

Learning and Sequential Decision Making¹

A. G. Barto

Department of Computer and Information Science
University of Massachusetts, Amherst MA 01003

R. S. Sutton

GTE Laboratories Incorporated
Waltham, MA 02254

C. J. C. H. Watkins

Philips Research Laboratories
Cross Oak Lane, Redhill Surrey, RH1 5HA, England
COINS Technical Report 89-95

¹The authors acknowledge their indebtedness to C. W. Anderson, who has contributed greatly to the development of the ideas presented here. We also thank S. Bradtke, J. E. Desmond, J. Franklin, J. C. Houk, A. I. Houston, and E. J. Kehoe for their helpful comments on earlier drafts of this report, and we especially thank J. W. Moore for his extremely detailed and helpful criticism. A. G. Barto acknowledges the support of the Air Force Office of Scientific Research, Bolling AFB, through grant AFOSR-87-0030, and the King's College Research Centre, King's College Cambridge, England, where much of this report was written. A version of this report will appear as a chapter in the forthcoming book *Learning and Computational Neuroscience*, M. Gabriel and J. W. Moore, editors, The MIT Press, Cambridge, MA.

Learning and Sequential Decision Making[†]

A. G. Barto

Department of Computer and Information Science
University of Massachusetts, Amherst MA 01003

R. S. Sutton

GTE Laboratories Incorporated
Waltham, MA 02254

C. J. C. H. Watkins

Philips Research Laboratories
Cross Oak Lane, Redhill Surrey RH1 5HA, England

COINS Technical Report 89-95
September 1989

Abstract—In this report we show how the class of adaptive prediction methods that Sutton called “temporal difference,” or TD, methods are related to the theory of sequential decision making. TD methods have been used as “adaptive critics” in connectionist learning systems, and have been proposed as models of animal learning in classical conditioning experiments. Here we relate TD methods to decision tasks formulated in terms of a stochastic dynamical system whose behavior unfolds over time under the influence of a decision maker’s actions. Strategies are sought for selecting actions so as to maximize a measure of long-term payoff gain. Mathematically, tasks such as this can be formulated as Markovian decision problems, and numerous methods have been proposed for learning how to solve such problems. We show how a TD method can be understood as a novel synthesis of concepts from the theory of stochastic dynamic programming, which comprises the standard method for solving such tasks when a model of the dynamical system is available, and the theory of parameter estimation, which provides the appropriate context for studying learning rules in the form of equations for updating associative strengths in behavioral models, or connection weights in connectionist networks. Because this report is oriented primarily toward the non-engineer interested in animal learning, it presents tutorials on stochastic sequential decision tasks, stochastic dynamic programming, and parameter estimation.

[†]The authors acknowledge their indebtedness to C. W. Anderson, who has contributed greatly to the development of the ideas presented here. We also thank S. Bradtke, J. E. Desmond, J. Franklin, J. C. Houk, A. I. Houston, and E. J. Kehoe for their helpful comments on earlier drafts of this report, and we especially thank J. W. Moore for his extremely detailed and helpful criticism. A. G. Barto acknowledges the support of the Air Force Office of Scientific Research, Bolling AFB, through grant AFOSR-87-0030, and the King’s College Research Centre, King’s College Cambridge, England, where much of this report was written. A version of this report will appear as a chapter in the forthcoming book *Learning and Computational Neuroscience*, M. Gabriel and J. W. Moore, editors, The MIT Press, Cambridge, MA.

Contents

1	Introduction	3
2	Sequential Decision Tasks	5
3	Solving Sequential Decision Tasks	11
4	Mathematical Framework	12
4.1	Systems and Policies	12
4.2	Return and Evaluation Function	13
4.3	Optimality	16
5	Stochastic Dynamic Programming	16
5.1	Computing State Evaluations for a Given Policy	17
5.2	Computing an Optimal Policy	18
6	Parameter Estimation	22
6.1	Feature Vectors, Decision Rules, and Models	24
6.2	Parameter Estimation Methods	25
7	Learning and Sequential Decision Tasks	28
7.1	Learning an Evaluation Function	29
7.2	Learning an Optimal Decision Policy	36
8	Conclusion	44

1 Introduction

In addition to being studied by life-scientists, learning is studied by engineers and computer scientists interested in developing useful devices and programs. The study of synthetic learning has produced an extensive collection of methods and mathematical theories pertaining to such tasks as pattern classification, prediction, and the adaptive control of dynamical systems. Within these theories learning is usually formulated as a search conducted in an abstractly defined space, and a large collection of mathematical concepts can be brought to bear on the problems of understanding and designing procedures, or algorithms, for enabling a device or program to improve its performance over time. What is the nature of the correspondence, if any, between the behavior of an animal in a classical conditioning experiment and the mathematical theories and computational procedures developed for synthetic learning? Is this behavior trivial from a computational perspective, or is it complex and subtle? The answers to these questions that we attempt to justify in this report are that the behavior observed in classical conditioning experiments is far from being computationally trivial; its strongest ties are to mathematical theories and computational procedures that are exceedingly useful in practice and surprisingly complex. By relating the behavior of an animal undergoing classical conditioning to perspectives developed for understanding synthetic learning systems, we hope to provide a framework that may lead to increased understanding of animal behavior as well as novel computational procedures for practical tasks.

Our analysis is based on the Temporal Difference, or TD, model of classical conditioning described by Sutton and Barto in refs. [64, 62]. In this report we view the TD model as a computational method that can be useful in solving engineering problems. Sutton [61] has shown how computational methods making use of "temporal differences," including the TD model of conditioning, are useful for adaptive prediction, and additional publications illustrate how TD methods can be used as components of synthetic learning systems (e.g., refs. [3, 9, 60]). Here, we restrict attention to a slightly simplified version of the TD model, which we call the TD *procedure* in this report. We show how the TD procedure is related to theoretical principles which serve both to explain the operation of TD methods and to connect them to existing theories of prediction, control, and learning. Some of the observations we make are further elaborated by Sutton [61] and Watkins [67], and some of the connections to existing theory were previously described by Werbos [68, 69].

We show how a TD method can be understood as the synthesis of concepts from two existing theoretical frameworks: the theory of *stochastic dynamic programming*, which addresses sequential decision tasks in which both short-term and long-term consequences of decisions must be considered, and the theory of *parameter estimation*, which provides the appropriate context for studying learning rules in the form of equations for updating associative strengths in behavioral models, or connection weights in connectionist networks. Although a clear explanation of how the relevant theoretical ideas fit together requires a certain amount of mathematical notation, it is not our purpose to present a mathematically rigorous account of these ideas, and there are many issues that we do not pursue very deeply or do not discuss at all. Our goals are to explain the main ideas clearly and to provide some aid in accessing the wide body of relevant theoretical literature.¹

¹Because of its connection to theories and computational methods that are in widespread use in many disciplines, it is not possible to describe all the literature relevant to the TD procedure. However, we are not aware that methods are currently in use which combine parameter estimation and dynamic programming in the way they are combined

Drawing precise parallels between behavioral models and computational procedures facilitates the exchange of ideas between researchers studying natural learning and those studying synthetic learning. Sutton and Barto [63] pointed out that the Rescorla-Wagner model of classical conditioning [47] is identical (with some minor caveats) to the equation presented by Widrow and Hoff [72] as a procedure for approximating solutions of systems of linear equations. As a behavioral model, this equation provides a remarkably simple account of a range of stimulus context effects in classical conditioning; as a computational procedure, it has proved exceedingly useful in technological applications, where it is called the LMS (Least Mean Squares) algorithm. This and closely related algorithms are widely used in the fields of signal processing and pattern classification, and are currently playing a major role in the emerging field of connectionist modeling (see, for example, Anderson and Rosenfeld [4], Hinton and Anderson [23], McClelland and Rumelhart [37], and Rumelhart and McClelland [37]). The connection between the experimental and computational literatures due to the parallel between the Rescorla-Wagner model and the LMS algorithm has allowed a fruitful exchange between researchers having widely differing ranges of expertise. In this report we hope to expand this basis for exchange by extending the parallels pointed out by Sutton and Barto [63]. Within the framework adopted here, the Rescorla-Wagner model—and hence also the LMS algorithm—appears as a specialization of the TD procedure.

The relationship between the TD procedure and dynamic programming outlined in this report also has the potential for fostering communication between animal learning theorists and

in the TD procedure, although there has been much research on related problems (Ross [51] and Dreyfus and Law [15] provide good expositions of dynamic programming, and Footnotes 5 and 6 provide references to some of the related engineering research on the adaptive control of Markov processes). To the best of our knowledge, the earliest example of a TD procedure is a technique used by Samuel in a checker-playing program in the late nineteen-fifties [54, 55]. Samuel's program used the difference between the evaluation of a board configuration and the evaluation of a likely future board configuration to modify the equation used to evaluate moves. The evaluation equation was adjusted so that its value when applied to the current board configuration came to reflect the utility of configurations that were likely to arise later in the game. Using this method, it was possible to "assign credit" to moves that were instrumental in setting the stage for later moves that directly captured opponent pieces. Minsky [40, 41] discussed the credit assignment problem and methods similar Samuel's in terms of connectionist networks and animal learning. Mendel and McLaren [38] discussed similar methods in the context of control problems, and the learning method of Witten [77], presented in the context of Markov decision problems, is closely related to the method we describe here. Werbos [68] independently suggested a class of methods closely related to the TD procedure and was the first, to the best of our knowledge, to explicitly relate them to dynamic programming. Werbos calls these "heuristic dynamic programming" methods. A similar connection was made recently by Watkins [67], who uses the term "incremental dynamic programming."

In his PhD dissertation [60], Sutton developed an algorithm, calling it the "adaptive heuristic critic" algorithm, that is closely related to Samuel's method but extended, improved, and abstracted away from the domain of game playing. This work began with the interest of Sutton and Barto in classical conditioning and the exploration of Klopf's idea of "generalized reinforcement" [27, 28], which emphasized the importance of sequentiality in a neuronal model of learning. The adaptive heuristic critic algorithm was used (although in slightly different form) in the reinforcement-learning pole balancer of Barto, Sutton, and Anderson [9], where it was incorporated into a neuron-like unit called the "adaptive critic element." This system, which was inspired by the "Boxes" system of Michie and Chambers [39], was further studied by Selfridge, Sutton, and Barto [57] and Anderson [2, 3]. Since then, Sutton [61] has extended the theory and has proved a number of theoretical results. His results suggest that TD procedures can have advantages over other methods for adaptive prediction.

A number of other researchers have independently developed and experimented with methods that use TD principles or closely related ones. The "bucket brigade" algorithm of Holland [24] is closely related to the TD procedure as discussed by Sutton [61] and Liepins, Hilliard, and Palmer [32]. Booker's [11] learning system employs a TD procedure, as does Hampson's [22] proposed learning system, which is very similar to the one we discuss here. Other related procedures have been proposed as models of classical conditioning in publications cited in refs. [64, 62].

behavioral ecologists. Dynamic programming has been used extensively in behavioral ecology for the analysis of animal behavior (see, for example, Krebs, Kacelnik, and Taylor [29], Houston, Clark, McNamara, and Mangel [25], and Mangel and Clark [36]). In these studies dynamic programming is used to determine decision strategies meeting certain definitions of optimality to which animal behavior is compared. Behavioral ecologists do not suggest that the animals themselves perform dynamic programming—indeed, most of the forms of behavior studied are regarded as being innate, and dynamic programming would appear to provide a poor model of learning. In a sense that will be made clear in what follows, dynamic programming methods work backwards from the end of a decision task to its beginning, calculating information pertinent to decision making at each stage based on information previously calculated from that stage to the task's end. As a result of this back-to-front processing, it is difficult to see how dynamic programming can be related to learning processes that operate in real-time as a system interacts with its environment. However, we show how the TD procedure can accomplish much the same result as a dynamic programming method by repeated forward passes through a decision task, that is, through repeated trials, instead of explicit back-to-front computation. This relationship of the TD procedure to dynamic programming suggests a range of research questions involving links between animal behavior in carefully controlled learning experiments and the less restricted forms of behavior studied by behavioral ecologists.

This report is organized into three major parts. The first part, consisting of Sections 2 and 5, presents the basic ideas of stochastic dynamic programming. Section 2 describes, in both informal and mathematical terms, a general class of tasks known as *stochastic sequential decision tasks* to which the methods of stochastic dynamic programming apply, and Section 5 describes some of these methods. This material is tutorial in nature and is based on the formulation of Ross [51]. The second major part of the report, consisting of Section 6, is a tutorial on parameter estimation based on the view taken in the field of adaptive control as described by Goodwin and Sin [20]. Some of this material also appears in Barto [5]. The report's third major part, consisting of Section 7, shows how the TD procedure emerges as a synthesis of the ideas from dynamic programming and parameter estimation covered in the first two parts of the report. Here we also show how the TD procedure can be used in conjunction with another procedure and applied to stochastic sequential decision tasks to produce an analog of instrumental learning. The combination of these two procedures corresponds to the use of the "adaptive critic element" and the "associative search element" in the pole balancer of Barto, Sutton, and Anderson [9]. The framework of stochastic sequential decision theory helps explain the interaction of these two procedures and suggests other learning methods for this and related tasks. We conclude with a discussion of what the theoretical basis of the TD procedure suggests about animal learning and some of the directions that can be taken in extending this approach.

2 Sequential Decision Tasks

Animals face many situations in which they have to make sequences of actions to bring about circumstances favorable for their survival. We are interested in tasks in which the consequences of an action can emerge at a multitude of times after the action is taken, and we shall be concerned with strategies for selecting actions based on both their short- and long-term consequences. Tasks of this kind can be formulated in terms of a dynamical system whose behavior unfolds over time under the influence of a decision maker's actions. Modeling the behavior of such a system is greatly simplified by the concept of *state*. The state of a system at a particular time is a description of

the condition of the system at that time that it is sufficient to determine all aspects of the future behavior of the system when combined with knowledge of the system's future input. Whatever happened to the system in the past that is relevant to its future behavior is summed up in its current state—future behavior does not depend on how the system arrived at its current state, a property sometimes called “path-independence” of the system description. The concept of state is also useful in describing systems which operate according to probabilistic rules. In this case, the system state together with future input determine the probabilities of all aspects of the future behavior of the system independently of how the state was reached. This is the *Markov property* of a stochastic dynamical system.

Consider a decision making agent, which we simply call an *agent*, facing the following task. The agent interacts with a system in such a way that at the beginning of each of a series of discrete time periods, it observes the system and is able to determine the system state at that time. Based on the observed state, the agent performs an action, thereby causing the system to deliver to the agent a ‘payoff’, which we think of as a number whose value depends on the system state, the agent's action, and possibly random disturbances. The system then makes a transition to a new state determined by its current state, the agent's action, and possibly random disturbances. Upon observing the new state, the agent performs another action, receives another payoff, and the system changes state again. This cycle of state-observation, action, payoff, and state-change repeats for a sequence of time periods. The agent's task, described mathematically in Section 4, is to select the actions that maximize the total, or cumulative, amount of payoff it receives over time. This is more difficult than merely trying to maximize each individual payoff. Some actions may be useful in producing a high immediate payoff, but these same actions may cause the system to enter states from which *later* high payoffs are unlikely or impossible. Hence performing these actions would result in less total amount of payoff than might be possible otherwise. Conversely, some actions may produce low payoff in the short-term but are necessary to set the stage for greater payoff in the future. The agent's decision making method must somehow account for both the short- and the long-term consequences of actions.

The total amount of payoff received by the agent over many time periods depends on the number of time periods over which this total is determined, the sequences of actions and states that occur over these time periods, and the outcomes of whatever random factors influence the payoffs and state transitions. The number of time periods over which the total amount of payoff is determined is called the *horizon* of the decision task. If the horizon is finite, the total amount of payoff is simply the sum of the individual payoffs received at each time period until the task's horizon is reached. In the horizon is infinite, however, this sum may not be finite, a difficulty remedied by introducing a *discount factor* that allows payoffs to be weighted according to when they occur. In this case, what we mean by the total amount of payoff over an infinite number of time periods is a weighted sum of the infinite number of individual payoffs received, where the weights decrease with increasing temporal remoteness of the payoffs (we define this precisely in Section 4). If the discount factor is chosen appropriately, then this weighted sum will always have a finite value despite its dependence on an infinite number of payoffs. Sutton and Barto [62] refer to this as *imminence weighting*. In this report, we restrict attention to infinite horizon tasks where a discount factor is used in determining the relevant measure of the total amount of payoff.

Describing a decision task in terms of system states permits a relatively simple statement of how action and state sequences determine the total amount of payoff an agent receives. Suppose the agent uses a rule to select actions depending on system state. This rule, called the agent's decision policy, or simply its *policy*, associates an action with each system state. The agent's

action upon observing a state is the action associated with that state by the policy. If no random factors are involved in the task, then the sequences of actions and states depend only on the agent's policy and the system state at the beginning of the task, i.e., the task's initial state. Consequently, the total amount of payoff received until the task's horizon is reached (where the total amount is determined by discounting if the task has an infinite horizon) also depends only on the agent's policy and the task's initial state. By a policy's *return* for a given system state we mean the total amount of payoff the agent receives until the task's horizon is reached, assuming the task's initial state is the given state and the agent uses the given policy. For infinite horizon tasks where a discount factor is used, a policy's return for a state is the weighted sum (where the weights depend on the discount factor) of the payoffs the agent would receive over an infinite number of time periods for the given initial state if the agent were to use the given policy to select an infinite sequence of actions.² Thus, when no random factors are involved in a sequential decision task, the payoff for a system state depends on a single action of the agent, but the return for a state depends on the consequences of the agent's decisions as specified by its policy for the duration of the task.

When a decision task involves random factors, a policy's return for each system state is a random variable. In this case, one can define the *expected return* for each policy and system state. For the infinite horizon case with discounting, which is our concern here, the expected return for a policy and system state is the mathematical expectation, or mean, of the random variable giving the return for that policy and state. The expected return depends on the distribution functions of all the random factors influencing the task and can be thought of as the average of an infinite number of instances of the decision task, where the agent uses the same policy and the system starts in the same initial state in each instance. As formulated here, the objective of a sequential decision task is defined in terms of expected return: The objective is to find a policy that maximizes the expected return for all possible initial system states. Such a policy is called an optimal policy.

Although we discuss tasks requiring maximizing expected return, this class includes as special cases tasks in which the objective is to obtain any payoff at all. For example, suppose that for all but one system state the payoffs received by the agent are zero no matter what action the agent selects. Also suppose that the task ends when a nonzero payoff is obtained. One can think of the state from which nonzero payoff is available as the *goal state*. In this case, a policy's return for each initial state is zero unless its use by the agent brings about the goal state. Hence, in this case, selecting actions to maximize return is the same as selecting actions that cause the system to enter the goal state. If a discount factor is used, it turns out that the return is maximized by selecting actions that bring about the goal state in the fewest time periods. Tasks such as this, in which the objective is to reach a designated goal, are included in the theory we describe, and the example used throughout this report is an instance of this type of task.

