

An Empirical Evaluation of Interval Estimation for Markov Decision Processes

Alexander L. Strehl
Department of Computer Science
Rutgers University
Piscataway, NJ USA
strehl@cs.rutgers.edu

Michael L. Littman
Department of Computer Science
Rutgers University
Piscataway, NJ USA
mlittman@cs.rutgers.edu

Abstract

This paper takes an empirical approach to evaluating three model-based reinforcement-learning methods. All methods intend to speed the learning process by mixing exploitation of learned knowledge with exploration of possibly promising alternatives. We consider ϵ -greedy exploration, which is computationally cheap and popular, but unfocused in its exploration effort; R-Max exploration, a simplification of an exploration scheme that comes with a theoretical guarantee of efficiency; and a well-grounded approach, model-based interval estimation, that better integrates exploration and exploitation. Our experiments indicate that effective exploration can result in dramatic improvements in the observed rate of learning.

1. Introduction

One of the distinguishing features of reinforcement learning is the control the learner has over its own inputs. During learning, the decision maker must often choose between decisions that appear to be best at achieving high reward and those that may lead it to parts of the state space where even higher reward can be achieved. Choosing wisely can dramatically decrease the time needed to learn to behave effectively in a new environment.

In this work, we evaluate agents following three exploration schemes, ϵ -greedy, R-Max, and model-based interval exploration (MBIE), and show how each can trade off exploration to identify promising parts of the state space with exploitation to increase the rate the agent gathers reward. Each algorithm includes a parameter that controls the exploration–exploitation tradeoff and we measure cumulative reward on a set of simple Markov decision processes for

a variety of parameter settings. All show a distinctive “inverted u” behavior—for low parameter values, algorithms quickly converge to suboptimal behavior, while for high parameter values, algorithms converge much more slowly, but to more efficient policies. In between, good policies are found relatively quickly, thus maximizing cumulative reward over the period of experimentation.

The idea of model-based interval estimation (MBIE) was first introduced by Wiering (1999). It is inspired by interval estimation (Kaelbling, 1993), shown to produce probably approximately optimal policies in bandit problems (Fong, 1995). In our formulation, it builds a confidence interval on the parameters of a Markov decision process model of the environment, then finds the combination of parameters and agent behavior that leads to an upperbound on the maximum estimated reward, subject to the constraint that parameters lie in the confidence interval. In contrast to the original MBIE approach, our algorithm is simpler, comes with a formal convergence guarantee, and uses a more accurate calculation for the confidence intervals.

Section 2 introduces the exploration–exploitation tradeoff, discusses bandit problems and Markov decision processes, and presents several basic exploration algorithms. Our MBIE approach is presented in Section 4 and experimental comparisons in Section 5.

2. Approaches to Exploration

This section presents several approaches to exploration that have been the target of research in bandit problems and more complex Markov decision processes (MDPs). The MDP approaches are all model based.

2.1. Bandit Problems

Bandit problems (Berry & Fristedt, 1985) provide a minimal mathematical model in which to study the exploration–exploitation tradeoff. Here, we imagine we are faced Each

arm i has a stationary payoff distribution with finite mean μ_i and variance. With a slot-machine with a set of k arms, only one of which can be selected on each timestep.

The optimal long-term strategy for a k -armed bandit is to always select the arm with the highest mean payoff, $\arg\max_i \mu_i$. However, these expected payoffs are not known to the agent, who must expend a number of trials in exploration. By the law of large numbers, increasingly accurate estimates of the mean payoff of each arm can be made by repeated sampling. However, an agent that continues to sample all arms indefinitely will be wasting opportunities to exploit what it has learned to maximize its payoffs.

While a number of positive results are available for optimizing discounted reward (Gittins, 1989), the pursuit of simpler and more general algorithms has led researchers to other objective criteria. One that we would like to highlight is the design of algorithms that are *probably approximately optimal* (PAO). After a relatively small amount of experience (typically polynomial in problem parameters such as k , $1/\epsilon$, $1/\delta$, and the maximum possible reward R_{\max}), a PAO algorithm very likely produces a near-optimal strategy (no more than ϵ less than optimal with probability $1 - \delta$).

