES:
IF range(X,Y) ≠ close AND
   heading(Y,X) ≠ headon
THEN avoids(X,Y).

ADVICE:
IF range(agent,obstacle) ≠ close THEN
ACHIEVE range(agent,target) = close.
IF range(agent,obstacle) = close THEN
ACHIEVE avoids(agent,obstacle).
ACHIEVE speed(agent) = high.
```

The *same* 10 SKB nonop rules used for Evasion are again used for this domain and are successful in the compilation procedure. This confirms the portability of our SKB for these two domains, thus addressing our third question. Also, although the definition of *avoids* was intended to be domain-specific, it actually applies

[2]These results differ from the results presented in [6]. This discrepancy results from an improvement in the initial rule strengths. Future work will focus on studies analyzing system sensitivity to parameter variations. Our initial studies indicate that performance is more robust with respect to changes in some parameters, e.g., variations in the qualitative to quantitative mappings, than to changes in others, e.g., the initial rule strengths.

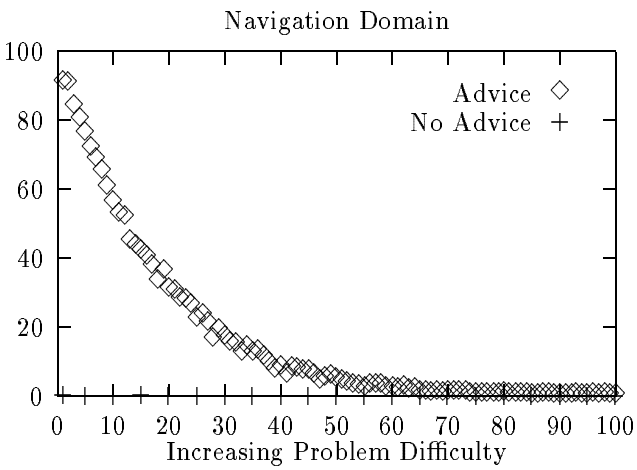

Figure 6: Results on the Navigation domain without GA refinement. Y axis denotes success rate.

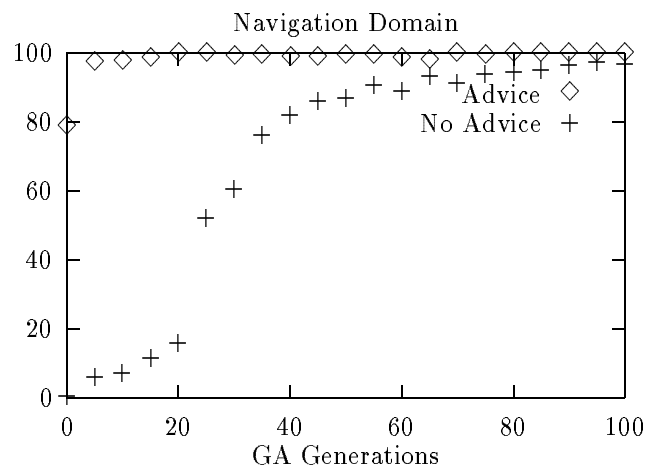

Figure 7: Results on the Navigation domain with the GA. Y axis denotes success rate.

to both of our domains. A total of 42 op rules are generated.[3]

Again, we address the first question by using SAMUEL without the GA. We increase the problem difficulty on this domain by increasing the number of obstacles. The success rate is an average over 1000 episodes. Without advice, the success rate is 0% because this is a very difficult domain. Figure 6 shows how we improve the success rate to as much as 90% by using our advice on this domain. We increase problem difficulty by increasing the number of obstacles. At all but the last few points, the differences between the means are statistically significant ($\alpha = 0.05$). When we vary the number of obstacles, performance follows a different trend than for the Evasion domain. By far the greatest benefit of the advice occurs when there are few obstacles. The success rate is 90% with advice and 0% without advice when there is only one obstacle, for example. The advantage drops as the problem complexity increases. After difficulty level 80, the particular advice we gave the agent no longer offers any benefit. For this level of problem difficulty, different advice is probably more appropriate.

Our experiments on both domains confirm that our advice compiler can be effective, however, they also indicate that the usefulness of advice may be restricted to a particular range of situations. Another learning task, which we are currently exploring, would be to identify this range and add additional conditions to the advice.

We address the second question by comparing the performance of the GA with and without advice. Noise is added. Figure 7 shows the results. All differences between the means are statistically significant ($\alpha = 0.05$). Here, the number of obstacles is fixed at five (chosen arbitrarily). For this domain, the convergence time is the number of GA generations required to maintain a 95% success rate.