There are numerous examples of sequential decision tasks, many of which have great practical significance. The task of finding the least-cost route from one place to another is perhaps the most generic example. Choice points along a route correspond to the states of the system, actions determine what place is reached next, and the magnitude of the payoff received in response to an action is inversely related to the cost of the path travelled (so that by maximizing the total amount of payoff, one minimizes total path cost). More complex tasks involving resource

²Othe formulations of sequential decision tasks result from different definitions of a policy's return. One formulation that is commonly studied defines a policy's return as the average amount of payoff per-time-step over a task's duration. Ross [51] also discusses this formulation.

allocation, investment, gambling, and foraging for food are also examples of sequential decision tasks. Most of the planning and problem-solving tasks studied by artificial intelligence researchers are sequential decision tasks. Other examples are studied by control engineers, such as the problem of placing a spacecraft into a desired orbit using the least amount of fuel. In some sequential decision tasks, the distinction between the agent and the system underlying the decision task may not be as clear-cut as our discussion would lead one to believe. For example, in a foraging model in behavioral ecology, the state of the system may be the forager's energy reserves [36], a quantity apparently describing an aspect of the agent itself instead an external system. It can be misleading to identify an agent with an entire organism.

We use a simple route-finding task to illustrate the concepts and methods described in this report. We emphasize that this is merely an example and that these concepts and methods are applicable to tasks that are much more complex than this route-finding task.

An Example—Figure 1 shows a grid representing a region of space. Each intersection of the grid lines is a 'location', and the region contains a C-shaped barrier and a goal location. For the 8-by-12 grid shown, there are 96 locations. Two locations are adjacent if they are connected by a grid line that does not pass through any other locations. A path is a set of line segments tracing a route through the region, where each segment connects two adjacent locations in the region. The length of a path is the number of distinct line segments it contains. The task we consider is to find, for each location, a path to the goal that begins at the given location, does not cross the barrier, and is has the smallest possible length. Each such shortest path is an optimal path.

This task can be formulated as a sequential decision task by considering an agent which can move from its current location to an adjacent location in each time period. The spatial environment defines the system underlying the decision task. For each time period the state of the system is the current location of the agent, and the state at the next time period—the new location of the agent—is determined by the the current location of the agent and the action chosen. We let the agent choose any one of the four actions North (N), South (S), East (E), and West (W) at each time period. The effect of an action depends on the current system state, i.e., the current location of the agent. For most locations, the action causes the agent to move to the location adjacent to its current location in the direction indicated by the action. However, for locations from which a move is blocked by a barrier or would take the agent out of bounds, the effect of any 'disallowed' action is to leave the agent's location unchanged. If the agent is located at the goal, it stays there no matter what action it performs. Thus, the set of actions available to the agent is always the same, but actions can have differing consequences depending on the agent's location.³

A policy in this example is a rule that assigns an action to each location. One could think of a policy as one of the many possible patterns of placing a sign post (indicating N, S, E, or W) at each location which the agent is compelled to follow. The objective of the task is to form a policy, i.e., to place a pattern of 96 sign posts, that directs the agent from each location to the goal in the fewest possible time periods, i.e., along an optimal path. To formulate this as the problem of finding a policy that maximizes expected return, we effectively punish the agent every time period in which it is not at the goal. The agent always receives a payoff of -1 unless the agent is located at the goal, in which case it receives a payoff of zero for any action. Therefore, the sum

³Note that an action in a sequential decision task is not the same as a component of the agent's observable behavior. Observable behavior is a joint consequence of the agent's action and the state of the system underlying the decision task.




Figure 1: Plan view of a spatial environment for the route-finding example. The intersections of the lines are possible locations for the agent. The C-shape is a barrier, and the goal location is indicated.

of payoffs over a path from a starting location to the goal, i.e., the return produced over the path, is the negative of the number of time periods taken to reach the goal (assuming no discounting). Selecting actions to maximize return therefore minimizes the number of time periods taken to reach the goal. An optimal policy directs the agent along an optimal path from each location.

In what follows, we discuss several versions of this task that differ in terms of the amount of knowledge we assume for the agent. Although all of these tasks are relatively easy instances of the tasks encompassed by the theory, some of the sources of additional complexity can be appreciated clearly using the route-finding example. For example, the payoff received for a move need not always be -1 but can instead reflect the distance of a path, the degree of difficulty encountered in traversing it, etc. Additionally, the payoffs and the state transitions need not be deterministic. Actions may only influence the probabilities that specific places are reached. Finally, additional complexity, which we do not address at all in what follows, occurs if the agent does not have access to complete state information, a situation that would appear in the route-finding example when the agent is unable to distinguish all possible locations.

Before discussing solution methods for sequential decision tasks, we make several observations about the theoretical framework implied by this class of tasks. All of these observations are reminders that it can be misleading to take the abstractions involved in this framework too literally.

- *Discrete time*—What do the discrete time periods of a decision task mean in terms of real time? In the kinds of discrete-time models to which we restrict attention, these time periods are called *time steps* and are merely *computational instants* that may correspond to instants of real time separated by some interval, or to separate collections of events, such as trials. In modeling conditioning behavior, for example, a time step may represent some small interval of real time, as in the TD model presented by Sutton and Barto [64, 62]. Alternatively,

each time step might be an entire conditioning trial, as in a trial-level model such as the Rescorla-Wagner model [47]. In the discussion to follow, a time step merely refers to a time period of an abstract sequential decision task.

- *Receiving payoff*—The framework described above is sometimes interpreted to mean that payoff is delivered to the agent as if there were a single sensory pathway dedicated to this function. If one identifies the agent with an entire animal—which can be misleading as we emphasized above—this view suggests that an animal has a single sensory input dedicated to the function of receiving all primary reinforcement, which is obviously not the case. It is better to think of the payoff at each time period as a concise way of summarizing the affective significance of the immediate consequences of performing an action. Because the immediate payoff and the system's next state are both functions of the current action and system state, it is a special case to regard each payoff simply as a function of the system's current state. One can think of this function as being computed by the agent itself and not by the system underlying the decision task.
- *Complete state information*—The assumption is made that at each time step a complete description of the state of the system underlying the task is available to the agent. Clearly, this is a strong assumption. Although the consequences of weakening this assumption are theoretically significant and relevant to the study of natural learning, they are complex and beyond the scope of the present report. However, it is important to keep in mind the following two observations. First, one need not think of the current system state as something that must be read directly from the agent's sense organs. More generally, the state of the system can be provided through the aid of complex world models and memories of past sensations and behavior. The theory requires the availability of state information, but it is not sensitive to how this information is obtained. Second, knowing the current state of the system underlying a decision task is not the same as knowing beforehand how the system will behave in response to actions; that is, it is not the same as having an accurate model of the decision task.
- *Optimality*—The objective of a decision task considered here is to find an optimal decision policy—a policy that maximizes the expectation of the discounted sum of future payoffs. Some theories of animal behavior invoking optimality principles are “molar” optimality theories which propose that behavior can be understood in terms of optimization, but they do not specify mechanistic computational procedures by which the conditions of optimality can be achieved. “Molecular” optimization theories, on the other hand, specify procedures that operate from moment-to-moment but which only approximate optimal solutions in most tasks (see, for example, refs. [45, 59]). Although here it is not our goal to provide a comprehensive theory of animal behavior, we can describe the kind of theory with which the perspective taken in this report is consistent. The framework of sequential decision theory adopted here provides a molar view of behavior involving a specific optimality criterion, and the TD and policy adjustment procedures we describe provide molecular accounts of how optimal behavior might be approximated. At the molar level, because this view defines optimality in terms of a dynamical system underlying a decision task, it avoids the circularity inherent in theories that define optimality only as revealed by an organism's behavior. At the molecular level, the computational procedures outlined provide one account of how the maximization of cumulative payoff over time can be approximated without recourse to teleological principles. Moreover, because these procedures usually only approximate

optimal policies, they do not imply that optimal behavior will be achieved. The performance of these procedures in specific decision tasks depends on the specification of many details, such as the manner of representing system states. Choices made in specifying these details strongly influence the degree of optimality achievable.

3 Solving Sequential Decision Tasks

Because so many problems of practical interest can be formulated as sequential decision tasks, there is an extensive literature devoted to the study of solution methods for this type of task, the large majority of which require the agent to have a complete model of the decision task. Even if one has a complete model of the decision task, which means knowing all the state-transition and payoff probabilities of the system underlying the task, determining the best policy can require an extremely large amount of computation because it effectively requires a search through the set of all possible state sequences generated by all possible sequences of actions. Except for very specialized tasks in which analytical methods can be used instead of search, the required amount of computation increases so rapidly with increases in a task's size (as determined by its horizon, number of states, and number of actions) that it is not feasible to perform this search for large tasks. Dynamic programming, a term introduced in 1957 by Bellman [10], consists of particular methods for organizing the search under the assumption that a complete model of the decision task is available. Although these methods are much more efficient than explicit exhaustive search of all possible state sequences, the amount of computation still grows so rapidly with the size of the task that large tasks remain intractable. Search methods specialized to take advantage of specific kinds of "heuristic" knowledge can be applied to some types of larger tasks, but these methods also require a complete model of the decision task and can still require prohibitive amounts of computation.⁴

Methods for estimating optimal policies in the absence of a complete model of the decision task are known as adaptive or learning methods. Because the most difficult aspect of applying dynamic programming is often the accurate modeling of the decision task, adaptive methods have great practical importance. Additionally, if an adaptive method can improve a decision policy sufficiently rapidly, then the amount of computation required may be less than would be required by an explicit solution via dynamic programming. How can an optimal policy be constructed when a complete model of the decision task is not available? Instead of being able to generate a solution by manipulating a task model, it is necessary to learn about the system underlying the task while interacting with it. Two general approaches are possible, one of which has been much more widely studied than the other. Most widely studied is the *model-based* approach, which requires constructing a model of the decision task in the form of estimates of state-transition and payoff probabilities. These probabilities can be estimated by keeping track of the frequencies with which the various state transitions and payoffs occur while interacting with the system underlying the decision task. Assuming that these estimates constitute an accurate model of the decision task, one can then apply a computational technique for finding an optimal policy, such as a dynamic programming technique, which requires an accurate model of the decision task.⁵

⁴A large component of artificial intelligence research concerns search strategies of this type, called "heuristic search" strategies, although their objective is usually not to maximize a measure of cumulative payoff. See Pearl [44].

⁵Most of the methods for the adaptive control of Markov processes described in the engineering literature are model-based. Examples are provided by Borkar and Varaiya [12], El-Fattah [17], Kumar and Lin [30], Lyubchik

In this report our concern is with other approaches to learning how to solve sequential decision tasks, which we call *direct* approaches. Instead of learning a model of the decision task, that is, instead of estimating state-transition and payoff probabilities, a direct method adjusts the policy as a result of its observed consequences. Actions cannot be evaluated unless they are actually performed. The agent has to try out a variety of decisions, observe their consequences, and adjust its policy in order to improve performance. We call this process *reinforcement learning* after Mandel and McLaren [38] who describe its relevance to adaptive control.⁶ To facilitate the direct learning of a policy, it is possible to adaptively improve the criteria for evaluating actions so that the long-term consequences of actions become reflected in evaluations that are available immediately after an action is performed. The TD procedure appears as a method for accomplishing this.

In terms of theories of animal learning, the payoffs delivered to the agent at each stage of the decision task correspond to primary reinforcement, and the TD procedure provides a gradient of secondary reinforcement by anticipating the events that provide primary reinforcement. Secondary reinforcement is therefore defined in terms of estimates for the expected *sum of future primary reinforcement* (possibly discounted). Adjusting a decision policy on the basis of this acquired reinforcement is related to learning in instrumental tasks—actions that increase this estimated sum more than others are favored in the process of learning what actions to perform. Although this corresponds to a stimulus-response (S-R) view of instrumental learning based on a version of the Law of Effect [65], it is not our aim to argue for the validity of this view of animal instrumental learning, which probably involves more than can be accounted for by an S-R model (see, for example, Dickinson [14] and Rescorla [46]). In Section 7 we provide an example of how the TD procedure can be used with this kind of reinforcement learning method, but the TD procedure can also be combined with model-based methods in a variety of ways, as was done, for example, by Samuel [54], whose system performed look-ahead searches. However, it is beyond the scope of the present report to discuss these more complex learning systems. Watkins [67] discusses some of the issues that arise in combining direct and model-based learning methods.

4 Mathematical Framework

4.1 Systems and Policies

We assume that the system underlying the decision task is a discrete-time dynamical system with a finite set of states, X . At any time step $t = 0, 1, 2, \dots$, the system can be in a state $x \in X$. After observing the system state at time step t , the agent selects (and performs) an action from a finite set, A , of possible actions. Suppose that at time step t , the agent observes state x and selects action a . Then, independently of its past history, the system makes a transition from state x to state y with a probability that depends on the action a . We denote this probability $P_{xy}(a)$. When this state transition occurs, the agent receives a payoff, denoted r , which is determined randomly

and Poznyak [34], Mandl [35], Riordon [48], and Sato, Abe, and Takeda [56]. Most of these methods apply to the case in which return is the average payoff per-time-step and the underlying system is an ergodic Markov chain for each possible policy. They differ in how the policy is adjusted over time based on the current estimates for the transition and payoff probabilities.

⁶Examples of various direct methods for learning how to solve sequential decision tasks are those of Lyubchik and Poznyak [34], Lakshmivarahan [31], Wheeler and Narendra [71], and Witten [77]. Most of these rely on results about the collective behavior of stochastic learning automata and ergodic Markov chains (see also Narendra and Thathachar [43]).

in a manner depending on x and a . This sequence of events repeats for an indefinite number of time steps. Because we shall be concerned only with the *expectation* of the total amount of payoff accumulated over time, it is not necessary to specify the details of the probabilistic process by which a payoff depends on states and actions. It suffices to specify only how the expected value of a payoff depends on actions and states. We let $R(x, a)$ denote the expected value of a payoff received as a consequence of performing action a when observing state x . We assume that the payoff whose expected value depends on the system state and agent action at time step t is by the agent at the next time step: step $t + 1$. Although this formalism can be modified to let the set of actions available to the agent depend on the system state, for simplicity we have chosen to let the set of actions, A , be the same for each state.

The objective of the decision task is to find a policy for selecting actions that is optimal in some well-defined sense. In general, the action specified by the agent's policy at a time step can depend on the entire past history of the system. Here we restrict attention to policies, called a *stationary* policies, that specify actions based only on the current state of the system. A stationary policy can be represented by a mapping, denoted π , that assigns an action to each state. The action specified by policy π when the system is in state x is denoted $\pi(x)$. For the route-finding example, a policy specifies what direction to move from any location. In Section 7.2, we consider policies that specify actions probabilistically based on the current system state.

Letting x_t denote the system state at time step t , if the agent is using policy π , then the action it takes at step t is $a_t = \pi(x_t)$. The system state changes according to

$$\text{Prob}\{x_{t+1} = y | x_0, a_0, x_1, a_1, \dots, x_t = x, a_t = a\} = \text{Prob}\{x_{t+1} = y | x_t = x, a_t = a\} = P_{xy}(a).$$

Letting r_{t+1} denote the payoff received by the agent at time step $t + 1$, we have

$$E[r_{t+1} | x_t, a_t] = R(x_t, a_t), \tag{1}$$

where E indicates expected value. Figure 2 illustrates this formulation by showing when the relevant quantities are available to the agent, how they change over time, and on what information they depend. Because we assume the policy π is stationary, the sequence of states forms a Markov chain with transition probabilities $P_{xy} = P_{xy}(\pi(x))$. For this reason, decision tasks of this kind are called *Markov decision tasks*.

For the route-finding example, the state set, X , consists of the 96 different locations, one of which is the goal; the action set, A , is $\{N, S, E, W\}$; $R(x, a) = -1$ for all states x and actions a , except for the goal state, g , for which $R(g, a) = 0$ for all actions a . Because the system in this example is deterministic, $R(x, a)$ is the actual payoff received by the agent for performing action a in location x , and the transition probabilities equal either one or zero: $P_{xy}(a) = 1$ if action a moves the agent from location x to location y , and is zero otherwise.

4.2 Return and Evaluation Function

It is now possible to define a policy's return, which for our purposes will be the infinite horizon discounted return. Suppose an agent begins interacting with a system at time step $t = 0$ and is able to continue for an infinite number of time steps. Using discount factor γ , the measure of the total amount of payoff that the agent will receive over this infinite time period is

$$r_1 + \gamma r_2 + \gamma^2 r_3 + \dots + \gamma^n r_{n+1} + \dots \tag{2}$$

When $0 < \gamma < 1$, the powers of γ weighting the payoffs form a decreasing sequence so that later payoffs are weighted less than earlier ones, with payoffs in the far distant future contributing

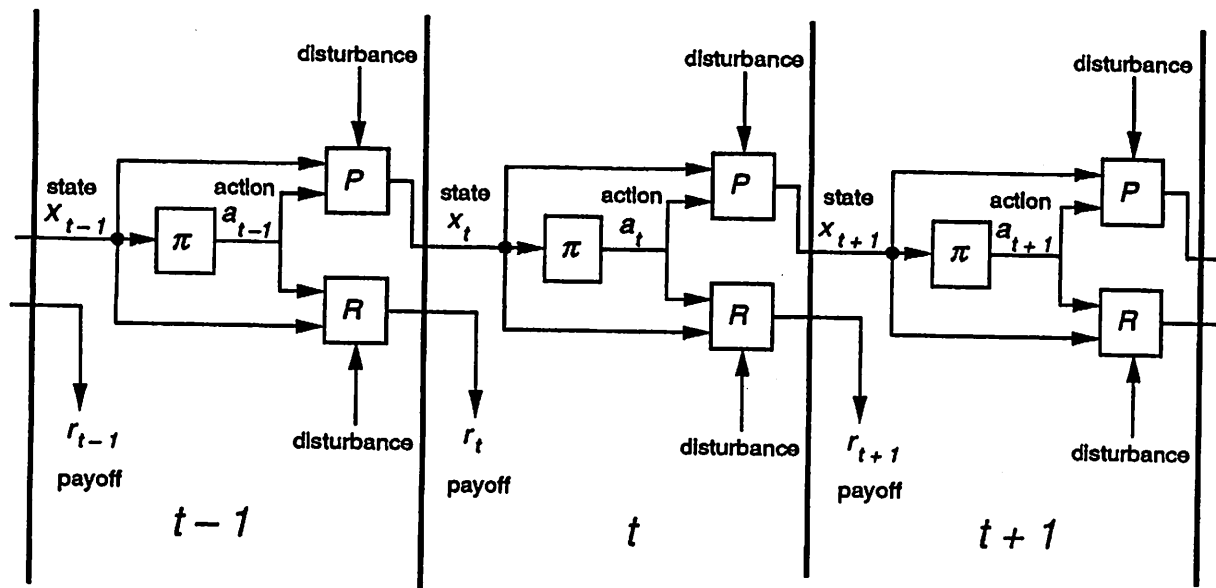


Figure 2: How the state, action, and payoffs change over time in a sequential decision task. The computation of the action, state, and payoff are shown for three successive time steps. The squares labeled π , P , and R respectively represent the decision policy, state-transition computation, and payoff computation. The stochastic aspects of the system are illustrated as disturbances influencing the latter two computations. In "space-time" diagrams of this type, several of which appear below, the quantities between two bold vertical lines are labeled with the same time subscript (exceptions to this which occur below are explained in the text). The repetition of a rectangle representing a system component indicates the same component viewed at different times.

negligibly to the return. If $\gamma = 0$, the sum in Expression 2 is simply the immediate payoff, r_1 , due to the first action, a_0 ; if $\gamma = 1$, it is the sum of all the payoffs received and generally is not finite.⁷ The discount factor adjusts the degree to which the long-term consequences of actions must be accounted for in a decision task and influences the rate at which learning occurs. Although we do not discuss the problem of determining a value of the discount factor appropriate for a particular application, it is clear that a value close to one is appropriate when there is a high degree of certainty about how the system will evolve over time. One would only want to sacrifice short-term return for long-term return that is highly certain. We also do not discuss procedures for dynamically adjusting γ with experience, even though this would be useful in many applications.