Consider the following simple strategies:

- ϵ -greedy: With probability ϵ , choose an arm at random (explore), otherwise greedily pick the arm with the highest estimated payoff (exploit). To insure that all arms receive at least one pull, arms that have not yet been tried are given an estimate of R_{\max} . A closely related approach, *Boltzmann exploration*, weights the probabilities according to how close actions are to being the best.
- “naïve”: Choose each arm a fixed number of times, C , then choose the action with the highest estimated payoff thereafter. Equivalently, estimate the payoff for any arm that has not yet been chosen C times as R_{\max} and always choose greedily with respect to this estimate.
- IE (interval estimation): Construct a $p\%$ confidence interval for the mean of the payoff distribution for each arm. At each step, choose the action with the highest upper confidence interval. Equivalently, estimate that each arm has a payoff equal to the estimated mean plus its $p\%$ confidence interval and always choose greedily with respect to this estimate.

Fong (1995) analyzed “naïve” and IE and showed that both are PAO for appropriate choices of the constants C and p . Building on these results, we can show that a choice of ϵ exists so that ϵ -greedy is PAO as well.

Proposition 1 *The ϵ -greedy exploration scheme is PAO in k -armed bandit problems.*

Proof sketch: We seek a strategy that, after a polynomial number of interactions with probability at least $1 - \delta$, executes a policy that is within ϵ of optimal. The results of Fong (1995) can be used to show that there is a number C , polynomial in magnitude, such that once each arm has been chosen C times, always choosing the arm with highest estimated payoff will be $\epsilon/2$ -optimal with probability $1 - \delta$.

In ϵ -greedy exploration, every arm has a chance of at least ϵ/k of being selected on any trial. Thus, after $\log(\delta/(ck))/\log(1 - \epsilon/k)$ trials, we are $1 - \delta/(ck)$ certain that an arm has been selected at least once. So, after $ck(\log(\delta/(ck))/\log(1 - \epsilon/k))$ trials, we are $1 - \delta$ certain that every arm has been selected at least C times. Let $\epsilon = \epsilon/(R_{\max} - R_{\min})$, where R_{\max} is the highest mean payoff and R_{\min} is the lowest mean payoff. Then, after a polynomial amount of experience, the ϵ -greedy policy is ϵ optimal. \square

Although all three approaches are PAO, in empirical studies, IE is often able to rule out actions more quickly, resulting in faster learning times. In the next section, we transfer these results to Markov decision processes.

3. Exploration in MDPs

A Markov Decision Process (MDP) is a four tuple $\langle S, A, T, R \rangle$, where S is the state space, A is the action space, $T : S \times A \times S \rightarrow \mathbb{R}$ is a transition function, and $R : S \times A \rightarrow \mathbb{R}$ is a reward function. The function $T(s, a, s')$ is interpreted as the probability of ending in state s' by taking action a from state s , and $R(s, a)$ is the expected reward for taking action a from state s . Throughout this paper, we assume that the state and action spaces are finite. Furthermore, we assume that a discount factor $0 \leq \gamma < 1$ and a vector of start-state probabilities are associated with each MDP. The discount factor defines the performance metric.

If the dynamics of the MDP are known, a policy can be found mapping states to actions that maximizes the expected discounted reward by solving the Bellman equations,

$$Q(s, a) = R(s, a) + \gamma \sum_{s'} T(s, a, s') \max_{a'} Q(s', a') \quad (1)$$

and selecting action $\arg\max_a Q(s_t, a)$ at time t . The equations can be solved by *value iteration*, which iteratively applies the equations as an update to an arbitrary initial guess of the values (Puterman, 1994; Sutton & Barto, 1998).

Note that MDPs generalize k -armed bandit problems. Any k -armed bandit problem can be thought of as an MDP with one state and k actions. The rewards for performing each action are the expected value of the rewards for pulling each arm. MDPs can be more complex than k -armed bandit problems because an action not only yields a reward, but also determines the next state (Duff & Barto, 1997).

The model-based reinforcement-learning algorithms we studied parallel the bandit algorithms introduced in the previous section and include a number of common elements. Let $n(s, a, s')$ be the number of times that the agent has ended in state s' after taking action a from state s during a run, and define $n(s, a) = \sum_{s'} n(s, a, s')$. The maximum likelihood model for the transition function T is $\hat{T}(s, a, s') = n(s, a, s')/n(s, a)$. Similarly, we can estimate $\hat{R}(s, a) = r(s, a)/n(s, a)$, where $r(s, a)$ is the total reward received by the agent from taking action a in state s . In our study, we use “optimistic initialization” (Moore & Atkeson, 1993) in that $\hat{R}(s, a) = R_{\max}$ if $n(s, a) = 0$. We also use $\hat{T}(s, a, s) = 1$ if $n(s, a) = 0$. A greedy policy with respect to the learned “certainty equivalent” model (Dayan & Sejnowski, 1996) can be computed by solving Equation 1 with $T = \hat{T}$ and $R = \hat{R}$.