Figure 7 shows that the addition of advice yields an enormous performance advantage on this domain. Figure 7 also shows that given a moderate amount of time (10 generations), the GA provides a 10% increase in the success rate. Furthermore, the addition of advice produces an 18-fold improvement in the convergence time over using GAs alone. Not only does advice improve the convergence *time*, but it also appears to improve the *level* of convergence: after 80 generations, the GA with advice holds a 99% or above success rate whereas after all 100 generations the GA without advice still cannot get above a 97% success rate.[4]

To further test our compilation method, we have recompiled our Navigation advice into op rules for a Nomad 200 mobile robot that is equipped with very noisy sonar and infrared sensors and can adjust its turning rate and speed. The sensors are so noisy that the robot sometimes mistakes two boxes four feet apart for a wall. The op rules that result from compilation have not been refined by the GA to develop a tolerance to noise; therefore, this noise poses a severe challenge.

The op rules are linked to a vendor-provided interface that translates the language of the SAMUEL rules (e.g., `IF asonar4 [17..85] THEN SET turn -400`) into joint velocity and servo motor commands. From high-level advice to avoid obstacles and rendezvous with a goal point, our method has compiled rules that enable the robot to succeed approximately a third of the time in avoiding three large boxes and reaching a goal point on the other side of a room. The same SKB rules are used for compilation. With the random

[3] We were able to decrease the number of op rules to 9 by making one careful qualitative to quantitative mapping choice.

[4] On both this domain and the one described in the previous section, our results are for 100 GA generations. In future work we plan to run the GA longer to determine whether our conclusions hold over a longer time period.

rule alone, it is extremely unlikely to successfully complete this task. Our next step will be to refine these robot rules using GAs within a simulation of this actual domain.[5]

In conclusion, our multistrategy system offers two advantages. First, it provides an initial boost from seeding with high-level knowledge. On both simulated domains we see a significant reduction in the convergence time - an order of magnitude on the Navigation domain. Second, the multistrategy system provides the robustness and improvement gained from GA refinement. Refinement yields a 10% increase in the success rate on Navigation and a 100% increase on Evasion. The fact that on both domains the advice biases the GA into a favorable region and significantly reduces its convergence time is very important in complex domains where a GA approach alone may not be feasible or may be very inefficient. Finally, the effectiveness of our advice-taking method has been confirmed on a robot with poor sensory capabilities.

To determine whether the presence of noise influences our conclusions, we have repeated our experiments on the two simulated domains with sensor readings that are not noisy. Although the success rates differ slightly, our conclusions still hold.

5 Related Work

This work relates most strongly to the following topics in machine learning: advice taking, combining projective and reactive planning, methods for compiling high-level goals into reactive rules, learning in fuzzy controllers, and multistrategy learning. This work also relates to research in differential game theory. We discuss each in turn.

5.1 Machine Learning

Advice taking has been considered by McCarthy as early as 1958 [15] and later by Mostow [19] and others. To date, research on assimilating advice in embedded agents has been limited but encouraging. Previous research has focused mainly on providing low-level knowledge. For example, Laird *et al.* [11] and Clouse and Utgoff [3] have had good success providing agents with information about which action to take. Chapman gives his agent high-level advice [2] . Our advice taker differs from Chapman's because it can operationalize advice long before the advice is applied and because it refines the advice with a GA. Most important of all, our advice taking method is unique because it involves a multistrategy approach that couples a knowledge-intensive *deductive* precompilation phase with an empirical *inductive* refinement phase.

[5]We wish to use SAMUEL both to handle the noise and because we had to manually refine the qualitative to quantitative mappings somewhat - SAMUEL could automate this.

We assume that high-level knowledge is operationalized but not applied immediately. Methods for operationalizing advice that will be applied immediately include STRIPS-like planners [20] and explanation-based learning (EBL) planners e.g., [26]. A closely related system is Mitchell's [17]. This system combines EBL projective planning with reactive planning. Our method for compiling goals is similar to that of EBL because it uses the notion of operationality. It differs because we do not assume that the advice will be applied immediately, and therefore our compilation method has no current state on which to focus plan generation. All of the above-mentioned methods create a projective plan to achieve a goal from the current state. We precompile advice for multiple possible states.

Because our method precompiles plans from possible states rather than from a current state, it is very similar to the methods of Schoppers [23] and Kaelbling [9] for compiling high-level goals into low-level reactive rules. Our method differs from those of Schoppers and Kaelbling because it includes the EBL notion of operationality. Also unlike Schoppers and Kaelbling, we use a refinement method following compilation.

Considerable prior work has focused on knowledge refinement. Others have used GAs to refine qualitative to quantitative mappings. For example, Karr uses GAs to select fuzzy membership functions for a fuzzy controller [10]. Lin, Mahadevan and Connell, and Singh initialize their systems with modular agent architectures then refine them with reinforcement learning [12], [14], and [27]. Lin trains a robot by giving it advice in the form of a sequence of desired actions. Mahadevan and Connell initialize their reinforcement learner with a prespecified subsumption architecture, and Singh guides his reinforcement learner by giving it abstract actions to decompose.