In addition to depending on the system underlying the decision task, the initial system state, and the decision policy used by the agent, the sum given by Expression 2 depends on unknown or random factors influencing the state transitions and the payoffs. The expected value of this sum over all possible decision tasks starting with initial state x when the agent uses policy π is:

$$E_{\pi} \left[\sum_{t=0}^{\infty} \gamma^t r_{t+1} | x_0 = x \right], \quad (3)$$

where E_{π} indicates that the expected value is taken under the assumption that policy π is used to select actions. We can think of this quantity as the value of a function, denoted V^{π} , assigning an expected return to each system state x :

$$V^{\pi}(x) = E_{\pi} \left[\sum_{t=0}^{\infty} \gamma^t r_{t+1} | x_0 = x \right]. \quad (4)$$

This function is the *evaluation function* for policy π . For each state x , $V^{\pi}(x)$ is the expected return over the infinite number of time steps beginning at $t = 0$ under the conditions that the system begins in state x and the agent uses policy π , where it is understood that the discount factor, γ , has some specified value. We call $V^{\pi}(x)$ the *evaluation* of state x ,

Because a system state has the property that the future behavior of the system does not depend on how or when the system arrives at a state, the evaluation of a state is in fact the expected return over an infinite number of time steps beginning *whenever* the system enters that state under the condition that the agent continues to use policy π thereafter. The evaluation of a state is a *prediction* of the return that will accrue throughout the future whenever this state is encountered. If one can determine the evaluation of a state merely from observing the system when it is in that state, then this prediction is effectively available *immediately* when the system enters that state, even though the prediction contains information about the entire future.

For the route-finding example, the evaluation of location x for policy π depends on the expected number of time steps the agent takes to reach the goal from location x using policy π . If π always brings the agent to the goal in, say, n time steps from a location x , then $V^{\pi}(x) = -1 - \gamma - \gamma^2 - \dots - \gamma^{n-1}$. The series has n terms because the payoff is zero for all time steps after the goal is reached. When $\gamma = 1$, this series sums to $-n$, and as γ decreases, the sum becomes less negative, approaching -1 as γ approaches zero. The evaluation of the goal, $V^{\pi}(g)$, is zero for any policy. In this task, because a location's evaluation is directly related, depending on γ , to the negative of the number of time steps to the goal, the larger a location's evaluation

⁷One situation in which it is finite when $\gamma = 1$ is when the structure of the task ensures that all payoffs are zero after a certain point in the decision task. However, restricting γ to be greater than zero and strictly less than one ensures a finite return under all circumstances (assuming the payoffs themselves are finite).

for policy π , the fewer time steps are required to move from it to the goal using π . If π does not generate a path from x to the goal in any finite number of steps, which can happen if π produces a looped path, then $V^\pi(x) = -\sum_{t=0}^{\infty} \gamma^t$, which converges to $-1/(1-\gamma) \leq -1$ if $0 \leq \gamma < 1$.

4.3 Optimality

Let π and π' be any two policies. Policy π' is an improvement over policy π if the expected return for π' is no smaller than that for π for any state, and is larger for at least one state. More precisely, π' is an improvement over π if $V^{\pi'}(x) \geq V^\pi(x)$ for all states x , with strict inequality holding for at least one state. A policy is an *optimal policy*, which we denote π^* , if no policy is an improvement over it. Because the optimality of policies depends on the discount factor, technically we should refer to γ -optimal policies. As γ changes, different policies become optimal because a policy best for short-term versions of a task will generally not be best for long-term versions. However, for simplicity, we omit reference to γ if no confusion is likely to result, and we assume that whenever policies are compared, they are compared according to expected returns defined for the same γ . In the route-finding example, recalling that a location's evaluation is directly related to the negative of the number of time steps to the goal, an optimal policy takes the agent from any location to the goal in the fewest possible time steps.

There can be many optimal policies for a given decision task. For example, in the route-finding illustration, there are several different ways to move from most locations to the goal in the minimum number of time steps. However, the evaluation functions for any two optimal policies must be identical. In the route-finding example, this corresponds to the simple observation that all the minimum-time paths from any location to the goal must be traversable in the same number of time steps. We can therefore define the unique *optimal evaluation function*, which we denote V^* , for a given decision task as the evaluation function for any optimal policy. For any optimal policy π^* , $V^{\pi^*}(x) = V^*(x)$ for all states x . An agent using an optimal policy will maximize the expected return from any system state. The object of a sequential decision task is to find one of the optimal policies.

5 Stochastic Dynamic Programming

The definition of the evaluation function for a given policy given by Equation 4 does not immediately suggest how the values of this function can be computed. For policy π , the evaluation of a state depends on *all possible* state sequences that can occur when policy π is used, and the probabilities that these sequences occur. Even if one knows all the transition probabilities and the expected values of all payoffs, it is not immediately obvious how one can compute the evaluations of states. Similarly, it is not obvious how to determine the optimal evaluation function or an optimal policy. Stochastic dynamic programming provides methods for computing the evaluations of states, or for approximating their evaluations as closely as desired, as well as for approximating the optimal evaluation function and an optimal policy, under the assumption that one has an accurate model of the decision task. Having an accurate model of the decision task means knowing the payoff expectations, $R(x, a)$, and state-transition probabilities, $P_{xy}(a)$, for all states x and y and all actions a .

We first describe a dynamic programming method for finding the evaluations of states for a given policy, π , under the assumption that we have an accurate model of the decision task.⁸

⁸In this case, it is not necessary to know the transition probabilities and payoff expectations for state-action

Then we describe dynamic programming methods for finding the optimal evaluation function and an optimal policy, again assuming an accurate model of the decision task. Consequently, in this section we discuss the *computation* of these various functions given that we have all the relevant knowledge about the decision task. These computational methods provide the background for understanding how these various functions can be *learned* without assuming this kind of knowledge using the methods presented in Sections 7.2 and 7.1.

5.1 Computing State Evaluations for a Given Policy

The evaluation of state x for policy π , $V^\pi(x)$, gives the expected return if policy π is used throughout an infinite horizon decision task beginning in state x . One method for determining a state's evaluation is to approximate it as the evaluation for a finite horizon version of the task with a horizon chosen large enough to produce an approximation of sufficient accuracy. As the horizon of this finite horizon task increases, the approximation approaches the state's true evaluation. This is a simple method of successive approximations, although it can be computationally costly if there is a large number of states.

If the agent interacts with the system by using policy π for some finite number of time steps, n , then we can consider the expected return for this finite horizon task. This n -step expected return for state x , denoted $V_n^\pi(x)$, is the expected value of the sum of the first n terms of Equation 2. We call this the n -step evaluation of state x . The successive approximations method begins with $n = 1$. It is easy to find the one-step evaluation of state x , $V_1^\pi(x)$. This is the expected value of the payoff the agent will receive by using policy π for one step starting in state x . According to the framework described in Section 4, we know that the action the agent will make is determined by applying policy π to the state x . This action is $a = \pi(x)$. We also know from Equation 1 that the expected value of the payoff as a result of performing action a from state x is $R(x, a)$. Then we have

$$V_1^\pi(x) = R(x, a) = R(x, \pi(x)).$$

Applying this equation to all system states completely determines the one-step evaluation function, V_1^π , for policy π .

Suppose this has been accomplished so that V_1^π is known, and now suppose that the agent faces the two-step task starting in state x . After the agent performs action $a = \pi(x)$, it receives payoff with expected value $R(x, a)$, and the next system state is y with probability $P_{xy}(a)$. If we call the actual next state y , then the agent faces the one-step task from state y , and it already knows that the expected return for using policy π for one step beginning in state y is the one-step evaluation of y , $V_1^\pi(y)$. Thus if the next system state is y , the two-step evaluation of x , i.e., the two-step expected return for state x , would be the expected value of the immediate payoff, $R(x, a)$, plus the one-step evaluation, discounted by γ , of y , which is $\gamma V_1^\pi(y)$. However, in general, the transition to a next state, y , occurs with probability $P_{xy}(a)$. Thus, instead of using the one-step evaluation of the actual next state, one must use the expectation over all possible next states of the one-step evaluations of these states, which is $\sum_y P_{xy}(a) V_1^\pi(y)$. We therefore have

$$V_2^\pi(x) = R(x, a) + \gamma \sum_{y \in X} P_{xy}(a) V_1^\pi(y), \quad (5)$$

for all states x , where $a = \pi(x)$. In a similar manner we can determine V_3^π from V_2^π , V_4^π from V_3^π , and so on. The general rule for determining V_n^π in terms of V_{n-1}^π is a basic equation of stochastic

pairs that never occur with the given policy.

dynamic programming:

$$V_n^\pi(x) = R(x, \pi(x)) + \gamma \sum_{y \in X} P_{xy}(\pi(x)) V_{n-1}^\pi(y), \quad (6)$$

for all system states x .

This iterative process for computing the evaluations of system states makes it clear why dynamic programming methods are said to proceed from a task's end to its beginning. Computing V_1^π can be viewed as considering the decision task when one step remains before the horizon is reached; computing V_2^π can be viewed as considering the decision task when two steps remain before the horizon is reached, etc. For an infinite horizon task, which is our major concern here, it turns out that the desired evaluation function is the unique limiting function generated by continued successive application of Equation 6 (provided $\gamma < 1$) and, further, this evaluation function is the unique function, V^π , satisfying the following equation for each state x :

$$V^\pi(x) = R(x, \pi(x)) + \gamma \sum_{y \in X} P_{xy}(\pi(x)) V^\pi(y). \quad (7)$$

In fact, the function satisfying this equation is the unique limiting function generated by successive application of Equation 6 beginning with an *arbitrary* real-valued function of system states, and not just with V_1^π .⁹

The successive application of Equation 6 can be illustrated with the route-finding example. First, we know from the description of the task in Section 2 that $V_1^\pi(x) = -1$ for all locations x except the goal, where is zero. Now apply Equation 6 with $n = 2$ to determine $V_2^\pi(x)$ for each location x . To do this, note that because the task is deterministic, for any location x , $P_{xy}(\pi(x)) = 1$ only for the one location y , say $y = x'$, that is always reached from x by performing action $\pi(x)$. Therefore, $\sum_y P_{xy}(\pi(x)) V_1^\pi(y) = V_1^\pi(x')$. So we need only consider in turn each location, x , and its immediate sequel, x' , determined by the action $\pi(x)$. Applying Equation 6 with $n = 2$, if neither x nor x' is the goal, g , we have $V_2^\pi(x) = -1 + \gamma \times (-1) = -1 - \gamma$; if $x \neq g$ but $x' = g$, then $V_2^\pi(x) = -1 + \gamma \times 0 = -1$; and if $x = g$ then it must be that $x' = g$ and $V_2^\pi(x) = 0 + \gamma \times 0 = 0$. We can now apply Equation 6 with $n = 3$ to each location to determine V_3^π . Again consider in turn each location, x , and its immediate sequel, x' , determined by the action $\pi(x)$. Applying Equation 6 with $n = 3$, if neither x nor x' is the goal, we have $V_3^\pi(x) = -1 + \gamma \times (-1 - \gamma) = -1 - \gamma - \gamma^2$; if $x \neq g$ but $x' = g$, then $V_3^\pi(x) = -1 + \gamma \times 0 = -1$; and if $x = g$ then it must be that $x' = g$ and $V_3^\pi(x) = 0 + \gamma \times 0 = 0$. Continuing the application of Equation 6 eventually produces the location, or state, evaluations given above in Section 4.2.

5.2 Computing an Optimal Policy

We now discuss two methods for finding a policy that maximizes expected return under the assumption that a complete and accurate model of the decision task is available. Consequently,

⁹See Ross [51] for a detailed discussion and proof of these results. Equation 7 defines a set of simultaneous linear equations, one equation for each system state, which can be solved by matrix inversion. The successive application of Equation 6 is an iterative procedure for solving this system of equations. In practice, one continues to apply Equation 6 for increasing values of n until the difference between successive approximations becomes smaller than a predefined amount, or until some time limit is reached. Of course, any resulting function can also be checked via Equation 7 to see if it is the desired evaluation function. Choosing the initial function to be a good approximation of the desired evaluation function can considerably reduce the number of iterations required to achieve a given degree of accuracy, as can reducing γ . Although this method of successive approximations generally produces an approximation of the desired evaluation function, in what follows we loosely refer to the result of this computation as the actual evaluation function.

we assume that $R(x, a)$ and $P_{xy}(a)$ are known for all actions a and all states x and y . Recall that an optimal policy is a policy that cannot be improved upon in terms of expected return, that is, in terms of the expected total amount of payoff (suitably discounted and summed) it produces over the course of the decision task. The optimal evaluation function is the evaluation function corresponding to an optimal policy. Although there can be many optimal policies, there is only one optimal evaluation function.

If we know the optimal evaluation function, V^* , it is relatively easy to form an optimal policy. The optimal action when the system is in state x is the action, a , that maximizes

$$R(x, a) + \gamma \sum_{y \in X} P_{xy}(a) V^*(y). \quad (8)$$

This is the sum of the expected value of the immediate payoff, $R(x, a)$, and the expected value, taken over the states that can be reached in one step, of the infinite-horizon evaluations of possible next states, discounted by γ , under the assumption that an optimal policy is used throughout the future. Selecting an action that maximizes Expression 8 effectively answers the question "What action is best to perform *now* assuming that an optimal policy will be followed in selecting all future actions?" If we know V^* , it is relatively easy to select an action that maximizes the quantity given by Expression 8 because it is only necessary to examine the consequences of applying action sequences of length one. The consequences of all the longer action sequences are summarized in the optimal evaluation function, V^* . Because we are assuming a complete model of the decision task and knowledge of V^* , all the quantities involved in Expression 8 are known. Hence, an optimal action for state x can be found for by calculating Expression 8 for x and each action, a , and selecting the action that produces the largest result. Doing this for all system states x produces an optimal policy. For a discount factor close to zero, the optimal action for each state is most strongly influenced by the expected value of the immediate payoff, $R(x, a)$; for a discount factor equal to one, on the other hand, the estimated total amount of payoff that will be accumulated throughout the future, given by $\sum_{y \in X} P_{xy}(a) V^*(y)$, is given weight nearly equal to that of immediate payoff. Therefore the discount factor can be regarded as determining how the agent trades off an estimate of future return with an estimate—which is usually more accurate—of immediate payoff.

According to the above analysis, the task of finding an optimal policy can be reduced to the task of finding the optimal evaluation function. The optimal evaluation function can be approximated by a method of successive approximations similar to the one described above for approximating the evaluation function for a given policy (Equation 6). Here, however, it is necessary to determine the *maximum* expected return at each step. Let V_n^* denote the optimal evaluation function for the first n -steps of the infinite horizon decision task; that is, $V_n^*(x)$ is the maximum expected return if the decision task is terminated n steps after starting in state x . For $n = 1$, the maximum expected return is just the maximum of the expected values for the immediate payoff:

$$V_1^*(x) = \max_{a \in A} R(x, a). \quad (9)$$

Suppose, then, that we know V_1^* , and the agent now faces the two-step task starting in state x . If the agent performs an action a (whatever it may be), it will receive a payoff with expected value $R(x, a)$, and the next state of the system will be y with probability $P_{xy}(a)$. If we call the actual next state y , then the agent faces the one-step task from state y , and we already know that the maximum expected return for this one-step task is $V_1^*(y)$. Thus if the next system state is y , the maximum expected return for two steps starting in state x is obtained by performing the

action which maximizes the sum of the expected value of the immediate payoff, $R(x, a)$, and the one-step optimal evaluation, discounted by the factor γ , of y , which is $\gamma V_1^*(y)$. However, because the transition to the next state, y , occurs with probability $P_{xy}(a)$, instead of using the one-step optimal evaluation of the actual next state, one must use the expectation over all possible next states of the one-step optimal evaluations of these states, which is $\sum_y P_{xy}(a)V_1^*(y)$. We therefore have

$$V_2^*(x) = \max_{a \in A} \left[R(x, a) + \gamma \sum_{y \in X} P_{xy}(a) V_1^*(y) \right], \quad (10)$$

for all states x . In a similar manner we can determine V_3^* from V_2^* , V_4^* from V_3^* , and so on. The rule for determining a V_n^* in terms of V_{n-1}^* is:

$$V_n^*(x) = \max_{a \in A} \left[R(x, a) + \gamma \sum_{y \in X} P_{xy}(a) V_{n-1}^*(y) \right], \quad (11)$$

for all states x (cf. Equation 6).

The limiting values obtained by the repeated application of Equation 11 are the values of the desired optimal evaluation function, V^* . Moreover, this function is the unique function that satisfies the following equation, which is another basic equation of stochastic dynamic programming (cf. Equation 7):

$$V^*(x) = \max_{a \in A} \left[R(x, a) + \gamma \sum_{y \in X} P_{xy}(a) V^*(y) \right], \quad (12)$$

for all states x . The repeated application of Equation 11 is a method of successive approximations, called *value iteration*, for computing the optimal evaluation function. Here also it is not necessary to begin the process with V_1^* ; one can begin with an *arbitrary* real-valued function of the system state.¹⁰ Choosing this function to be a good approximation of the optimal evaluation function can reduce the number of iterations required to achieve a given degree of accuracy.

In order to find an optimal policy, then, one first approximates the optimal evaluation function, V^* , by value iteration (repeated applications of Equation 11) and then determines an optimal policy, π^* , by defining it to select an action that maximizes Expression 8 for each state x . If more than one action maximizes Equation 8, it does not matter which one is selected. If V^* is computed with sufficient accuracy, this process produces a policy (there may be others) that maximizes the expected return given by Equation 4.