3.1. Epsilon-Greedy

The ϵ -greedy approach can be used in MDPs by computing the greedy policy with respect to the maximum likelihood model. When it is time to act, with probability $1 - \epsilon$, the optimal action with respect to the estimated transition model is taken, and with probability ϵ , a random action is chosen.

Although ϵ -greedy exploration ensures convergence to an optimal policy in the limit, it displays a number of weaknesses. For one thing, the agent has no control over when or how to explore. As such, the agent is just as likely to take greedy actions before it has much reason to believe that these actions are appropriate as it would to take a random action from a state in which the best choice is crystal clear. A more subtle weakness can be seen in this example. Suppose we have an MDP in which Action 1 in State 1 yields a large reward and returns the agent back to State 1, while Action 2 yields no reward but frequently sends the agent to State 2. From State 1, a reasonable course of action is to take Action 2, giving up the possible reward for Action 1, in hopes that the agent will visit State 2 to commence exploration of that unknown state. This is an example of directed exploration, a behavior not present in the ϵ -greedy algorithm (apart from the brief effect of optimistic initialization). Because of these weaknesses, ϵ -greedy is not PAO—it can require exponentially many steps to find near-optimal behavior (Koenig & Simmons, 1993), for example in “combination lock” environments.

3.2. R-Max

Kearns and Singh (2002) showed that a PAO condition for MDPs can be defined and they provided an algorithm (E^3) that provably achieves this condition. Brafman and

Tennenholtz (2002) made use of a similar analysis to show that a simpler algorithm, R-Max, is PAO.

R-Max uses a different estimate of the transition model than does the ϵ -greedy approach. Following the “naïve” algorithm described earlier, R-Max can be parameterized by a count C , with an optimistic estimate used whenever fewer than C trials of experience are available. Specifically, R-Max modifies the Bellman equations to use Equation 1 only in states for which $n(s, a) \geq c$ and for the other states to use $Q(s, a) = R_{\max}/(1 - \gamma)$, an upper bound on the maximum achievable value.

In the R-Max work, the value of the parameter C is set as a function of a number of other parameters, including the number of states, R_{\max} , and the desired certainty of near optimality. In this paper, however, we take C to be a free variable, typically much lower than its formal bound.

Using R-Max, an agent will always choose the action that is optimal given its computed $Q(s, a)$ values. The optimistic values for state–action pairs visited fewer than C times receive a kind of “exploration bonus” and, therefore, when it is time to act, the agent will often be encouraged to visit such state–action pairs. Compared to ϵ -greedy exploration, R-Max will explore an unknown action multiple times before using its empirical distribution, and uses no randomness in its action choices. It also provides a formal guarantee of efficient learning.

3.3. Model-based Interval Estimation

R-Max can be viewed as a generalization of the “naïve” algorithm from bandit problems to MDPs in that both require a predefined number of trials before an empirical estimate is used. IE was generalized to MDPs by Kaelbling (1993) in the context of Q-learning. During learning, the value of each state–action pair is represented as a distribution, because of noise in the learning process. As in the bandit case, IE uses the upper confidence interval on the mean of this distribution as its value when making updates.

The core type of uncertainty encountered when solving an MDP by reinforcement learning, however, is uncertainty on the model parameters, not the values. Model-based interval-estimation (MBIE) uses formal confidence intervals on the *probability distributions* for each state–action pair. As in the ϵ -greedy and R-Max approaches, the algorithm keeps maximum likelihood estimates \hat{T} and \hat{R} of the transition and reward functions. Like R-Max, it provides an exploration bonus for insufficiently explored state–action pairs. The form of the exploration bonus is different, however, as it is based on the likelihood with which better outcomes than those already experienced are possible. The form of the algorithm, its motivation, and an efficient implementation are explained next. In our extended paper (Strehl & Littman, 2004), we present a proof that MBIE satisfies a

PAO-style condition, achieving near-optimal reward over a polynomial number of trials.

4. Computing Policies in MBIE

Interval estimation for bandit problems estimates the mean for each arm. Based on a finite sample, it cannot be sure of the true mean for any arm, so it constructs a confidence interval on the mean. Of all possible means consistent with these confidence intervals, IE assumes that the most optimistic one is true. It then chooses the arm that maximizes its expected payoff under this assumption—it picks the arm with the highest upper confidence interval.