One of the most similar approaches to ours is that of Towell and Shavlik [31]. They also couple rule-based input with a refinement method; however, their refinement method is neural networks. This multistrategy system converts rules into a network topology. The content of each rule is preserved; therefore, the transformation is syntactic. Our multistrategy system, on the other hand, focuses primarily on semantic transformations that use qualitative knowledge about movements in space to convert abstract goals into concrete actions. Ram and Santamaria [22] present a multistrategy refinement scheme for a parametric agent performing a navigation task. In contrast with our genetic algorithms based approach, their system acquires values for the defined parameters by using reinforcement learning. Different parameter value combinations are acquired for distinct environment subclasses. The interesting feature of their system is the explicit identification of different environmental configurations using case-based reasoning. The identification of reg-

ularities in the environment occurs implicitly in the genetic algorithms approach. The deductive compilation scheme (but not the refinement) is in common with Mitchell *et al.*'s derivation of a strategy for pushing objects in a tray using a qualitative theory of the process [18].

5.2 Differential game theory

Differential game theory is a branch of mathematical optimal control theory. It assumes that the behavior of the controlled system can be modeled as a system of ordinary differential equations (ODEs). The evasion problem considered in this paper is a typical example of a differential game. In particular, the problem is two-person zero-sum differential game with a *constant terminal time*. Both the pursuer and the evader move in a bounded rectangle in two dimensions. The evader has to avoid getting to within a certain distance of the pursuer for a certain length of time. In the minimax formulation of the problem, the optimal strategy of the evader is one that achieves its objective under the least favorable assumptions on the motion of the pursuer.

Differential games are formulated mathematically by specifying the motion equations of the pursuer and evader, the class of admissible controls for both systems (which identifies the way in which the pursuer and evader can change their motions), and the target or goal functional. A classic reference for this is [1]. The focus of work in differential game theory is to identify conditions under which optimal strategies for the evader can be derived. This assumes complete knowledge of the dynamics of the evader and pursuer, both of which are unavailable to us. The theory would be more useful to us if it had a qualitative counterpart which allowed us to determine the existence of solutions to the evader's problem from partial knowledge of the evader and pursuer's dynamics.

6 Discussion

We have presented a novel multistrategy learning method for operationalizing and refining high-level advice into low-level rules to be used by a reactive agent. Operationalization uses a portable SKB. An implementation of this method has been tested on two complex domains and a Nomad 200 robot.

We have learned the following lessons:

1. Our advice compiler can be effective on complex domains, and it will be important to identify the regions of greatest effectiveness for advice,

2. A portable SKB appears feasible, and

3. Coordinating a deductive learning strategy (advice compilation) with an inductive learning strategy (GA refinement) can lead to a substantial performance improvement over either method alone.

This success, however, depends on the how the advice biases the GA search. Future work will focus on identifying those characteristics of advice that bias this search favorably. We will also focus on further addressing our questions about performance using different advice and alternative domains (e.g., [28]).

Many other interesting directions are suggested by our experimental results. At present we do not consider the cost of incorporating advice. For larger scale problems and situations where advice is provided more frequently, the agent has to reason about the costs and benefits of compiling advice at a given point in time. Classical issues in trading off deliberation time for action time are relevant here.

Another issue for future study is the problem of operationalization theory incompleteness. For example, our SKB was sufficiently complete for both domains but had it not been, we would have been faced with the problem of supplementing this knowledge base with additional rules. We are considering an approach like that used in the DISCIPLE system [29] to elicit advice and learn rules for operationalizing the advice.

We have chosen the GA method for refinement because it was readily available to us. A comparison of a neural network and other approaches with our current GA approach, on the problems studied here, will provide valuable insights into the tradeoffs between different refinement strategies. We believe that multistrategy learning systems of the future must have a bank of operationalization and refinement methods at their disposal and have fast methods for selecting them. We have chosen a specific breakdown of effort between the advice compilation and refinement phases. How this coordinates with our choice of problem domains and refinement schemes is another question for future study.

Acknowledgements

We appreciate the helpful comments and suggestions provided by Bill Spears, Alan Schultz, Connie Ramsey, John Grefenstette, Alan Meyrowitz. We would especially like to thank John Greffenstette for generously allowing us to use the SAMUEL system upon which our system is built.

References

[1] Basar, T. and G. Olsder, Dynamic Noncooperative Game Theory. New York: Academic Press, 1982.

[2] Chapman, D., Vision, instruction and action. Ph.D. thesis. MIT, 1990.

[3] Clouse, J. and P. Utgoff, A teaching method for reinforcement learning. In *Proc. of the Ninth International Workshop on Machine Learning*, 1992.