To make value iteration more concrete, consider how it applies to the route-finding task. Assuming that we have no better approximation for the optimal evaluation function, we begin by determining the one-step optimal evaluation function using Equation 9. For the one-step task, no matter what action is chosen from any non-goal location, the one-step return is -1 . Hence, $V_1^*(x) = -1$ for all locations $x \neq g$, and clearly $V_1^*(g) = 0$. Applying Equation 11 with $n = 2$ to find the optimal two-step evaluation function can be visualized as follows. For each location x , determine the result of applying each of the four actions in terms of the sum of the immediate payoff and the discounted evaluation, γV_1^* , of the location that results from the action. The maximum of these four sums is $V_2^*(x)$. In this case, $V_2^*(x) = -1 - \gamma$ for all locations x except the goal, g , and those locations immediately adjacent to g : g retains an evaluation of zero, and

¹⁰ See Ross [51] for a detailed discussion and proof of these results. If there is a finite number of states and actions, as we are assuming here, this method converges after a finite number of iterations.

locations adjacent to g retain the evaluation -1 . In a similar manner, one can see that $V_3^*(x)$ is zero for $x = g$; it is -1 for locations immediately adjacent to g ; $-1 - \gamma$ for locations two steps away from g ; and $-1 - \gamma - \gamma^2$ for all other locations. As this computation is repeated, the evaluations of locations stabilize from the goal outwards, until the evaluations of all locations are determined. In the route-finding task, this happens in a finite number of applications of Equation 11 because all optimal paths must reach the goal, after which payoff no longer accumulates. An optimal policy always moves the agent to the adjacent location that has the largest evaluation according to the optimal evaluation function. If more than one adjacent location has this maximum evaluation, then it does not matter which one is chosen.

We now describe a different method of successive approximations for constructing an optimal policy that is more closely related to the learning method we describe in Section 7 than is value iteration. This method is called the *policy improvement method* or *policy iteration* [26, 51] because it generates a sequence of policies, each of which is an improvement over its predecessor.

Although one needs to know the optimal evaluation function in order to define an optimal policy, the policy improvement method is based on the observation that in order to *improve* a given policy one only needs to know the evaluation function for that given policy. Suppose the current policy is π and that we have constructed the evaluation function, V^π , for this policy. Further suppose that policy π' is some proposed policy, and we wish to see if π' is an improvement over π . Recall that a policy π' is an improvement over π if $V^{\pi'}(x) \geq V^\pi(x)$ for all system states x , with strict inequality holding for at least one state, where $V^{\pi'}$ and V^π are the evaluation functions, respectively, for policies π' and π . Thus, one way to determine if π' is an improvement over π would be to compute $V^{\pi'}(x)$ and to compare it with $V^\pi(x)$ for every state x . However, there is a much simpler way to do this that does not require constructing $V^{\pi'}$. Consider the expected return, starting in state x , that would result from using policy π' for one step, and then using policy π thereafter. The expected return for this composite policy is

$$R(x, \pi'(x)) + \gamma \sum_{y \in X} P_{xy}(\pi(x)) V^\pi(y). \quad (13)$$

If this quantity is greater than or equal to $V^\pi(x)$ for all states x , and strictly greater for at least one state, then it can be shown that π' is an improvement over π .

Instead of using Expression 13 just to check if a proposed policy is an improvement over the current policy, one can use it to define a new policy that improves over the current one. Define a new policy to be the one that for each state, x , specifies the action, $\pi'(x)$, that maximizes Expression 13. Thus, the new policy is an optimal policy with respect to the evaluation function of the current policy, V^π . It can be shown that this new policy is either an improvement over the current policy, or both the current and the new policies are optimal (e.g., Ross [51]). Because we are assuming that we know the evaluation function for the current policy, determining the new policy by maximizing Expression 13 requires looking only one step ahead from every state.

The policy improvement method is based on the foregoing ideas and works as follows. First select an *arbitrary* initial policy; call it π_0 . Then compute the evaluation function, V^{π_0} , for π_0 by solving the system of equations defined by Equation 7 by the repeated application of Equation 6 or some other means. Given V^{π_0} , determine a policy which is optimal with respect to it by defining that policy to select for each state, x , an action, a , that maximizes

$$R(x, a) + \gamma \sum_{y \in X} P_{xy}(a) V^{\pi_0}(y).$$

Call the resulting policy π_1 . Policy π_1 is either an improvement over π_0 , or they are both *optimal*.

If they are not optimal, the next step is to compute the evaluation function, V^{π_1} , for policy π_1 , and then define a policy that is optimal with respect to it. Call this policy π_2 . It will be an improvement over π_1 , or both π_1 and π_2 are optimal. The method continues in this way, as policies are formed, evaluation functions corresponding to them are computed, and then new, improved policies are formed. If the state set, X , of the system underlying the decision task is finite, this method will produce an optimal policy after some finite number of iterations.

This completes our exposition of the basic concepts and computational methods of stochastic dynamic programming, all of which are well-known to control and decision theorists. The procedures described above—for approximating the evaluation function for a given policy, the optimal evaluation function, and an optimal policy—require both a model of the decision task and can require considerable amounts of computation. These concepts and methods permit clear statements of some of the most important things we would like an agent to learn from its experiences interacting with a dynamical system, but they do not provide more than hints about how to devise learning methods that can operate in real-time and do not require a model of the decision task. For this we have to turn to another body of theory, which is also well-known within mathematical and engineering disciplines, that concerns techniques for estimating parameter values. We shall see in Section 7 that the TD procedure is a method for estimating the parameters specifying the evaluation function for a given policy in the absence of a model of the decision task. We shall also describe how the TD procedure can be incorporated into real-time procedures for improving decision policies.

6 Parameter Estimation

The TD procedure is method for estimating an evaluation function by means of *parameter estimation*. Methods for parameter estimation are central to the fields of pattern classification (e.g., Duda and Hart [16], Sklansky and Wassel [58]), adaptive signal processing (e.g., Widrow and Stearns[73]), and adaptive control (e.g., Goodwin and Sin [20]), as well as the field of connectionist modeling (e.g., refs. [4, 23, 37, 53]). For example, one way to construct a pattern classifier is to assume that the classification rule, or decision rule, is a member of a specific class of rules, where each rule in the class is specified by selecting values for a set of parameters. Learning refers to the process of adjusting the parameter values to improve the decision rule's performance based on its performance for a set of training patterns whose correct classifications are supplied by a "teacher."

In signal processing and control, one is often seeking mathematical formulae capable of producing numerical values that match values measured from some physical system. One may wish, for example, to have a formula giving a patient's blood pressure as a function of the amount of a drug administered [73]. Such a formula constitutes a *model* of a functional relationship implemented by the physical system, and it can be used for a variety of prediction and control purposes. We call this functional relationship the *modeled function*.¹¹ Given an encoding of some input information (e.g., amount of drug administered), the model output should match the output of the modeled function for that same input (e.g., the patient's resulting blood pressure). A parameter estimation method adjusts parameter values based on data in the form of examples of how the

¹¹In adaptive control, where the modeling process is called *system identification*, the functional relationship being modeled is a dynamical system which is usually more complicated than a function from input to output due to the influence of internal states [20]. For our purposes, however, it suffices to discuss only the modeling of input-output functions.

modeled function behaves for a collection of cases.¹²

Among the most important elements in formulating a parameter estimation task is the selection of a method for representing the physical aspects of the problem as signals that act as input and output of the decision rule or model. What aspects of objects, inputs, states, etc. should be measured to provide data for decision making or modeling? If basic physical variables are obvious candidates (e.g., blood pressure), should they be represented simply as real-valued measurements, or should they be coded in some other way? Similarly, how is the output of the decision rule or model to be interpreted? There is little useful theory to guide these aspects of formulating a parameter estimation task. Another aspect of formulating parameter estimation tasks for which there is little theory is the selection of the parameterized class of decision rules or models, a choice that includes deciding how many parameters to use and how their values define the rule or model. Generally, one must make a compromise between rule or model complexity and its adequacy for the application. In this report, we restrict attention to one of the simplest classes of models and use a very simple representation for the route-finding example, but we indicate how more complex choices can be accommodated.

Other important elements in formulating a parameter estimation task are the choice of performance measure, the use of *a priori* knowledge, and the type of estimation algorithm most appropriate.¹³ Estimating parameters depends on comparing the performance of the different decision rules or models specified by different parameter values. The best measure of performance might evaluate the overall, final performance of the system that makes use of the decision rule or model. However, in practice one has to devise simpler criteria that are closely correlated with overall performance and yet can be measured easily and quickly. It is common to measure performance by the average of the squared error between the output of the decision rule or model and some target values.

The parameter estimation process can be improved by the use of *a priori* knowledge, not only as it is incorporated into the selection of the task representation and the class of decision rules or models, but also in the form of initial parameter values and constraints among parameter values. The use of parameter estimation methods by no means entails a *tabula rasa* view of learning and memory. On the contrary, the incorporation of prior knowledge can accelerate, and otherwise improve, the learning process.

Two major types of parameter estimation methods are distinguished: *off-line* and *on-line* methods. Off-line methods operate when all the training information, e.g., a set of classified patterns, is available at one time. These patterns are collected before performing the analysis that determines suitable parameter values. On-line methods, on the other hand, process training information as it becomes available to the learning system. In this report, we shall be concerned exclusively with on-line methods because we assume that most animal learning occurs on-line (although processes of memory consolidation might be construed as the result of some type of off-line method).

¹²Nearly all of the learning methods for connectionist networks are parameter estimation methods, where the form of the decision rule or model is determined by the structure of the network, which includes the types of units, how many of them there are, and how they are interconnected. The parameter values correspond to the synaptic weights. Similarly, many models of animal learning, such as the Rescorla-Wagner model [47] and the TD model [64, 62], can be regarded as parameter estimation methods where the parameter values correspond to the associative strengths of the components of stimuli.

¹³See Goodwin and Sin [20] whose discussion of these issues forms the basis for our remarks.

6.1 Feature Vectors, Decision Rules, and Models

The framework in which parameter estimation techniques are applied is essentially the same for classification tasks and tasks requiring the modeling of functional relationships. For a classification task, let x denote any one of the objects we wish to classify. For a function modeling task, let x denote any of the collections of data that can be provided as input to the modeled function. In what follows, we call all such collections of data *patterns*, even though this is not the usual terminology for many function modeling tasks. Suppose that a set of measurements can be made of pattern x to produce a description in the form of a list of attribute or feature values. If there are n measurements, let ϕ_i , $i = 1, \dots, n$, denote the "measuring instruments," and if we assume that the measurements are real numbers, as is usual, each ϕ_i is a function from the set of possible patterns to the real numbers.¹⁴ Writing $\phi_i(x)$ to denote the value of the i^{th} measurement of x , the vector $\phi(x) = (\phi_1(x), \dots, \phi_n(x))^T$ is the *feature vector* description of x (we assume here that vectors are column vectors, so the superscript T indicates that the row form shown is the transpose of the default column form). One can think of $\phi(x)$ as the parameter estimation system's internal representation of pattern x .

For a classification task suppose that each decision rule among the rules to be considered is defined by a vector of $n + 1$ parameters $v = (v_1, \dots, v_{n+1})^T$, where each v_i , $i = 1, \dots, n + 1$, is a real number. For example, one might assume that the decision is made by weighting each feature value $\phi_i(x)$ by v_i and comparing the weighted sum to a threshold given by $-v_{n+1}$ (the minus sign is present for a technical reason to be made clear below). The object is in class 1 if the weighted sum exceeds zero; it is in class 2 if the weighted sum is less than zero. Each decision rule of this type has the form:

$$x \text{ is } \begin{cases} \text{class 1} & \text{if } v_1\phi_1(x) + \dots + v_n\phi_n(x) > -v_{n+1} \\ \text{class 2} & \text{if } v_1\phi_1(x) + \dots + v_n\phi_n(x) < -v_{n+1}, \end{cases} \quad (14)$$

for some parameter vector v . This is the linear threshold rule often used in connectionist networks where it is implemented by an abstraction of a neuron. The parameter values (synaptic weights) specifying the best approximation to the desired decision rule are initially unknown and are estimated through the use of a parameter estimation method.

In a modeling task, one of the simplest classes of models consists of linear models defined by weighted sums of n feature values, plus the parameter v_{n+1} :

$$v_1\phi_1(x) + \dots + v_n\phi_n(x) + v_{n+1}. \quad (15)$$

This sum is the output of the model specified by the parameter vector v for input x . Although the development of efficient parameter estimation methods for nonlinear models is important for the utility of the methods discussed in this report, the issues we address here are very different, and for simplicity, we assume the use of linear models. It should be clear, however, that the learning procedures we discuss can be extended to nonlinear cases (see, for example, refs. [2, 3]).¹⁵

Two conventions are usually followed when decision rules or models in the form of Expressions 14 and 15 are used. First, to the list of n functions ϕ_i , one adds an $(n + 1)^{\text{st}}$ which is always 1;

¹⁴For some applications, it is common to let some, or all, measurements produce only two values to indicate, for example, whether a logical predicate applied to the object or state is true or false. This is a special case of the framework adopted here.

¹⁵Notice that even in the case of a linear model, the model output can be a nonlinear function of input patterns because the ϕ functions can be nonlinear functions of the patterns. But because these functions are not parameterized, they do not directly enter into the derivation of a parameter estimation method.

that is, for any pattern x , $\phi_{n+1}(x) = 1$. This permits the decision rule given by Expression 14 to be written as follows:

$$x \text{ is } \begin{cases} \text{class 1} & \text{if } v^T \phi(x) > 0 \\ \text{class 2} & \text{if } v^T \phi(x) < 0, \end{cases} \quad (16)$$

where $v^T \phi(x) = v_1 \phi_1(x) + \dots + v_{n+1} \phi_{n+1}(x)$ is the inner product, or scalar product, of the vectors v and $\phi(x)$. Similarly the parameterized linear model given by Expression 15 can be written

$$v^T \phi(x). \quad (17)$$

A second useful convention is to refer to the *time* at which pattern x appears for classification, or as input to the modeled function, instead of referring directly to x . Let ϕ_t and v_t respectively denote the feature vector and the parameter vector at time t . The set of values that the time variable, t , can take depends on whether one considers continuous or discrete-time parameter estimation. We restrict attention to discrete-time methods, so that t is a non-negative integer.

6.2 Parameter Estimation Methods

The methods for parameter estimation that we consider are error-correction methods because they adjust parameters based on errors between actual and desired outputs of the decision rule or model. Error-correction methods operate by receiving sequences of training pairs, where each pair consists of a feature vector and the output desired of the decision rule or model when it is presented with that feature vector as input. The collection of all training pairs that are available comprises the training data for the parameter estimation task. This paradigm is known as "supervised learning" or "learning with a teacher." In the classification task, a "teacher" provides the correct object classifications, whereas in the modeling task, the modeled function acts as a teacher. In what follows, we discuss parameter estimation in the context of modeling a functional relationship instead of forming a decision rule because the TD procedure can be best understood in this context.

An on-line estimation method operates as follows (see Figure 3). At time step t , there is a current estimate, v_t , for the unknown parameter vector that specifies the best model within the class of models under consideration. There is also available a feature vector, ϕ_t , describing the current input, x_t , to the modeled function. The feature vector is used as input to the model specified by the parameter estimate v_t , and the resulting output of this model is an estimate for the output of the modeled function given input x_t . For the linear model given by Expression 17, this estimate is $v_t^T \phi_t$. This estimate is then compared with the actual output of the modeled function, which we assume becomes available at the next time step, $t + 1$. Letting y_{t+1} denote this output of the modeled function, the error, ϵ_{t+1} , in the model's estimate is the scalar

$$\epsilon_{t+1} = y_{t+1} - v_t^T \phi_t. \quad (18)$$

In the most widely used methods, a new parameter vector, v_{t+1} , is determined from the current vector, v_t , as follows:

$$v_{t+1} = v_t + M_t \epsilon_{t+1} \phi_t, \quad (19)$$

where M_t is sometimes called the algorithm "gain" at step t . A number of on-line algorithms for parameter estimation follow the form of Equation 19 but differ in what M_t is, and how it is computed. In some algorithms, M_t is an $(n + 1) \times (n + 1)$ matrix computed in an iterative manner

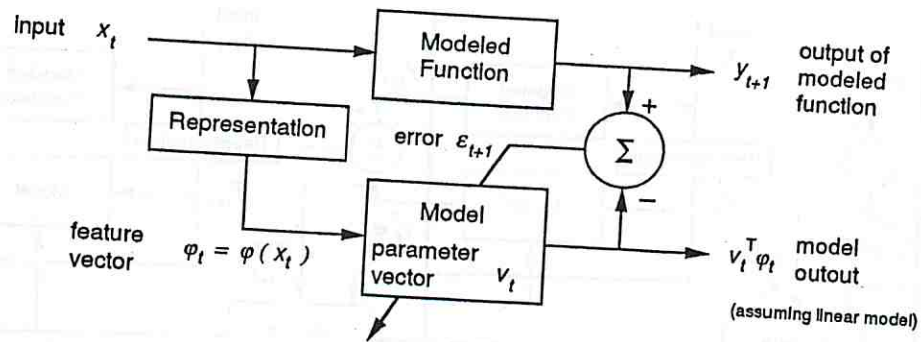


Figure 3: On-line parameter estimation for modeling a functional relationship. The error, ε_{t+1} , resulting from comparing the output of the modeled function, y_{t+1} , with the model's estimate of this output, $v_t^T \phi_t$, is used to update the parameter vector, v_t . The input to the model is the feature vector, $\phi_t = \phi(x_t)$, determined from the input to the modeled function, x_t , by the box labeled "Representation." The output of the modeled function is assumed to be a linear function of the feature vectors. The circle labeled " Σ " in this and following figures is a summation unit, with the signs of the summed quantities labeling each input.

based on the training data [20]. Here we restrict attention to the simplest case in which M_t is replaced by a positive constant, β ,¹⁶ which, together with the magnitude of the error, determines the magnitude of changes in parameter values from one step to the next. In this case, we can write the following parameter-update rule:

$$v_{t+1} = v_t + \beta \varepsilon_{t+1} \phi_t. \quad (20)$$

Figure 4 is an elaboration of Figure 3 showing how the various quantities involved in on-line parameter estimation change over time and on what quantities their successive values depend. It shows these quantities for two successive time steps. The box labeled "Update Parameters" implements a parameter-update rule such as the one given by Equation 20. The inputs to this box correspond to the variables on the right-hand side of Equation 20.

Because we express parameter-update rules in the vector notation of Equation 20 throughout this report, it is important to understand how an equation written in this form specifies how the value of each parameter making up the parameter vector is updated. Referring to Equation 20, β and ε_{t+1} are single numbers (scalars) and v_t and ϕ_t are vectors having the same number of components. Hence, this equation indicates that at time step t , one adds $\beta \varepsilon_{t+1}$ times the feature vector ϕ_t to the parameter vector v_t to form the next parameter vector, v_{t+1} . Focusing on the i^{th} individual parameter value, this means that its value at step $t+1$ is its old value plus $\beta \varepsilon_{t+1}$ times the value of the i^{th} feature at step t .