This scheme for optimization under uncertainty has several beneficial properties. First of all, it is easy to compute. Second, by choosing the arm with the highest upper confidence interval, one of two possibilities occurs. Either the estimate is accurate and the decision maker receives a near-optimal reward (in expectation), or the estimate is inaccurate and the decision maker gets new experience that can be used to create a more accurate model.

Translated to the transition probabilities in the learned MDP model, the interval estimation idea suggests that, for each state–action pair, we find the probability distribution $\tilde{T}(s, a, \cdot)$ within a confidence interval of the empirically derived distribution $\hat{T}(s, a, \cdot)$ that leads to the policy with the largest reward. This can be defined with the equations

$$Q(s, a) = \hat{R}(s, a) + \max_{\tilde{T}(s, a, \cdot) \in CI} \gamma \sum_{s'} \tilde{T}(s, a, s') \max_{a'} Q(s', a'), \quad (2)$$

where CI represents the set of probability distributions that are within a $(1 - \delta)$ confidence interval of the observed distribution $\hat{T}(s, a, \cdot)$ given a sample size of $n(s, a)$. The next section provides a method for describing such a set. This is followed by a method for solving Equation 2 efficiently.

4.1. Confidence Intervals

Building on standard results, Weissman et al. (2003) showed that

$$\Pr(\|\hat{T}(s, a, \cdot) - \tilde{T}(s, a, \cdot)\|_1 \geq \varepsilon) \leq (n(s, a) + 1)^{k(s, a) - 1} e^{-n(s, a)\varepsilon^2/2},$$

where $k(s, a) = |\{s_{\max}\} \cup \{s' | \hat{T}(s, a, s') > 0\}|^1$, that is, the number of states observed to follow s under action a , along with a hypothetical state s_{\max} , and $\|x(\cdot)\|_1 = \sum_i |x(i)|$ is the L_1 norm. Manipulating this expression, we get a definition of the confidence interval set

$$CI = \{\tilde{T}(s, a, \cdot) \mid \|\tilde{T}(s, a, \cdot) - \hat{T}(s, a, \cdot)\|_1 \leq$$

¹ Erratum: It should be $k(s, a) = |S| + 1$. Due to our search for the optimal setting of parameter D , the empirical results are still valid.

$$\sqrt{\frac{2((k(s, a) - 1) \ln(n(s, a) + 1) - \ln(\delta))}{n(s, a)}} \}$$

In our experiments, however, we used a slightly weaker form of this expression:

$$CI = \{\tilde{T}(s, a, \cdot) \mid \|\tilde{T}(s, a, \cdot) - \hat{T}(s, a, \cdot)\|_1 \leq D \sqrt{k(s, a) \ln(n(s, a) + 1) / n(s, a)}\},$$

with the free parameter D , again controlling the tradeoff between exploration (small D) and exploitation (large D). The simplified expression has the property that D can be set as a function of ϵ and δ to satisfy (upper bound) the expression used to determine membership of CI .

4.2. Value Iteration with Confidence Intervals

We now argue that Equation 2 can be solved using value iteration, that is, by starting with an arbitrary estimate of $Q(s, a)$ and then repeatedly using Equation 2 as an assignment statement. This follows from the fact that Equation 2 leads to a contraction mapping.

Proposition 2 *Let $Q(s, a)$ and $W(s, a)$ be value functions and $Q'(s, a)$ and $W'(s, a)$ be the result of applying Equation 2 to Q and W , respectively. The update results in a contraction mapping in max-norm, that is, $\max_{s, a} |Q'(s, a) - W'(s, a)| \leq \gamma \max_{s, a} |Q(s, a) - W(s, a)|$.*

Proof sketch: The argument is nearly identical to the standard convergence proof of value iteration (Puterman, 1994), noting that the maximization over the compact set of probability distributions does not interfere with the contraction property. \square

Note that convergence is linear in the discount factor, so a good approximation can be found quickly, or an exact solution in polynomial time assuming a fixed discount factor (Givan et al., 2000).

We seek a probability distribution $\tilde{T}(s, a, \cdot)$ that leads to a high confidence upper bound on the value $Q(s, a)$. To ensure that our estimate is sufficiently large, we assert the existence of a state, s_{\max} , yielding maximum reward and a self loop. Although a transition to s_{\max} is never observed, it is not until the algorithm has seen a sufficient number of transitions from (s, a) that such a transition can be ruled out.