[4] Erickson, M. and J. Zytkow, Utilizing experience for improving the tactical manager. In *Proc. of the Fifth International Workshop on Machine Learning*, 1988.

[5] Gordon, D. and D. Subramanian, A multistrategy learning scheme for assimilating advice in embedded agents. In *Proc. of the Second International Workshop on Multistrategy Learning*, 1993.

[6] Gordon, D. and D. Subramanian, Assimilating advice in embedded agents. Unpublished manuscript, 1993.

[7] Grefenstette, J., Ramsey, C., and A. Schultz, Learning sequential decision rules using simulation models and competition. *Machine Learning,* Volume 5, Number 4, 1990.

[8] Holland, J. *Adaptation in Natural and Artificial Systems.* University of Michigan Press: Ann Arbor, 1975.

[9] Kaelbling, L., Goals as parallel program specifications. In *Proc. of the Seventh National Conference on Artificial Intelligence*, 1988.

[10] Karr, C., Design of an adaptive fuzzy logic controller using a genetic algorithm. In *Proc. of the Ninth International Conference on Genetic Algorithms*, 1991.

[11] Laird, J., Hucka, M., Yager, E., and C. Tuck, Correcting and extending domain knowledge using outside guidance. In *Proc. of the Seventh International Conference on Machine Learning*, 1990.

[12] Lin, L., Programming robots using reinforcement learning and teaching. In *Proc. of the Ninth National Conference on Artificial Intelligence*, 1991.

[13] Maes, P. and R. Brooks, Learning to coordinate behaviors. In *Proc. of the Eighth National Conference on Artificial Intelligence*, 1990.

[14] Mahadevan, S. and J. Connell, Automatic programming of behavior-based robots using reinforcement learning. In *Proc. of the Ninth National Conference on Artificial Intelligence*, 1991.

[15] McCarthy, J., Mechanisation of thought processes. In *Proc. Symposium*, Volume 1, 1958.

[16] Michalski, R., Inferential learning theory as a basis for multistrategy task-adaptive learning. In *Proc. of the Eighth International Workshop on Multistrategy Learning*, 1991.

[17] Mitchell, T., Becoming increasingly reactive. In *Proc. of the Eighth National Conference on Artificial Intelligence*, 1990.

[18] Mitchell, T., Mason, M., and A. Christiansen, Toward a learning robot. CMU Technical Report, 1987.

[19] Mostow, D. J., Machine transformation of advice into a heuristic search procedure. In R. Michalski, J. Carbonell, and T. Mitchell (Eds.), *Machine Learning: An Artificial Intelligence Approach* (Vol. 1). Tioga Publishing Co., Palo Alto, CA., 1983.

[20] Nilsson, N., *Principles of Artificial Intelligence.* Tioga: Palo Alto, 1980.

[21] Odetayo, M. and D. McGregor, Genetic algorithm for inducing control rules for a dynamic system. In *Proc. of the Third International Conference on Genetic Algorithms*, 1989.

[22] Ram, A. and J.C. Santamaria, A multistrategy case-based and reinforcement learning approach to self-improving reactive control systems for autonomous robotic navigation. In *Proc. of the Second International Workshop on MultiStrategy Learning*, 1993.

[23] Schoppers, M., Universal plans for reactive robots in unpredictable environments. In *Proc. of the Sixth National Conference on Artificial Intelligence*, 1987.

[24] Schultz, A. and J. Grefenstette, Using a genetic algorithm to learn behaviors for autonomous vehicles. In *Proc. of the Navigation and Control Conference*, 1992.

[25] Schultz, A. and J. Grefenstette, Improving tactical plans with genetic algorithms. In *Proc. of the IEEE Conference on Tools for AI*, 1990.

[26] Segre, A., *Machine Learning of Robot Assembly Plans.* Kluwer: Boston, 1988.

[27] Singh, S., Scaling reinforcement learning algorithms by learning variable temporal resolution models. In *Proc. of the Ninth International Workshop on Machine Learning*, 1992.

[28] Subramanian, D. and S. Hunter, Some preliminary studies in agent design in simulated environments. Unpublished manuscript, 1993.

[29] Tecuci, G. Cooperation in knowledge base refinement. In *Proc. of the Ninth International Workshop on Machine Learning*, 1992.

[30] Tesauro, G., Temporal difference learning of backgammon strategy. In *Proc. of the Ninth International Workshop on Machine Learning*, 1992.

[31] Towell, G. and J. Shavlik, Refining symbolic knowledge using neural networks. In *Proc. of the Eighth Workshop on Machine Learning*, 1991.

[32] Whitehead, S. and D. Ballard, Learning to perceive and act by trial and error. *Machine Learning,* Volume 7, Number 1, 1991.