If the parameterized linear model given by Expression 17 is assumed, then Equation 20 can be expanded via Equation 18 to yield:

$$v_{t+1} = v_t + \beta(y_{t+1} - v_t^T \phi_t) \phi_t. \quad (21)$$

¹⁶This is the special case in which each entry in M_t is equal to β for each t .

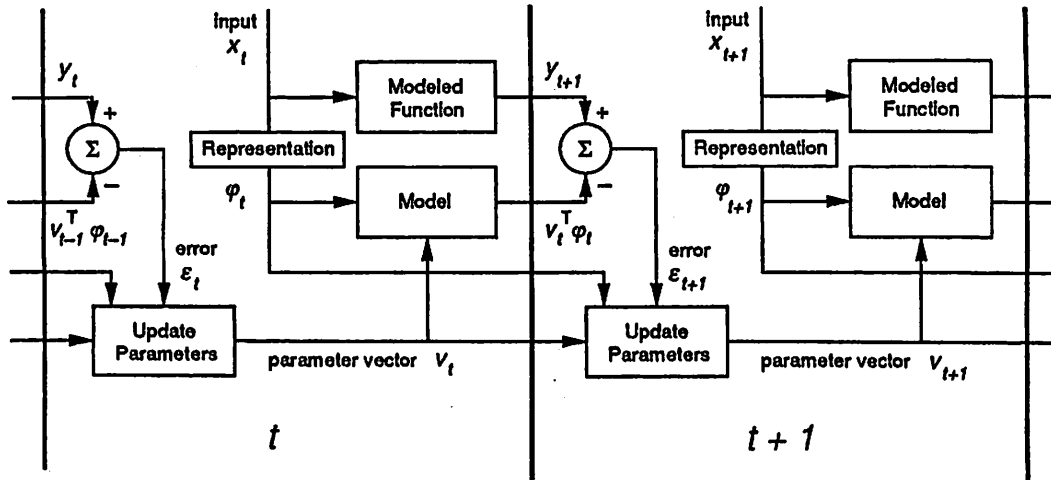


Figure 4: How the various quantities involved in on-line parameter estimation change over time. Values for two successive time steps are shown, and a linear model is assumed. The box (shown once for each time step) labeled “Update Parameters” implements a parameter-update rule such as the LMS rule.

This is the estimation method presented by Widrow and Hoff [72] in 1960, also known as the LMS rule, that is essentially the same as the Rescorla-Wagner model of classical conditioning [47, 63].¹⁷

Some on-line estimation procedures, such as the LMS rule, can be understood as gradient-descent procedures, which are methods for finding a local minimum of a surface by moving down the slope, or gradient, of the surface.¹⁸ According to this view, there is a surface, defined over the space of parameter vectors, whose height at each parameter vector gives a measure of the error over all patterns of the model specified by that parameter vector. The task is to search for a parameter vector at which this surface is minimum. Updating estimates according to Equation 19 tends to form new estimates that are steps down on this error surface. The gain matrix, M_t (which in the LMS rule (Equation 21) is replaced by the constant β) adjusts the direction and size of the step on the error surface. For the LMS rule, the error measure is the mean value of the squared error of the model over all patterns (hence its name).¹⁹

We briefly discuss how the LMS rule is derived in terms of gradient descent because the derivation requires a critical step similar to one we employ in explaining the TD procedure. The

¹⁷If one replaces the linear model of Equation 15 with the linear threshold decision rule given by Expression 16, one obtains the Perceptron learning rule of Rosenblatt [50], which is identical to Equation 21 except that the weighted sum $v_i^T \phi_i$ is replaced by the threshold of the weighted sum as defined by Expression 16. These correspondences are explained in more detail in refs. [6, 16, 63].

¹⁸A function, f , from an n -dimensional space to the real numbers can be viewed as a surface. If x is a point in the n -dimensional space, then the gradient of f at x is the vector $(\frac{\partial f}{\partial x_1}(x), \dots, \frac{\partial f}{\partial x_n}(x))^T$, where $\frac{\partial f}{\partial x_i}(x)$ is the partial derivative of f with respect to the i^{th} dimension evaluated at the point x . This vector points in the direction of the steepest increase of the surface at the point x .

¹⁹The error back-propagation method [52] is derived by computing this error gradient for a particular class of models in the form of layered connectionist networks.

objective is to adjust the parameter values of the linear model in order to minimize a measure of model error over all patterns. Letting, as above, $\varepsilon_{t+1} = y_{t+1} - v_t^T \phi_t$, one wants to minimize $E[\varepsilon_t^2]$, where E is the expectation over all patterns, and the error is squared to produce a quantity that is always positive. This quantity—the mean square error—depends on the parameter values. At each step of the estimation process, one would like to adjust the parameter values in such a way that the mean square error is reduced. Given that one knows the mathematical expression defining this error measure, it is possible to pre-compute the gradient of the error measure as a function of the parameters values. The result of this computation specifies that one has to add to the current parameter vector, v_t , a vector proportional to

$$E[\varepsilon_{t+1} \phi_t], \quad (22)$$

where, again, E is the expectation over all possible patterns.

Unfortunately, in an on-line procedure, one cannot determine at each time step the quantity given by Expression 22 because it depends on how well the current parameter values work for all possible patterns. The critical step, therefore, in deriving the LMS rule is to assume that $\varepsilon_{t+1} \phi_t$ for just the *current* feature vector, ϕ_t , is a good estimate of the quantity given by Expression 22. Making this assumption yields the parameter-update equation of the LMS rule given above (Equation 21). Because the expected value of $\varepsilon_{t+1} \phi_t$ over all patterns is the gradient of the error measure that is to be minimized (Expression 22), $\varepsilon_{t+1} \phi_t$ is an unbiased estimate of this gradient. For the class of parameterized linear models defined by Expression 17, the LMS rule can be shown to minimize the mean square error over all patterns under appropriate conditions and with certain caveats that need not concern us here (Widrow and Stearns [73]).

The theory of the LMS rule, and other parameter estimation methods, is very well-developed but beyond the scope to the present report. Theoretical treatments of parameter estimation can be found in refs. [16, 20, 33, 58, 73]. This theory specifies conditions under which a parameter estimation method converges to a final estimate and what criterion of best fit the final estimate satisfies. To obtain these theoretical results, the manner of presenting training data to the algorithm must be assumed to satisfy certain properties. For modeling tasks, for example, the output of the modeled function has to be observed for a sequence of inputs that is both sufficiently varied to reveal the function's form and sufficiently repetitive to make learning possible. A set of inputs satisfying these properties, is said to be *sufficiently stimulating*.²⁰ Theoretical results concerning parameter estimation for classification tasks similarly require a training regime that repeats the training pairs sufficiently often or uses as training data a sufficiently large sample of the set of possible patterns.

7 Learning and Sequential Decision Tasks

We are now in a position to combine the concepts from stochastic dynamic programming described in Section 5 with those just described concerning parameter estimation. We address tasks that require *learning* evaluation functions and optimal decision policies in the absence of the knowledge necessary to apply the dynamic programming techniques presented in Section 5. Although these dynamic programming techniques are therefore not applicable, the principles they involve can be applied to the design of learning methods. In Section 7.1, we consider the problem of learning the

²⁰This term is used in the context of modeling dynamical systems where it has a specific meaning depending on the system being modeled. The inputs to the system have to stimulate all of the system's "modes of activity." See, for example, ref. [20].

evaluation function for a given decision policy. The TD procedure (Equation 27) emerges from this analysis as an on-line method for this task. Although a method for learning an evaluation function for a given policy does not itself improve a policy in terms of expected return, the evaluation function estimates that result from applying such a method have great utility for improving policies. An evaluation function provides an immediately-accessible prediction of the long-term return a policy will achieve throughout the future. This information can be used in adjusting policies so that decisions are not dominated by the short-term consequences of behavior. In Section 7.2 we consider learning the optimal evaluation function and an optimal decision policy by a method that takes advantage of evaluation function estimates produced by the TD procedure.

For all of these learning tasks, we must assume that the agent has the opportunity to interact with the system underlying the task in a manner yielding stimulation sufficient to produce good parameter estimates. How the requirement for sufficient stimulation can be met generally depends on details of the system, and the issues that arise are beyond the scope of the present report. Here it suffices to assume that the system underlying the decision task is set to a randomly chosen state, either periodically or when certain conditions are met. For example, in the route-finding task, we place the agent at a randomly chosen location, let it interact with the system (i.e., move around) while updating its parameter values for a given number of time steps or until it reaches the goal. This constitutes a trial. We run a sequence of trials until some performance criterion is satisfied or a time limit is reached.

7.1 Learning an Evaluation Function

Here we consider the problem of learning an evaluation function for a given policy in a sequential decision task. The learning system is given the opportunity to interact for a sequence of trials with the system underlying the decision task. A fixed policy, π , is used throughout the sequence of trials. We wish to form the evaluation function, V^π , for policy π without having prior knowledge of the quantities $R(x, a)$ and $P_{xy}(a)$, for system states x and y and actions a . Recall that V^π assigns to each state, x , its value, which is the expected return if policy π is used to select actions after the system enters state x (Equation 4).

We formulate the problem of learning an evaluation function as a parameter estimation task. A set of features for representing system states must be selected. Let $\phi(x)$ denote the feature vector representing state x . A parameterized class of models for the evaluation function is then selected, and the task is to adjust the parameter values in order to best approximate the true evaluation function according to a measure of how closely the model's evaluations match the true evaluations of states.²¹ Here, for simplicity, we select the class of linear models given by Expression 15 or 17 and consider models that assign to each state, x , the evaluation estimate

$$v_1\phi_1(x) + \dots + v_n\phi_n(x) + v_{n+1} = v^T\phi(x), \quad (23)$$

for parameter vector $v = (v_1, \dots, v_{n+1})^T$.

Because we are seeking an on-line estimation method, we refer to evaluation estimates, parameter vectors, and feature vectors by the times at which they occur. We let V_t denote the estimate at time step t for the evaluation function. Letting v_t denote the parameter vector at time step t , the estimate at time step t for the evaluation of state x is

$$V_t(x) = v_t^T\phi(x). \quad (24)$$

²¹Note that the model considered here is a model of the evaluation function, not of the decision task; that is, it does not model the state-transition and payoff probabilities of the system underlying the decision task.

If we were to follow the general form for parameter estimation described in Section 6.2, given the current system state x_t , an error, ε_{t+1} , would be determined by comparing the evaluation estimate of x_t , $V_t(x_t)$, with its actual evaluation, $V^\pi(x_t)$, given by the true evaluation function for policy π . Unfortunately, we do not know the true evaluation function V^π . However, we do know a condition that V^π must satisfy. From Equation 7 we know that:

$$V^\pi(x_t) = R(x_t, a_t) + \gamma \sum_{y \in X} P_{x_t, y}(a_t) V^\pi(y),$$

and we can define an error based on how much the current estimate for the evaluation function departs from this condition for state x_t :

$$\varepsilon_{t+1} = \left[R(x_t, a_t) + \gamma \sum_{y \in X} P_{x_t, y}(a_t) V_t(y) \right] - V_t(x_t). \quad (25)$$

If it were possible to adjust the parameter vector v_t to reduce this error, the new evaluation function estimate would tend to be closer to the actual evaluation function. Updating the parameter vector this way would be similar to applying a step of the successive approximations method specified by Equation 6, but only to state x_t . The quantity $R(x_t, a_t) + \gamma \sum_{y \in X} P_{x_t, y}(a_t) V_t(y)$ should be a better estimate for $V^\pi(x_t)$ than is $V_t(x_t)$.

However, it is still not possible to determine the error given by Equation 25 in an on-line estimation method because it depends on the unknown payoff expectations, $R(x_t, a_t)$, and the unknown transition probabilities $P_{x_t, y}(a_t)$, for a_t , x_t , and all states y . The TD procedure is a consequence of replacing these unknown quantities with approximations that are available during on-line learning. First, we substitute the payoff actually received at step $t+1$, which is r_{t+1} , for the expected value of this payoff, $R(x_t, a_t)$. Because $E[r_{t+1} | x_t, a_t] = R(x_t, a_t)$ (Equation 1), r_{t+1} is an unbiased estimate of $R(x_t, a_t)$. Second, we substitute the current evaluation estimate of the state *actually reached* in one step for the expectation of the evaluation estimates over the states *reachable* in one step. That is, we use $V_t(x_{t+1})$ in place of $\sum_y P_{x_t, y}(a_t) V_t(y)$ in Equation 25. $V_t(x_{t+1}) = v_t^T \phi_{t+1}$ is the evaluation estimate of the *next* system state using the *current* parameter values. This substitution does not necessarily result in an unbiased estimate of the desired quantity because the evaluation estimate of x_{t+1} may not be an unbiased estimate of the true evaluation of x_{t+1} . Nevertheless, we assume that the substitution results in a satisfactory approximation, and we approximate the error given by Equation 25 by the following expression, called the TD error, involving only quantities that are available at time step $t+1$:

$$\bar{\varepsilon}_{t+1} = [r_{t+1} + \gamma V_t(x_{t+1})] - V_t(x_t). \quad (26)$$

Finally, substituting this error estimate into the equation for an on-line estimation method (Equation 20), we arrive at the TD procedure:

$$v_{t+1} = v_t + \beta [r_{t+1} + \gamma V_t(x_{t+1}) - V_t(x_t)] \phi_t. \quad (27)$$

This is identical to a special case of the TD model of classical conditioning [62] obtained by setting the parameters α_i of that model to one and the parameter δ to zero. If δ were present within the framework developed here, setting it to a nonzero value would have the effect of replacing ϕ_t in Equation 27 with ϕ_t plus a weighted sum of the feature vectors representing states through which the system passed in time steps preceding step t . Using such a stimulus trace in the TD

procedure can accelerate the learning process as shown by Sutton [61], and Watkins [67] shows how this kind of stimulus trace can be viewed within the framework of dynamic programming.

The assumptions made in justifying the TD procedure are similar to the assumption made in deriving the LMS rule as a gradient-descent procedure discussed in Section 6.2. The LMS rule is a consequence of assuming that the error in responding to a single training pattern can be used to approximate the model's expected error over all patterns. Similarly, the TD procedure is a consequence of assuming that the payoff actually received at a time step is a good estimate for the expected value of the payoff at that step, and that the evaluation estimate of the next state is a good estimate of the expected value of the evaluations of the possible next states. Despite this similarity to the LMS rule, the situation is considerably more complex for the TD procedure, and it is not as easy to interpret the TD procedure in terms of gradient descent. One source of complexity is that the training examples experienced by the TD procedure can be influenced by the agent's actions so as to bias the estimate of $\sum_y P_{x_t,y}(a_t)V_t(y)$. Sutton [61] has been able to prove that under certain conditions, the TD procedure given by Equation 27 converges to form the desired evaluation function with repeated trials, each consisting of a finite number of time steps. Among the conditions required for convergence is that the parameter values are not updated during a trial. The parameter adjustments determined by the TD procedure are accumulated during each trial and used to update the parameter values only at the end of that trial. In practice, however, it is often adequate to update parameter values at each time step.

It is beyond the scope of this report to discuss the more theoretical aspects of the TD procedure. The theory of the TD procedure is neither as simple, nor as fully developed, as is the theory of the LMS rule. However, it is noteworthy that a special case of the TD procedure is thoroughly understood because it *is* the LMS rule. When γ is set equal to zero in Equation 27, the result is

$$v_{t+1} = v_t + \beta [r_{t+1} - V_t(x_t)] \phi_t, \quad (28)$$

which is exactly the LMS rule given by Equation 21. We can think of the LMS rule used in this way as a one-step-ahead adaptive predictor, which is a standard application of the LMS rule [73]. Applied to a sequential decision task, this rule adjusts the parameter values so that the evaluation estimate, $V_t(x_t)$, of the state at time step t , predicts the payoff to be received at the next time step: r_{t+1} . Setting γ equal to zero in Equation 27 therefore renders the TD procedure capable of estimating only the most "short-sighted" evaluation function.

Let us summarize the application of the TD procedure given by Equation 27. We refer to Figure 5, which shows two successive time steps of the application of the TD procedure. This figure is a combination of the diagram of the system underlying the decision task shown in Figure 2 (appearing at the top of Figure 5, although here the 'disturbances' shown in Figure 2 are omitted to simplify the figure) with the diagram of on-line parameter estimation shown in Figure 4, where the latter has been modified to according to the special features of the TD procedure.

An agent employing the TD procedure interacts with a dynamical system that can be in any one of a number of states at each time step. The agent uses some fixed policy, π , for selecting its actions, which influence the behavior of the system. At present, it does not matter what this policy is: it could be a good policy or a bad one in terms of its expected return for system states. The agent is attempting to estimate the function that associates to each state an estimate for the expected return of policy π .

Although the expected return depends on the current state of the system, the agent does not necessarily receive enough information to accurately determine the state. At time step t , the agent has access to the feature vector $\phi_t = \phi(x_t)$, determined from state x_t by the box

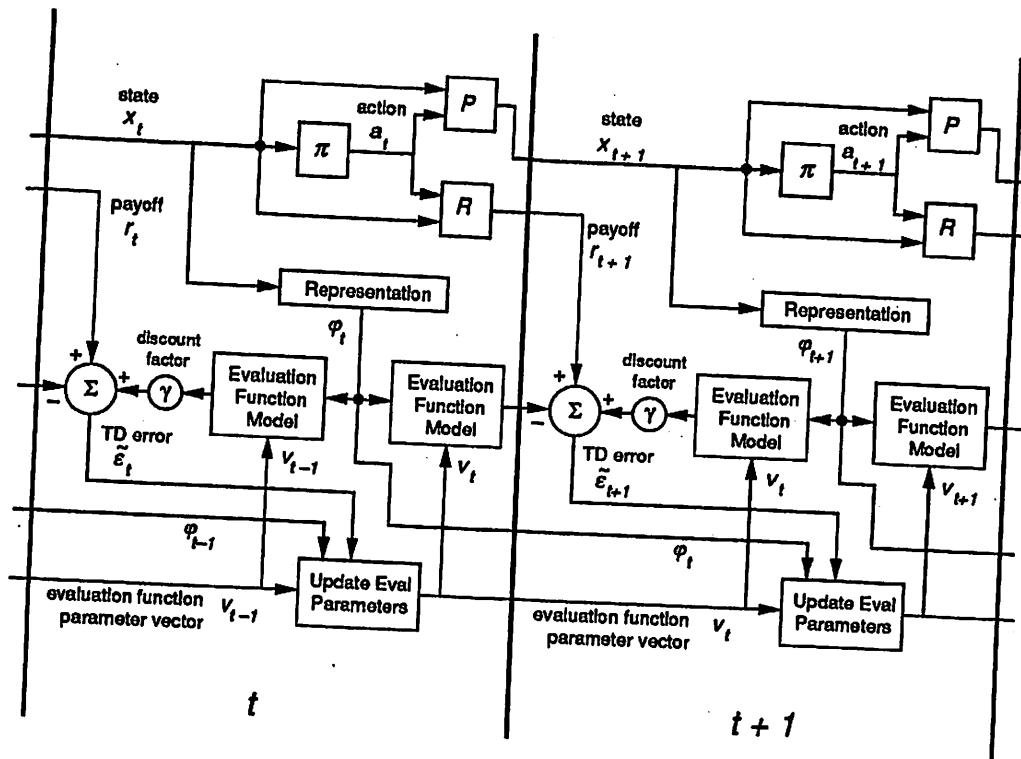


Figure 5: Two successive time steps of application of the TD procedure. This figure is a combination of the diagram of the system underlying the decision task shown in Figure 2 (here appearing at the top of the figure) with the diagram of on-line parameter estimation shown in Figure 4, where the latter is modified to according to the special features of the TD procedure. Notice that if $\gamma = 0$, the left-hand box labeled "Evaluation Function Model" in each time period becomes disconnected, and the lower part the figure takes the same form as Figure 4. In this figure and Figure 6, crossing lines do not represent connections unless there is a dot on the intersection.

labeled "Representation" in Figure 5. One can think of this feature vector as a pattern of sensory stimulation produced by the current state of the system, or as a representation provided by some mechanism, not discussed here, which estimates the system's state. Although states with different evaluations may give rise to the same feature vector, and states with similar evaluations may give rise to different feature vectors, it is obviously better if the representation assigns similar feature vectors to states with similar evaluations. Better estimates and faster learning are possible if the representation is suited to the task. Ideally, one would like features that are highly correlated with the evaluations of the states. For example, a wind-borne odor may be strongly correlated with high return if the policy is to travel up-wind. In any case, assume that the representation is given as well as the parameterized class of models for the evaluation function. Here, we assume that the evaluation function is a weighted sum of the feature values, and the agent has the task of estimating the weights.