4.3. Solving the CI Constraint

To implement value iteration based on Equation 2, we need to solve the following computational problem. Maximize

$$M_V = \sum_{s'} \tilde{T}(s') V(s'),$$

where $\tilde{T}(\cdot)$ is a probability distribution subject to the constraint $\|\tilde{T}(\cdot) - \hat{T}(s, a, \cdot)\|_1 \leq \varepsilon$, and

$\varepsilon = D\sqrt{k(s,a)\ln(n(s,a)+1)/n(s,a)}$ as in Equation 3.

The following procedure can be used to find the maximum value. First, note that if $\max_{s'} \hat{T}(s, a, s') \leq \varepsilon/2$, then $M_V = \max_{s'} V(s')$. This is because the L_1 distance from $\hat{T}(s, a, \cdot)$ to one that puts all its weight on the maximum value of $V(\cdot)$ is less than ε . Otherwise, let $\tilde{T}(s') = \hat{T}(s, a, \cdot)$ to start. Let $s^* = \operatorname{argmax}_s V(s)$ and increment $\tilde{T}(s^*)$ by $\varepsilon/2$. At this point, $\tilde{T}(\cdot)$ is not a probability distribution, since it sums to $1 + \varepsilon/2$; however, its L_1 distance to $\hat{T}(s, a, \cdot)$ is $\varepsilon/2$. We need to remove exactly $\varepsilon/2$ weight from $\tilde{T}(\cdot)$. The weight is removed iteratively, taking the maximum amount possible from the state $s^- = \operatorname{argmin}_{s|\tilde{T}(s)>0} V(s)$, since this decreases M_V the least.

Proposition 3 *The procedure just described maximizes M_V .*

Proof sketch: First, the weight on s^* must be $\hat{T}(s, a, s^*) + \varepsilon/2$. If it's larger, $\tilde{T}(\cdot)$ violates either the sum to 1 constraint or the L_1 constraint. Let $\tilde{T}(s') - \hat{T}(s, a, s')$ be called the *residual* of state s' . If $\|\tilde{T}(\cdot) - \hat{T}(s, a, \cdot)\|_1 = \varepsilon$, then the sum of the positive residuals is $\varepsilon/2$ and the sum of the negative residuals is $-\varepsilon/2$. For two states s_1 and s_2 , if $V(s_1) > V(s_2)$, then M_V is increased by moving positive residual from s_2 to s_1 or negative residual from s_1 to s_2 . Therefore, putting all positive residual on s^* and moving all the negative residual toward states with smallest $V(s)$ values maximizes M_V among all distributions $\tilde{T}(\cdot)$ with the given L_1 constraint. \square

Givan et al. (2000) derived analogous results for a closely related problem where constraints are expressed in a weighted L_∞ norm. Both algorithms can be implemented to run in nearly linear time in the number of observed next states.

Combining the various results from this section together, we can efficiently find a policy that maximizes the expected reward in an MDP, assuming that transition probabilities are optimistic but within a fixed confidence interval of the observed transitions. An agent following the MBIE exploration algorithm chooses actions according to this policy to jointly explore the environment and exploit its experience. Like R-Max, it also provides a guarantee of efficient learning (Strehl & Littman, 2004).

5. Experimental Results

Many approaches to exploration in MDPs have been proposed (Thrun, 1992; Wiering & Schmidhuber, 1998; Darden et al., 1999; Meuleau & Bourgin, 1999; Wyatt, 2001). In this section, we present a comparison of MBIE to two other representative schemes for which their learning effi-

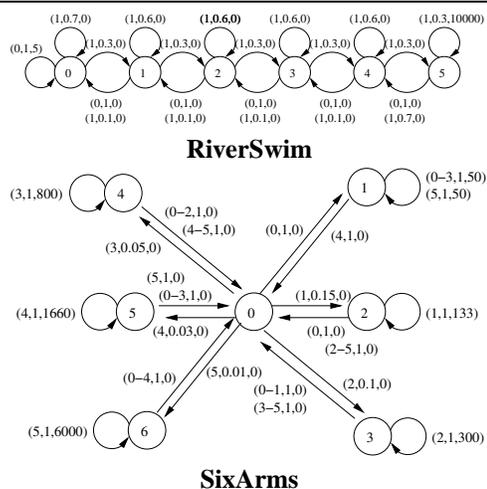


Figure 1. Two example environments: RiverSwim (top), and SixArms (bottom). Each node in the graph is a state and each edge is labeled by one or more transitions. A transition is of the form (a, p, r) where a is an action, p the probability that action will result in the corresponding transition, and r is the reward for taking the transition. For SixArms, the agent starts in State 0. For RiverSwim, the agent starts in either State 1 or 2 with equal probability.

ciency is known, R-Max and ϵ -Greedy, in two simple MDPs for the purpose of validating the basic MBIE approach and illustrating some of its strengths.