Upon receiving the feature vector ϕ_t , the agent determines its estimate for the expected return using its current parameter vector, v_t . This is the evaluation estimate of state x_t given by $V_t(x_t) = v_t^T \phi_t$. It is computed by the *right-most* box labeled "Evaluation Function Model" within the time- t part of Figure 5. (At this point, the *left-most* box labeled "Evaluation Function Model" within the time- t part of the figure has already been used to provide information for updating v_{t-1} to form the current parameter vector v_t , a process we explain below in terms of time step $t + 1$.) Policy π is used to select an action based on the current system state, x_t . After the action is performed at step t , the system makes a transition to a new state, x_{t+1} , and the agent receives another feature vector, ϕ_{t+1} , and a payoff, r_{t+1} .

Refer now to the time- $(t + 1)$ part of Figure 5. Using parameter vector v_t (which has not been updated yet), the agent computes an estimate for the return expected from time step $t + 1$ onwards. This is $V_t(x_{t+1})$, which is computed by the *left-most* box labeled "Evaluation Function Model" within the time- $(t + 1)$ part of Figure 5. Now, if the linear model of the evaluation function is adequate and the parameter vector is correct, then on the average it should be true that the evaluation estimate of the state at time t , $V_t(x_t) = v_t^T \phi_t$, equals r_{t+1} plus γ times the evaluation estimate of the state at time $t + 1$; that is, $V_t(x_t)$ should, on the average, equal $r_{t+1} + \gamma V_t(x_{t+1})$. The discrepancy between these estimates is the error, $\tilde{\epsilon}_{t+1}$ (determined by the summation unit on the left-hand side of the time- $(t + 1)$ part of Figure 5) which is used to update the parameter vector v_t . This update process is shown in Figure 5 as the box in the $t + 1$ area labeled "Update Eval Parameters" that implements Equation 27 to produce the parameter vector v_{t+1} . Using this parameter vector, state x_{t+1} is *re-evaluated* by the *right-most* box labeled "Evaluation Function Model" in Figure 5, and the result is available at time step $t + 2$ for a repeat of the steps described above.²²

Figure 5 makes a subtlety of the TD procedure clear. At each time step, two estimates of the evaluation of the system state are computed—the first using the parameter vector before it is updated, and the second using the parameter vector after it has been updated. This is why there are two boxes labeled "Evaluation Function Model" for each time step in Figure 5 (although, of course, one can think of the computation for each time step as using the same box twice, using a different parameter vector each time).²³

²²Note that when γ is zero—the case of the TD procedure equivalent to a one-step-ahead LMS predictor—the *left-most* "Evaluation Function Model" box for each time step becomes disconnected, and the flow of computation is identical to that shown in Figure 4 for the task of modeling an unknown function whose true evaluations become available to the estimation process in one time step.

²³The procedure used by Barto, Sutton, and Anderson [9] for the "Adaptive Critic Element" in their pole balancer

The Route-Finding Example—To illustrate the TD procedure applied to the route-finding example, we use one of the simplest state representations possible. We assume the 96 states (spatial locations) are numbered 1 to 96, and we represent each state, x , by a feature vector of 96 components: $\phi(x) = (\phi_1(x), \dots, \phi_{96}(x))^T$ where

$$\phi_i(x) = \begin{cases} 1 & \text{if } x = i \\ 0 & \text{otherwise.} \end{cases} \quad (29)$$

The feature vector $\phi(x)$ therefore is a list of zeros with a single one in the position corresponding to state x . Given this set of features, the linear model of V^π given by Equation 23 takes a very simple form:

$$v_1\phi_1(x) + \dots + v_{96}\phi_{96}(x) + v_{97} = v_x + v_{97},$$

for parameter vector $(v_1, \dots, v_x, \dots, v_{97})^T$, where v_x is the parameter, or weight, associated with state x . For simplicity, we let $v_{97} = 0$,²⁴ and we let v_x denote the value at time step t of the parameter associated with state x . Thus, the estimate at time step t for the evaluation of state x is simply

$$V_t(x) = v_x. \quad (30)$$

Notice that whereas $V_t(x)$ denotes the estimate at time step t of the evaluation of state x , it is not necessarily the case that x is the system state at time step t .

Suppose the agent uses a fixed policy to decide how to move around the space. Assume it moves from location x at time step t to location y at step $t + 1$. If x is not the goal, then a payoff of minus one is received. Applying the TD procedure in the form it takes given the simple way we represent locations, only the parameter v_x , which is the current estimate of the evaluation of location x , is updated after the agent makes this move. This evaluation estimate, v_x , is adjusted by adding to it $\beta[-1 + \gamma v_y - v_x]$. To simplify things even more, let us assume that $\gamma = 1$ and $\beta = 1$. Then the new evaluation estimate of x , i.e., the new value of the parameter v_x , is $v_x + [-1 + v_y - v_x] = -1 + v_y$.

Given the foregoing manner of representing states and evaluation estimates of states, adjusting parameter v_x directly adjusts the evaluation estimate of state, or location, x . A feature vector can be thought of as indexing into a lookup table that lists the evaluation estimates of all the locations.²⁵ Because $\phi(x)$ contains a single one in the position corresponding to state x , and appears multiplicatively on the right-hand side of Equation 27, at each time step, the TD procedure updates only the evaluation estimate of the agent's current location. We use feature vectors defined by Equation ?? because they make the application of the TD procedure easy to understand, and they allow us to represent an arbitrary evaluation function using a linear model.

differs from the TD procedure described here in that it computes a single evaluation for each feature vector at each time step. In terms of Figure 5, the left-most "Evaluation Function Model" for each time step is missing and the result of the right-most computation is fed leftward into the summation unit. If changes in parameter values remain small, this is a reasonable approximation to the TD procedure described here, but it can become unstable if the parameter values change rapidly (see refs. [60, 61]). The Adaptive Critic Element also uses stimulus traces as in the TD model of conditioning described in refs. [64, 62].

²⁴Setting $v_{97} = 0$ in the linear model does not limit the model's ability to represent arbitrary evaluation functions of the 96 states, but it does prevent a simple form of generalization in which v_{97} might be adjusted to equal, for example, the average state evaluation.

²⁵A lookup-table representation of a function is simply a list of its values. To evaluate the function for a particular item in its domain, one simply accesses the entry in the place in the table corresponding to that item. Here, $\phi(x)$

can be thought of as accessing the place where v_x is stored.

However, this is a very special case of the range of possibilities encompassed within the theoretical framework presented here, and many interesting phenomena involving the transfer of evaluation information among the states do not appear if one uses these feature vectors.²⁶

Normally, one would begin the process of estimating the evaluation function for a given decision policy by setting the parameters to provide as good an approximation to the true evaluation function as is possible. For simplicity, however, suppose we begin by setting all the parameters to zero and then run a series of trials, where each trial begins by placing the agent in a randomly chosen location and letting it take some reasonably large number of steps. As the agent moves in the first trial, it will leave a trail of minus ones in its wake as the parameters are updated that correspond to each location visited. If it reaches the goal, the parameter determining the goal's evaluation, v_g , will remain zero. On the second trial, the same thing will happen except when the agent's path joins and moves along a path taken in a previous trial (once the agent's path intersects a path previously traversed, it will always follow that path if the decision policy is deterministic). The parameter for each location visited one step before visiting a location already assigned a -1 will be set to -2 . If the goal is reached, the location visited one step earlier will remain at -1 . On the third trial, some locations will be assigned -1 , some will be assigned -2 , as in trial two, and some will be assigned -3 if the next place visited had been assigned -2 in the second trial.

Notice what happens if, in this third trial, the agent approaches the goal on a path traveled in both preceding trials. It will reach a location, call it x , assigned an evaluation estimate of -2 , then move to a location assigned an evaluation estimate of -1 , call it y , and then reach the goal, g , which still has an evaluation estimate equal to zero. Moving from x to y generates a payoff of -1 , and the TD rule will update v_x to produce the new evaluation estimate $-1 + v_y = -1 - 1 = -2$, which is the same as its previous value. Similarly, v_y will not change when g is reached because its current evaluation estimate already equals the payoff received upon moving to the goal, -1 , plus the goal's evaluation: zero. As various paths are followed on subsequent trials, evaluation estimates stabilize in a similar fashion, from the goal outwards, to minus the number of steps required to reach the goal using the given decision policy. If the policy prevents the goal from being reached from some location, then the evaluation estimate of that location can become more negative with each trial, being decremented by the TD rule each time the location is visited. With an unbounded number of trials, evaluation estimates for such locations grow negative without bound, which is correct since the number of steps to the goal using the policy is indeed infinite, but which may be considered an undesirable form of instability. The use of a discount factor with value less than one eliminates the possibility that evaluation estimates can grow without bound.

²⁶Much of the study of parameter estimation, especially in the context of connectionist networks, is concerned with the problem of finding representations and classes of models that facilitate useful forms of transfer or extrapolation, often called generalization. See, for example, refs. [4, 23, 37, 53]. Most of these methods for representation and modeling obviously can be used to advantage with the TD procedure in the route-finding example, but the lookup-table representation allows us to address in a simple form the issues most relevant to the focus of this report. Some of the issues that arise when more complicated representations are used in sequential decision tasks are discussed by Sutton [60], and examples involving other representations are provided by Anderson's [3] use of layered networks in the pole-balancing problem and Watkin's [67] use of the representation method proposed by Albus [1] in a model of the cerebellum.

7.2 Learning an Optimal Decision Policy

Section 7.1 addressed the problem of learning the evaluation function for a fixed policy in the absence of a model of the decision task. The TD procedure is a parameter estimation method for estimating such evaluation functions. Here we consider one way of using the TD procedure to facilitate the approximation of an optimal policy, again in the absence of a model of the decision task. Although the resulting method is not guaranteed to converge to an optimal policy in all tasks, it is relatively simple and is closely related to traditional views of animal learning. Throughout the present section we assume that we do not know the quantities $R(x, a)$ and $P_{xy}(a)$, for system states x and y and actions a , which are required for the application of dynamic programming methods.

The approach described here for learning an optimal policy is related to the policy improvement method described in Section 5.2. In the policy improvement method one computes the evaluation function for some initial policy, and then defines a new policy that specifies the actions which maximize expected return as measured by that evaluation function. One then computes the evaluation function for this new policy, defines another policy so as to maximize expected return as measured by this second evaluation function, computes another evaluation function, and so on. The policies defined in this sequence keep improving until an optimal policy is obtained. The method for forming an optimal policy described below can be thought of as combining, on a step-by-step basis, the process of forming an evaluation function with the process of forming a new policy. Whereas the policy improvement method computes a *complete* evaluation function for a given policy, that is, it computes the evaluation of *all* states for that policy, and then uses these evaluations to define an improved policy, the method described below performs one step of an on-line procedure for estimating an evaluation function for a given policy, and at the same time performs one step of an on-line procedure for improving that policy. Consequently, both the current evaluation function estimate and the current policy change at each time step while the agent interacts with the system underlying the decision task. Because each policy adjustment is based on an estimate of the evaluation function for the current policy, the policy will not necessarily improve at each step as it does in the policy improvement method. However, because the evaluation estimates improve over time, the policy also tends to improve.

Two on-line procedures are therefore required for the approach to learning an optimal policy considered here. The procedure for estimating the evaluation function is the TD procedure, and the procedure for improving the decision policy is an example of a reinforcement learning procedure (see Section 3). Because our goal in this report is to describe the computational aspects of the TD procedure, we do not provide a detailed exposition of reinforcement learning. Instead, we describe the general characteristics of these methods and discuss a specific example as it applies to the route-finding task. The processes involved are diagrammed in Figure 6, to which we refer throughout this section. This figure is an expansion of Figure 5: whereas in that figure, the fixed policy is represented by the boxes labeled π , in Figure 6 the policy is adjusted by the procedure diagrammed in the top part of the figure. The bottom part of Figure 6 shows the operation of the TD procedure and is identical to Figure 5.

In Section 7.1 the TD procedure was developed as an on-line parameter estimation method for estimating an evaluation function. Similarly, the procedure for improving policies that we describe is an on-line parameter estimation method. Referring to the concept of a parameterized model described in Section 6, each policy available to the agent is assumed to be given by some parameterized model, called the policy model. Selecting specific values for the parameters, which we call the policy parameters, determines a specific policy. Policy parameters are adjusted in an

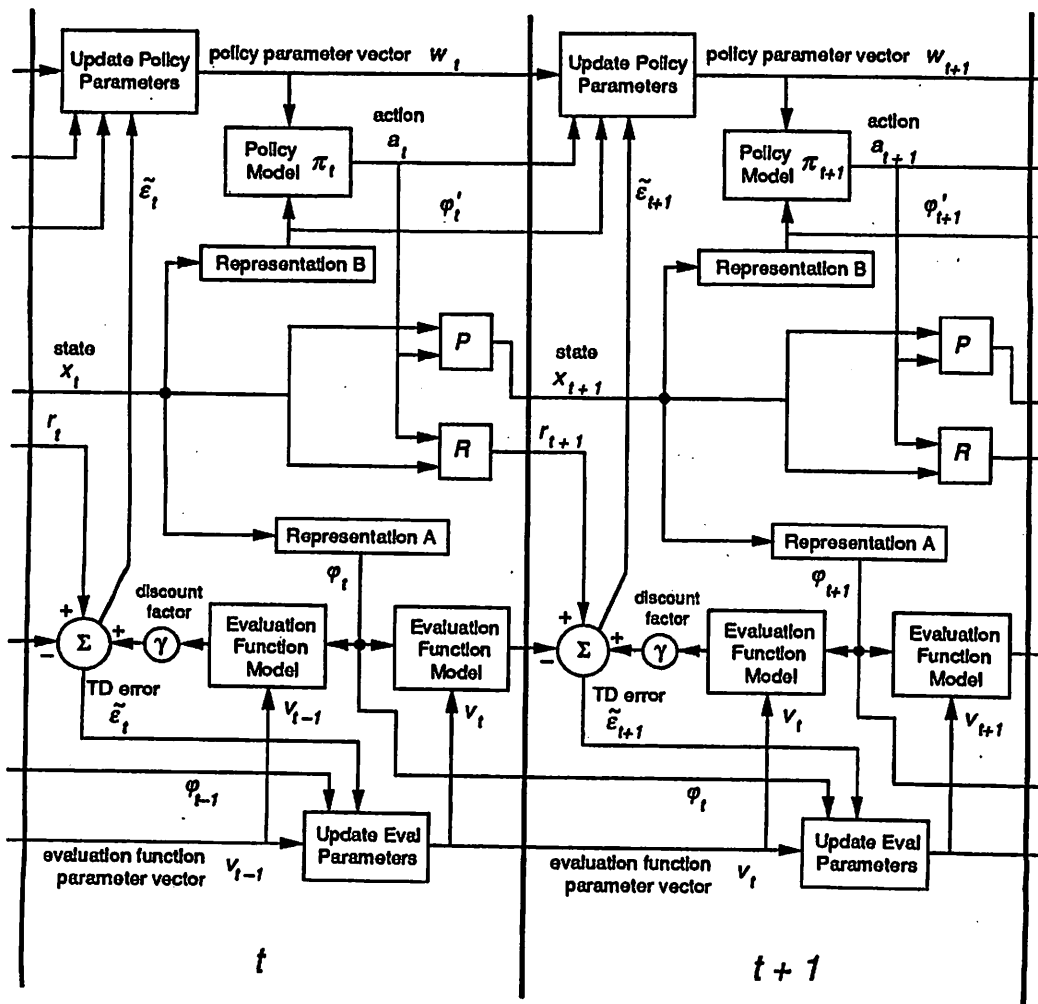


Figure 6: An elaboration of Figure 5 showing the operation of an on-line method for adjusting a decision policy. The policy is adjusted through the process diagrammed in the top part of the figure. The bottom part of the figure shows the operation of the TD procedure and is identical to Figure 5.

attempt to improve the policy they specify (accomplished by the boxes labeled "Update Policy Parameters" in Figure 6). At the same time that the policy parameters are adjusted, the TD procedure adjusts different parameters specifying the evaluation function model.

At time step t , the state of the system underlying the decision task is represented by the feature vector $\phi_t = \phi(x_t)$, which acts as input to the model of the evaluation function. In Figure 6, this representation is produced by the box labeled "Representation A" in the time- t portion of the figure. Another feature vector representing the state at step t acts as input to the model of the policy. This feature vector is denoted ϕ'_t and is shown in Figure 6 as the output of the box labeled "Representation B" in the time- t portion of the figure. Although it may be not always be necessary to use different feature vectors for the evaluation function and policy models, in general these two models do not place equal demands on state representations. However, two different parameter vectors are necessary: The parameter vector v specifies the the evaluation function, as described in Section 7.1, and the vector w specifies the policy. For simplicity, we consider only linear models for evaluation functions and decision policies.

As the agent interacts with the system, it senses the feature vectors of successive states and receives a payoff at each time step. The TD procedure (Equation 27) is applied exactly as described in Section 7.1 to update the parameter vector v . Here, however, the policy changes over time as the parameter vector w is adjusted. Let π_t denote the policy used to select the action at time step t . This policy is specified by w_t , the value of the parameter vector w at time step t for the policy model. Although at each time step, t , the TD procedure is approximating the evaluation function for policy π_t , over many time steps it is approximating the policy to which π_t is tending as t increases. Because under ideal circumstances this limiting policy should approximate an optimal policy, we regard the TD procedure used in this way as approximating the optimal evaluation function.

How does w_t specify the policy π_t , and how is w_t updated? A policy is a mapping, or function, from system states to actions that can be parameterized in many different ways. If a policy can assign to each system state one of only two possible actions, then one might represent a policy as a linear decision rule having the same form used in pattern classification (Equation 16). If more than two actions are possible, as in the case of the route-finding example, then a more complicated parameterization is required. Several possibilities for parameterizing multi-category decision rules are suggested in the pattern classification literature (e.g., refs. [16, 58]). One method, which we adapt for the route-finding example, is to use a separate parameterized function for each possible action. Each of these functions assigns a number to each state. An action is then selected by comparing the numbers produced by these functions for the current state and selecting the action whose number is largest. This is a kind of competition among the actions that can be implemented by connectionist networks with lateral inhibition (e.g., Feldman and Ballard [19]).

An additional consideration in parameterizing policies is that some policy adjustment methods require stochastic policies instead of deterministic policies. In this case, one has to parameterize functions which assign action probabilities to system states instead of specific actions. We shall describe an example of a parameterized class of multi-action stochastic policies in the context of the route-finding example. In the meantime, however, it is not misleading to think in terms of the simpler case of deterministic, two-action policies, the simplest example being policies defined by the linear threshold rule given by Expression 16.

If it were possible to know what action is desired for each state, it would be possible to adjust the policy parameters by applying a parameter estimation method, such as the LMS rule given by Equation 21, using the supervised learning paradigm. If we knew the expected values, $R(x_t, a)$,

of all the payoffs and transition probabilities, $P_{x_t,y}(a)$, for the current system state, x_t , and all possible next states y under all the actions a , then we could use the current evaluation function, V_t , to select the desired action. It would be the action, a , that maximizes

$$R(x_t, a) + \gamma \sum_{y \in X} P_{x_t,y}(a) V_t(y). \quad (31)$$

A policy selecting this action for state x_t , and otherwise following policy π_t , either improves policy π_t or leaves its performance unchanged, under the assumption that the current evaluation function estimate, V_t , is an accurate estimate of the true evaluation function for policy π_t . Unfortunately, lacking knowledge of the expected payoff and transition probabilities that appear in Expression 31, and also lacking estimates for these quantities, we cannot directly specify a desired action. Therefore, we cannot update the current policy using a rule for supervised learning, like the LMS rule, which requires a known desired response for each input vector. It is necessary to employ a method capable of learning to perform the desired actions, here the actions that maximize Expression 31, when these actions are not known by a "teacher." In order to do this, we use a reinforcement learning method instead of a supervised learning method.

The agent might overcome the unavailability of a teacher capable of providing desired responses by forming estimates of the quantities in Expression 31 while it interacts with the system underlying the decision task. This means that it would effectively have to construct a model of the decision task. However, reinforcement learning is possible without such a model: Think of the agent as beginning an instrumental conditioning trial at each time step. At time step t , it emits an action, a_t , based on a policy that shows some form of variation. At time step $t + 1$, it receives payoff r_{t+1} and estimates the evaluation of the state reached at $t + 1$ by applying the evaluation function estimate specified by the parameter vector v_t (see Equation 24). This evaluation estimate, $V_t(x_{t+1})$, which is an estimate of expected return, is added to the immediate payoff r_{t+1} , after discounting by γ , to produce a measure of the utility of emitting a_t , i.e., a measure of how useful a_t is in causing high return. This utility measure is

$$r_{t+1} + \gamma V_t(x_{t+1}). \quad (32)$$

This measure, which is identical to the sum of the first two terms of the TD error $\tilde{\epsilon}_{t+1}$ (Equation 26), can be used to alter the strength of the associative link between the feature vector ϕ'_t and a_t in order to change the probability that a_t will be performed whenever state x_t , or a state represented by a feature vector similar to ϕ'_t , is encountered in the future. The expected value of the measure given by Expression 32 is largest for the action that, if performed at time step t , would be the best action for improving policy π_t under the assumption that V_t is an accurate estimate of the evaluation function for the current policy. We need a learning rule that increases the probability of performing action a_t when the utility measure given by Expression 32 is large, and decreases this probability when it is small.

But what values of this utility measure should be considered "large" or "small"? How can a measure of the utility of an action, such as that given by Expression 32, be translated into a quantity that can be used effectively as a "reinforcement factor" to adjust action probabilities? If one had access to the expected value of the utility measure over all actions performed when the system is in a given state, it would make sense to treat measures larger than this expected value as indicating performance improvements, and measures smaller than this expected value as indicating performance decrements.²⁷ According to this approach, a reinforcement factor would be

²⁷Sutton [60] calls this the "reinforcement comparison" approach to reinforcement learning.

obtained by subtracting the expected value of the utility measure, or an estimate of this expected value, from the actual utility measure obtained for a given action. Using this reinforcement factor, a learning rule could favor actions leading to better-than-expected performance while selecting against actions leading to worse-than-expected performance.

Conveniently, within the framework described here, for each system state we already have an estimate for the expected value of the utility measure given by Expression 32 over actions performed in that state: We can use the evaluation estimate of a state given at time step t by the evaluation function estimate V_t (Equation 24). To see why the evaluation estimate of a state can serve this purpose, consider what happens to the TD error $\tilde{\epsilon}_{t+1} = r_{t+1} + \gamma V_t(x_{t+1}) - V_t(x_t)$ (Equation 26) as the TD procedure adjusts the parameter vector, v_t , defining V_t : The expected value (over system states) of $\tilde{\epsilon}_{t+1}$ should tend to zero, and we would expect $V_t(x_t)$ to approach the expected value of $r_{t+1} + \gamma V_t(x_{t+1})$, which is the utility measure of performing action a_t given by Expression 32. Consequently, a reinforcement factor can be formed by subtracting this evaluation estimate, $V_t(x_t)$, from $r_{t+1} + \gamma V_t(x_{t+1})$. The result is just the TD error itself (Equation 26), repeated here for emphasis:

$$\tilde{\epsilon}_{t+1} = [r_{t+1} + \gamma V_t(x_{t+1})] - V_t(x_t). \quad (33)$$

The TD error, $\tilde{\epsilon}_{t+1}$, therefore provides information about how good action a_t is compared to how good previous actions taken in response to ϕ'_t (see Figure 6) have been "on the average." If $\tilde{\epsilon}_{t+1}$ is positive, then a_t is better than average and should be made to occur more frequently in response to ϕ'_t in the future; if $\tilde{\epsilon}_{t+1}$ is negative, then a_t is worse than average and should be made to occur less frequently in response to ϕ'_t . It is relatively easy to use a quantity having these properties as a reinforcement factor for adjusting the policy parameters. This is accomplished by the component labeled "Update Policy Parameters" in Figure 6, which receives the TD error as one of its inputs

Several types of rules have been studied for adjusting a policy on the basis of a reinforcement factor, such as $\tilde{\epsilon}_{t+1}$, instead of direct specifications of desired actions required for supervised learning. In order for these reinforcement learning rules to work effectively, different actions must be performed in response to each system state so that so that the relative merits of performing different actions can be assessed. One way to maintain the necessary variability in activity is to let the agent "explore" by using and improving stochastic policies. It is beyond the scope of this report to describe stochastic reinforcement learning in detail, but we give an example of this approach in the context of the route-finding example.²⁸

The Route-Finding Example—A deterministic policy for this task assigns to each location one of the four actions N, S, E, or W. Although we would like to end up with a deterministic policy selecting a shortest path to the goal from each location, the learning process we describe modifies

²⁸Some of the rules for adjusting stochastic policies have been inspired by the theory of stochastic learning automata developed by cyberneticians and engineers [43, 66]. A precursor of this theory is the statistical learning theory developed by psychologists [13, 18]. Barto and Anandan [8], and Barto [6, 7] discuss a learning rule of this kind called the Associative Reward/Penalty, or A_{R-P} , rule. Sutton [60], Anderson [2, 3], and Gullapalli [21] describe the results of computer simulations of related methods. Williams [74, 75] provides theoretical analysis of a more general class of stochastic learning rules. There are other approaches to reinforcement learning that we do not address at all in this report. These include the estimation of evaluations for state-action pairs instead just states (Watkins [67]); the computation of evaluation gradients using the model of the evaluation function (Munro [42], Werbos [69, 70], Robinson and Fallside [49], and Williams [76]); and the use of systematic, instead of stochastic, variation in activity.

a stochastic policy. A stochastic policy selects actions probabilistically so that over time many actions are tried from each location. This exploratory behavior is refined as learning continues, favoring the actions appearing to yield the highest return. In other words, the stochastic policy is adjusted, by modifying the parameters specifying it, so that for each location the probability of one action—which should be the optimal action—approaches one, while the probabilities of the other actions approach zero.

Of the many different ways to define stochastic policies for the route-finding example, we used the following. The set of policy parameters, i.e, the parameters defining a stochastic policy, consists of a separate parameter for each location, x , and action, $a \in \{N, S, E, W\}$. We let $w(x, a)$ denote the value of the parameter for location x and action a . This value can be an arbitrary real number. The probability that any one of the actions is emitted when the agent visits location x is determined by the parameters corresponding to location x : $w(x, N)$, $w(x, S)$, $w(x, E)$, and $w(x, W)$. Actions are selected by means of the following probabilistic competitive process. At each time step t , four positive random numbers are drawn from an exponential distribution so that smaller values are more probable than larger values. If the agent is positioned at location x_t , a different one of these random numbers is added to each of $w(x_t, N)$, $w(x_t, S)$, $w(x_t, E)$, and $w(x_t, W)$, and the action for which this sum is the largest is performed by the agent, which moves to a new location and receives payoff r_{t+1} . In other words, the action, a_t , performed at time step t is the action a such that $w(x, a) + \eta_a$ is largest, where η_a is the random number drawn for a . The parameter values $w(x, a)$ are not altered in this selection process. Selecting actions using this competitive process has the advantage that no special precautions have to be taken to ensure that valid action probabilities are maintained throughout learning.

Now, the parameter value $w(x_t, a_t)$, corresponding the action actually selected, a_t , is updated so that the new value is

$$w(x_t, a_t) + \rho \bar{\epsilon}_{t+1}, \quad (34)$$

where x_t is the agent's location at time step t , ρ is a positive constant, and $\bar{\epsilon}_{t+1}$ is the reinforcement factor (Equation 33) corresponding to performing action a_t when the system state is x_t . Thus, if $\bar{\epsilon}_{t+1}$ is positive, the move that was taken from location x_t will be chosen with a higher probability when location x_t is visited again. If $\bar{\epsilon}_{t+1}$ is negative, a_t will be less likely to be chosen from location x_t . The parameters corresponding to the actions not selected are not changed, although the probabilities of these actions being selected are changed as a result of the change in the values $w(x_t, a_t)$ and the competitive action-selection process described above. As the parameter values, $w(x, a)$, for the various places visited and actions selected, change according to Expression 34, the decision policy changes, becoming more deterministic as specific actions become selected with higher probabilities. At the same time that the parameter values $w(x, a)$ are updated, the TD procedure (Equation 27) is applied to adjust the values of the parameters specifying the evaluation function.

Because the reinforcement factor, $\bar{\epsilon}_{t+1}$, used to adjust the policy parameters according to Expression 34 is the same as the error used by the TD procedure (Equation 27), $w(x_t, a_t)$ is updated in exactly the same way that the TD procedure updates the parameter v_{x_t} specifying the evaluation estimate of state x_t . Indeed, there is only one difference between the TD procedure and the rule used to update the policy parameters. Whereas the TD procedure updates the evaluation function parameters in a manner that is insensitive to the action chosen at time step t , the rule given by Expression 34 updates only those policy parameters corresponding to the action chosen at step t . More generally, if the decision policy is represented in a manner more complicated than the one illustrated here—for example, if actions are represented as patterns of

activity over a set of connectionist units—then the updating of policy parameters is modulated by the representation of the chosen action so as to appropriately influence the action probabilities. In Figure 6, notice that the component “Update Policy Parameters” receives actions as input, whereas the component “Update Eval Parameters” does not. Aside from this dependence of the policy-update procedure on the action performed, and possibly different values for the constants β in Equation 27 and ρ in Equation 34, the parameter-update procedures for evaluation estimation and policy improvement are exactly the same.

Figure 7 shows the estimates for the optimal evaluation function for the route-finding example after various numbers of trials in which both the TD procedure and the policy-update procedure just described were applied at each time step. A trial begins when the agent is placed at a randomly chosen location and ends when the agent reaches the goal. Each evaluation function estimate is shown as a surface whose height at each location is the estimate of the minimum number of steps to the goal from that location. Recalling that evaluation functions in this task give the *negative* of the number of steps to the goal, each surface shown in Figure 7 is actually a graph of the negation of an estimate of the optimal evaluation function. The goal is marked on each surface, and the reader should refer to Figure 1 for the position of the barrier. If a surface corresponds to the optimal evaluation function, as can be shown for the surface in Panel C, then an optimal policy is any policy selecting the actions leading most-steeply down the surface, with the exception, of course, of the actions blocked by the barrier. By the definition of the optimal evaluation function, the surface corresponding to the optimal evaluation function can have no local minimum that is not at the goal.

Surface A is the result after 50 trials. Surfaces B and C show the results after 500 and 5000 trials respectively. Examination of Surface C reveals that it is a close approximation of the (negated) optimal evaluation function for this task. Additionally, the agent’s policy improves over trials until it becomes an optimal policy with a slight degree of nondeterminism. The average over 100 trials of the number of steps required to reach the goal before learning was 2690 steps per trial. After 50 learning trials, the average over 100 non-learning trials was 41; after 500 trials, it was 9.14, and after 5000 trials, it was 6, which is near the minimum possible for this task. From these data, one can see that although many trials were required to optimize performance, rapid initial improvement occurred. This improvement would not have occurred if the agent instead spent these trials estimating only state-transition and payoff probabilities.

Although the route-finding example illustrates many of the properties of the learning methods we have described, it is too simple to illustrate others. For example, the applicability of the methods to stochastic decision tasks is not illustrated. In the example, an action performed in a given location always causes the same payoff and the same state transition. This determinism in the state-transition and payoff processes makes the application of the TD procedure and policy learning procedure easier to understand, but it does not test the ability of this learning method to estimate an optimal evaluation function and improve a decision policy in the presence of uncertainty. Sutton [61] and Watkins [67] present examples of the TD procedure applied to stochastic tasks. Another simplification present in the route-finding example is due to the particular representation used. The estimates of the optimal evaluation function and optimal policy are represented in lookup-table form, with a separate set of parameters for each system state. As mentioned in Section 7.1, this representation prevents generalization: a location has to be visited by the agent in order for it to update the estimates associated with that location.

The method described in this section for improving a policy consists of two concurrent adaptive processes: evaluation estimation and policy adjustment. The agent continually estimates an




Figure 7: Estimates of the optimal evaluation function for the route-finding example: Surface A: after 50 trials, Surface B, after 500 trials, Surface C, after 5000 trials. The surface height at each location is the estimate of the minimum number of steps to the goal from that location (and hence is the graph of the negative of the corresponding estimate of the optimal evaluation function). Because the surface in Panel C corresponds to the optimal evaluation function, an optimal policy is any policy selecting the actions leading most-steeply down this surface, with the exception, of course, of the actions blocked by the barrier (see Figure 1 for the position of the barrier).

evaluation function for its current policy, and also uses its current evaluation function to improve the policy. To improve its policy, the agent must experiment with different actions; in the course of its visits to each state, the agent must repeatedly try out all the actions possible at that state to establish which action is best. In the learning method described here, the agent's current policy is stochastic so that the agent will try a variety of actions at each state while following its policy. As time goes on, the stochastic policy is adjusted so that in each state the probabilities that less-than-optimal actions are chosen should, under ideal circumstances, decrease toward zero.

Methods such as the one described in this section involve a tradeoff between acquiring information to aid decision making in the future and using the information already acquired in the attempt to select optimal actions. Always performing actions that appear optimal based on current estimates (of evaluation functions, etc.) can prevent the experimentation required to improve these estimates. Although adjusting a stochastic policy is one way to decrease the severity of this problem, it is not easy to design methods that work well in all decision tasks. The amount of experimentation that an agent using a stochastic policy can do is limited by the degree of randomness of the policy. In the later stages of learning, the policy becomes almost deterministic, and the rate of experimentation, and hence of policy improvement, becomes very low. At intermediate stages of learning, the policy may become less-than-optimal and almost deterministic for some states. Because experimentation for these states is conducted at such a low rate, changes in the policy for these states may be very slow. Although the method of adjusting a stochastic policy while estimating an evaluation function described in this section does perform effectively in a variety of tasks, especially if learning involves a series of trials with random starting states, we know of no one who has worked out conditions under which it always converges to an optimal policy.

Watkins [67] has suggested a general approach to minimizing the difficulties caused by the estimation/control tradeoff. The idea is to separate exploration from evaluation estimation by letting the agent deviate from the policy for which it is estimating an evaluation function—its *estimation policy*—as long as it switches off its evaluation estimation procedure when such a deviation is made. The agent may do anything it likes, as long as it does not modify its estimate of the evaluation function based on information acquired when it deviates from its estimation policy. The agent is only correct to use its experience to estimate the evaluation function for the estimation policy when it is actually following that policy, but it can adjust its estimation policy all the time, whether it is following it or not. This enables an agent to experiment while maintaining the integrity of its current estimation policy and evaluation function estimate. The agent might choose experiments with more randomness than specified by its estimation policy; it might have specific innate curiosities that lead it to try certain actions in certain circumstances; or it might imitate the actions of others. Beneficial outcomes of any of these experiments can be incorporated into the estimation policy.

8 Conclusion

In this report we show how the TD procedure emerges as a combination of techniques from stochastic dynamic programming and parameter estimation. Embedding the TD procedure within these theoretical frameworks connects it with very rich traditions of theoretical research and helps explain what kinds of computational tasks it can help solve. Placing the TD procedure in this computational perspective suggests how this procedure might be applied to practical engineering problems as well as how it might help explain aspects of animal behavior.

We think the best way to relate the sequential decision framework and the TD procedure to a real-time view of classical conditioning is as follows. Think of the process of updating an evaluation function estimate as occurring continuously throughout each trial in response to continuously varying stimulus patterns and reinforcement values. The stimulus patterns provide information, possibly incomplete, about how the state of some underlying system is changing continuously over time, where the dynamics of this system are determined by the experimental procedure. The application of the TD procedure at discrete time intervals is a discrete-time simulation of a process having no explicitly demarcated steps. According to this view, the learning mechanisms that play the major role in classical conditioning construct a gradient of secondary reinforcement by facilitating connections from stimuli, and perhaps from internally stored representations, that are predictive cues for events that provide primary reinforcement. The values r_{t+1} and $\gamma V_t(x_{t+1}) - V_t(x_t)$ respectively correspond to the primary and secondary reinforcement received at time $t + 1$, with the total reinforcement being the sum of these given by the TD error, $\tilde{\epsilon}_{t+1}$. The evaluation $V_t(x_t)$ is an estimate of the expected discounted sum of future primary reinforcement available after time t . Given this interpretation, secondary reinforcement is determined by changes over time in the prediction of cumulative primary reinforcement, a view developed more completely in ref. [62].

The task of improving a decision policy within the framework developed here is similar to the task faced by an animal in instrumental conditioning. The amount of payoff the agent receives depends on its actions, and feature vectors representing the states of the system underlying the decision task act as discriminative stimuli. The method of updating a decision policy that we outlined in the context of the route-finding example is a rather straightforward implementation of an S-R view of instrumental conditioning based on the Law of Effect: Actions are selected according to their consequences in producing payoff by altering the strengths of connections between those actions and stimuli present when they were emitted. In the method we described, the consequences of actions are assessed by means of the sum of the immediate primary and secondary reinforcement. Although this role of secondary reinforcement in instrumental learning is also a standard feature of traditional views of animal learning, here we define secondary reinforcement as the state-to-state change in an estimate of an evaluation function, i.e., $\gamma V_t(x_{t+1}) - V_t(x_t)$, where the estimate is updated through time by the TD procedure.

Although learning methods similar to the one we described in Section 7.2 for improving policies have been studied by means of computer simulation (see, for example, refs. [3, 6, 8, 60]), we do not claim that the method we described always produces an optimal policy or that it is a valid model of animal behavior in instrumental conditioning experiments. However, we think it is significant that the procedures comprising this method—one of which updates an evaluation function estimate (the TD procedure), the other of which updates the decision policy—are almost identical. The same parameter-update rule is used in these procedures, but whereas it is applied uniformly to all the parameters specifying the evaluation function, it is applied to the parameters specifying the decision policy in a fashion that is modulated by a representation of the action that was emitted by the agent.

Whatever method the agent uses to adjust its decision policy, a means for approximating the evaluation function corresponding to a policy, such as the TD procedure, has obvious utility. An evaluation function provides an immediately-accessible prediction of the long-term return a policy will achieve throughout the future. This information can be used in adjusting policies so that decisions are not dominated by the short-term consequences of behavior. If the evaluation function for a given policy is available, it is possible to select actions based on their future consequences

without waiting for the future to unfold. Because it is unlikely that exact evaluation functions are ever available to a behaving animal, evaluation function estimates must serve instead. The TD procedure is a simple on-line method for forming these estimates.

We distinguished the TD procedure from model-based methods which use a model of the decision task in the form of estimates for the state-transition and payoff probabilities of the system underlying the decision task. The TD procedure is a direct method for estimating an evaluation function. It does not rely on a model of the decision task. The lack of such a model rules out many alternative strategies for making decisions based on estimates of the long-term consequences of behavior. Without a model of the decision task, it is not possible for the agent to implement conventional dynamic programming methods, such as value or policy iteration, or to perform explicit look-ahead planning of the type widely studied in artificial intelligence. Lacking a model, the agent cannot evaluate actions or action sequences without actually performing them. Model-based methods are more advanced than direct methods and can provide capabilities impossible to obtain with direct methods.

However, the additional capabilities of model-based methods should not diminish the appeal of direct methods both as computational procedures and as models of animal learning. Direct methods are not only very simple compared to model-based methods, they can be used in natural ways as components of model-based methods. For example, it is as easy to apply the TD procedure to state sequences contemplated with the aid of a task model as to state sequences that are actually experienced. Additionally, direct methods applied to actual experience cannot be misled by modeling inaccuracies in the same way that model-based methods can be. The desired consequences of executing a carefully crafted plan based on an erroneous model of reality can bear little relation to what actually happens. Although direct methods are themselves subject to many sources of inaccuracy, they do not risk the compounding of errors due to applying iterative methods to inaccurate task models. By using, as it were, the real task as its own model, direct methods maintain contact with the actual consequences of behavior.

Although we presented the TD procedure as a synthesis of ideas from stochastic dynamic programming and parameter estimation, the route-finding example used throughout this report involves only the simplest elements of these theories. Being a task with very simple deterministic state-transition and payoff structures, it does not illustrate the generality of the dynamic programming concepts. Similarly, our applications of learning methods to this task do not adequately illustrate the wide range of possibilities for representing evaluation function estimates and policies. Lookup tables are easy to understand and can be useful in solving some tasks, but the learning procedures presented here apply also when evaluation function estimates and policies are represented by other means, such as connectionist networks parameterized by connection weights. Indeed, some of the most interesting open questions concerning the learning procedures described in this report concern their use with various kinds of state representations and parameterized classes of functions. For example, the kind of generalization produced by a representation can either aid the learning process or hinder it, and methods can be studied for adaptively re-representing functions during learning (as proposed, for example, by Michie and Chambers [39] and Riordon [48]). Good features for representing policies will tend to be highly correlated with optimal actions, and good features for representing evaluation functions will tend to provide descriptions of states that are highly correlated with the optimal return available from those states.

In addition to providing a framework in which to understand the operation of the TD procedure, the formalism of sequential decision making provides a framework in which to study a

variety of other problems involving learning. Tasks requiring the control of incompletely-known dynamical systems are ubiquitous in engineering applications, they are clearly related to problems that animals routinely solve by means of learning, and they offer a rich set of challenges to designers of learning algorithms. It is our hope that the connections described in this report encourage research which brings together ideas from animal learning theory, behavioral ecology, and adaptive control engineering. As we have attempted to show for the TD procedure, computational methods inspired by biological learning are not necessarily simple special cases of existing theory.

References

- [1] J. S. Albus. Mechanisms of planning and problem solving in the brain. *Mathematical Biosciences*, 45:247–293, 1979.
- [2] C. W. Anderson. *Learning and Problem Solving with Multilayer Connectionist Systems*. PhD thesis, University of Massachusetts, Amherst, MA, 1986.
- [3] C. W. Anderson. Strategy learning with multilayer connectionist representations. In *Proceedings of the Fourth International Workshop on Machine Learning*, pages 103–114, Irvine, CA, 1987. Morgan Kaufmann.
- [4] J. A. Anderson and E. Rosenfeld, editors. *Neurocomputing: Foundations of Research*. Bradford Books/MIT Press, Cambridge, MA, 1988.
- [5] A. G. Barto. Connectionist learning for control: An overview. In T. Miller, R. S. Sutton, and P. J. Werbos, editors, *Neural Networks for Control*. The MIT Press, Cambridge, MA. To appear.
- [6] A. G. Barto. Learning by statistical cooperation of self-interested neuron-like computing elements. *Human Neurobiology*, 4:229–256, 1985.
- [7] A. G. Barto. From chemotaxis to cooperativity: Abstract exercises in neuronal learning strategies. In R. Durbin, R. Maill, and G. Mitchison, editors, *The Computing Neuron*. Addison-Wesley, 1989.
- [8] A. G. Barto and P. Anandan. Pattern recognizing stochastic learning automata. *IEEE Transactions on Systems, Man, and Cybernetics*, 15:360–375, 1985.
- [9] A. G. Barto, R. S. Sutton, and C. W. Anderson. Neuronlike elements that can solve difficult learning control problems. *IEEE Transactions on Systems, Man, and Cybernetics*, 13:835–846, 1983. Reprinted in J. A. Anderson and E. Rosenfeld, *Neurocomputing: Foundations of Research*, The MIT Press, Cambridge, MA, 1988.
- [10] R. E. Bellman. *Dynamic Programming*. Princeton University Press, Princeton, NJ, 1957.
- [11] L. B. Booker. *Intelligent Behavior as an Adaptation to the Task Environment*. PhD thesis, University of Michigan, Ann Arbor, MI, 1982.
- [12] V. Borkar and P. Varaiya. Adaptive control of markov chains I: Finite parameter set. *IEEE Transactions on Automatic Control*, 24:953–957, 1979.

- [13] R. R. Bush and F. Mosteller. *Stochastic Models for Learning*. Wiley, New York, 1955.
- [14] A. Dickinson. *Contemporary Animal Learning Theory*. Cambridge University Press, Cambridge, 1980.
- [15] S. E. Dreyfus and A. M. Law. *The Art and Theory of Dynamic Programming*. Academic Press, New York, 1977.
- [16] R. O. Duda and P. E. Hart. *Pattern Classification and Scene Analysis*. Wiley, New York, 1973.
- [17] Y. El-Fattah. Recursive algorithms for adaptive control of finite markov chains. *IEEE Transactions on Systems, Man, and Cybernetics*, 11:135-144, 1981.
- [18] W. K. Estes. Toward a statistical theory of learning. *Psychological Review*, 57:94-107, 1950.
- [19] J. A. Feldman and D. H. Ballard. Connectionist models and their properties. *Cognitive Science*, 6:205-254, 1982.
- [20] G. C. Goodwin and K. S. Sin. *Adaptive Filtering Prediction and Control*. Prentice-Hall, Inc., Englewood Cliffs, N.J., 1984.
- [21] V. Gullapalli. A stochastic algorithm for learning real-valued functions via reinforcement feedback. Technical Report 88-91, University of Massachusetts, Amherst, MA, 1988.
- [22] S. E. Hampson. *A Neural Model of Adaptive Behavior*. PhD thesis, University of California, Irvine, CA, 1983.
- [23] G. E. Hinton and J. A. Anderson, editors. *Parallel Models of Associative Memory*. Erlbaum, Hillsdale, NJ, 1981.
- [24] J. H. Holland. Escaping brittleness: The possibility of general-purpose learning algorithms applied to rule-based systems. In R. S. Michalski, J. G. Carbonell, and T. M. Mitchell, editors, *Machine Learning: An Artificial Intelligence Approach, Volume II*. Morgan Kaufmann, Los Altos, CA, 1986.
- [25] A. Houston, C. Clark, J. McNamara, and M. Mangel. Dynamic models in behavioral and evolutionary ecology. *Nature*, 332:29-34, March 1988.
- [26] R. Howard. *Dynamic Programming and Markov Processes*. The MIT Press, Cambridge, MA, 1960.
- [27] A. H. Klopff. Brain function and adaptive systems—A heterostatic theory. Technical Report AFCRL-72-0164, Air Force Cambridge Research Laboratories, Bedford, MA, 1972. A summary appears in *Proceedings of the International Conference on Systems, Man, and Cybernetics*, 1974, IEEE Systems, Man, and Cybernetics Society, Dallas, TX.
- [28] A. H. Klopff. *The Hedonistic Neuron: A Theory of Memory, Learning, and Intelligence*. Hemisphere, Washington, D.C., 1982.
- [29] J. R. Krebs, A. Kacelnik, and P. Taylor. Test of optimal sampling by foraging great tits. *Nature*, 275:27-31, September 1978.

- [30] P. R. Kumar and Woei Lin. Optimal adaptive controllers for unknown markov chains. *IEEE Transactions on Automatic Control*, 25, 1982.
- [31] S. Lakshminvarahan. *Learning Algorithms and Applications*. Springer-Verlag, New York, 1981.
- [32] G. E. Liepins, M. R. Hilliard, and M. Palmer. Credit assignment and discovery in classifier systems. In preparation.
- [33] L. Ljung and T. Söderstrom. *Theory and Practice of Recursive Identification*. The MIT Press, Cambridge, MA, 1983.
- [34] L. M. Lyubchik and A. S. Poznyak. Learning automata in stochastic plant control problems. *Automatic and Remote Control (USSR)*, 35:777-789, 1974.
- [35] P. Mandl. Estimation and control in markov chains. *Advances in Applied Probability*, 6:40-60, 1974.
- [36] M. Mangel and C. W. Clark. *Dynamic Modeling in Behavioral Ecology*. Princeton University Press, Princeton, NJ, 1988.
- [37] J. L. McClelland and D. E. Rumelhart, editors. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, vol.2: Applications*. Bradford Books/MIT Press, Cambridge, MA, 1986.
- [38] J. M. Mendel and R. W. McLaren. Reinforcement learning control and pattern recognition systems. In J. M. Mendel and K. S. Fu, editors, *Adaptive, Learning and Pattern Recognition Systems: Theory and Applications*, pages 287-318. Academic Press, New York, 1970.
- [39] D. Michie and R. A. Chambers. BOXES: An experiment in adaptive control. In E. Dale and D. Michie, editors, *Machine Intelligence 2*, pages 137-152. Oliver and Boyd, 1968.
- [40] M. L. Minsky. *Theory of Neural-Analog Reinforcement Systems and its Application to the Brain-Model Problem*. PhD thesis, Princeton University, 1954.
- [41] M. L. Minsky. Steps toward artificial intelligence. *Proceedings of the Institute of Radio Engineers*, 49:8-30, 1961. Reprinted in E. A. Feigenbaum and J. Feldman, editors, *Computers and Thought*. McGraw-Hill, New York, 406-450, 1963.
- [42] P. Munro. A dual back-propagation scheme for scalar reward learning. In *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*, pages 165-176, Seattle, WA, 1987.
- [43] K. Narendra and M. A. L. Thathachar. *Learning Automata: An Introduction*. Prentice Hall, Englewood Cliffs, NJ, 1989.
- [44] J. Pearl. *Heuristics: Intelligent Search Strategies for Computer Problem Solving*. Addison-Wesley, 1984.
- [45] H. Rachlin, R. Battalio, J. Kagel, and L. Green. Maximization theory in behavioral psychology. *The Behavioral and Brain Sciences*, 4:371-417, 1981.
- [46] R. A. Rescorla. A pavlovian analysis of goal-directed behavior. *American Psychologist*, 42:119-129, 1987.

- [47] R. A. Rescorla and A. R. Wagner. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black and W. F. Prokasy, editors, *Classical Conditioning II*, pages 64–99. Appleton-Century-Crofts, New York, 1972.
- [48] J. S. Riordon. An adaptive automaton controller for discrete-time markov processes. *Automatica*, 5:721–730, 1969.
- [49] A. J. Robinson and F. Fallside. The utility driven dynamic error propagation network. Technical Report CUED/F-INFENG/TR.1, Cambridge University Engineering Department, Cambridge, England, 1987.
- [50] F. Rosenblatt. *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*. Spartan Books, 6411 Chillum Place N.W., Washington, D.C., 1961.
- [51] S. Ross. *Introduction to Stochastic Dynamic Programming*. Academic Press, New York, 1983.
- [52] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning internal representations by error propagation. In D. E. Rumelhart and J. L. McClelland, editors, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, vol.1: Foundations*. Bradford Books/MIT Press, Cambridge, MA, 1986.
- [53] D. E. Rumelhart and J. L. McClelland, editors. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, vol.1: Foundations*. Bradford Books/MIT Press, Cambridge, MA, 1986.
- [54] A. L. Samuel. Some studies in machine learning using the game of checkers. *IBM Journal on Research and Development*, pages 210–229, 1959. Reprinted in E. A. Feigenbaum and J. Feldman, editors, *Computers and Thought*, McGraw-Hill, New York, 1963.
- [55] A. L. Samuel. Some studies in machine learning using the game of checkers. II—Recent progress. *IBM Journal on Research and Development*, pages 601–617, November 1967.
- [56] M. Sato, K. Abe, and H. Takeda. Learning control of finite markov chains with unknown transition probabilities. *IEEE Transactions on Automatic Control*, 27, 1982.
- [57] O. Selfridge, R. S. Sutton, and A. G. Barto. Training and tracking in robotics. In *Proceedings of the Ninth International Joint Conference of Artificial Intelligence*, Los Angeles, CA, August 1985.
- [58] J. Sklansky and G. N. Wassel. *Pattern Classifiers and Trainable Machines*. Springer-Verlag, New York, 1981.
- [59] J. E. R. Staddon. Optimality analyses of operant behavior and their relation to optimal foraging. In J. E. R. Staddon, editor, *Limits to Action: The Allocation of Individual Behavior*, pages 101–141. Academic Press, New York, 1980.
- [60] R. S. Sutton. *Temporal Credit Assignment in Reinforcement Learning*. PhD thesis, University of Massachusetts, Amherst, MA, 1984.
- [61] R. S. Sutton. Learning to predict by the methods of temporal differences. *Machine Learning*, 3:9–44, 1988.

- [62] R. S. Sutton and A. G. Barto. Time-derivative models of pavlovian conditioning. In M. Gabriel and J. W. Moore, editors, *Learning and Computational Neuroscience*. The MIT Press, Cambridge, MA. To appear.
- [63] R. S. Sutton and A. G. Barto. Toward a modern theory of adaptive networks: Expectation and prediction. *Psychological Review*, 88:135-171, 1981.
- [64] R. S. Sutton and A. G. Barto. A temporal-difference model of classical conditioning. In *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*, Seattle, WA, 1987.
- [65] E. L. Thorndike. *Animal Intelligence*. Hafner, Darien, Conn., 1911.
- [66] M. L. Tsetlin. *Automaton Theory and Modeling of Biological Systems*. Academic Press, New York, 1973.
- [67] C. J. C. H. Watkins. *Learning from Delayed Rewards*. PhD thesis, Cambridge University, Cambridge, England, 1989. Pending.
- [68] P. J. Werbos. Advanced forecasting methods for global crisis warning and models of intelligence. *General Systems Yearbook*, 22:25-38, 1977.
- [69] P. J. Werbos. Building and understanding adaptive systems: A statistical/numerical approach to factory automation and brain research. *IEEE Transactions on Systems, Man, and Cybernetics*, 1987.
- [70] P. J. Werbos. Generalization of back propagation with applications to a recurrent gas market model. *Neural Networks*, 1, 1988.
- [71] R. M. Wheeler and K. S. Narendra. Decentralized learning in finite markov chains. *IEEE Transactions on Automatic Control*, 31:519-526, 1986.
- [72] B. Widrow and M. E. Hoff. Adaptive switching circuits. In *1960 WESCON Convention Record Part IV*, pages 96-104, 1960. Reprinted in J. A. Anderson and E. Rosenfeld, *Neuro-computing: Foundations of Research*, The MIT Press, Cambridge, MA, 1988.
- [73] B. Widrow and S. D. Stearns. *Adaptive Signal Processing*. Prentice-Hall, Inc., Englewood Cliffs, N.J., 1985.
- [74] R. J. Williams. Reinforcement learning in connectionist networks: A mathematical analysis. Technical Report ICS 8605, Institute for Cognitive Science, University of California at San Diego, La Jolla, CA, 1986.
- [75] R. J. Williams. Reinforcement-learning connectionist systems. Technical Report NU-CCS-87-3, College of Computer Science, Northeastern University, 360 Huntington Avenue, Boston, MA, 1987.
- [76] R. J. Williams. On the use of backpropagation in associative reinforcement learning. In *Proceedings of the IEEE International Conference on Neural Networks*, San Diego, 1988.
- [77] I. H. Witten. An adaptive optimal controller for discrete-time Markov environments. *Information and Control*, pages 286-295, 1977.