To illustrate the algorithms, we report results with two simple MDPs. We measured the cumulative reward over the first 5000 steps of learning. The algorithms sought to maximize discounted reward ($\gamma = .95$).

Although not reported here, we also performed tests in several other domains and with Boltzmann exploration, action elimination (Even-Dar et al., 2003), and the MBIE algorithm due to Wiering and Schmidhuber (1998); none were consistently competitive with MBIE in our experiments.

Each of the three algorithms we studied has one principle parameter (ϵ , C , and D), which, coarsely, trades off between exploration and exploitation. To better quantify the effect of the parameters, we repeated experiments with a range of parameter settings. Results indicate that the values of the parameters had a very significant effect on the performance.

The two MDPs we used in our experiments are illustrated in Figure 1. **RiverSwim** consists of six states. The agent starts in one of the states near the beginning of the row. The two actions available to the agent are to swim left

or right. Swimming to the right (against the current of the river) will more often than not leave the agent in the same state, but will sometimes transition the agent to the right (and with a much smaller probability to the left). Swimming to the left (with the current) always succeeds in moving the agent to the left, until the leftmost state is reached at which point swimming to the left yields a small reward of five units. The agent receives a much larger reward, of ten thousand units, for swimming upstream and reaching the rightmost state. This MDP requires sequences of actions in order to explore effectively.

SixArms consists of seven states, one of which is the initial state. The agent must choose from six different actions. Each action pulls a different arm, with a different payoff probability, on a k -armed bandit. When the arm pays off, the agent is sent to another state. Within that new state, large rewards can be obtained. The higher the payoff probability for an arm, the smaller the reward received (from the room the agent is sent to by the arm). In this MDP, a decision maker can make use of smaller amounts of experience on the low paying arms and still perform well.

Figure 2 shows the algorithms’ performances on the two different domains at different parameter settings. Each of the graphs plots the cumulative discounted reward over 5000 steps, averaging over 1000 runs for each algorithm. Since we are concerned with efficient exploration, cumulative reward is an appropriate measure of performance—algorithms must quickly discover and exploit the environment in which they are learning.

As expected, all three algorithms exhibit an “inverted u” behavior—for small parameter values, cumulative reward is low because the policy learned is suboptimal due to inadequate exploration. For large parameter values, cumulative reward is low because the algorithm waits a long time before exploiting. Each algorithm performs best with an intermediate setting of its parameter.

The sole exception to this pattern is the performance of ϵ -greedy on **SixArms**, which achieves its best, although quite poor, performance with $\epsilon = 0$. This behavior is due to the fact that the exploration scheme of ϵ -greedy relies on the execution of a generally poor policy. When ϵ is high, the algorithm discovers a decent policy, but never executes it. Note that in Figure 2, for a given domain, the performance of each algorithm is the same at the extreme left of the graph since the completely greedy policy is followed when the relevant parameter is set to zero, for each of the algorithms we tested.

In **RiverSwim**, R-Max and MBIE are more effective at using focused exploration to ferret out the high-reward state by sequencing exploration steps. ϵ -greedy can reach this state, but only at a value of ϵ that results in near-random behavior and therefore low cumulative reward. In **SixArms**, MBIE is quickly able to rule out low-reward actions, re-

sulting in near optimal behavior. For R-Max, however, low settings of C result in premature convergence to a suboptimal policy while high values result in many wasted trials on the clearly suboptimal actions. Similarly ϵ -greedy fails to do well in **SixArms**. For both domains, MBIE performs at least as well as the best of the other two and better than both in one domain.

In Figure 3, the graph plots the average cumulative reward verses time for the MBIE algorithm on the **SixArms** domain. The slope of this graph measures the value of the policy being executed at that time. At low parameter settings, a small slope is achieved very quickly due to greedy sub-optimal behavior. At high parameter settings, a large positive slope may be achieved but only at a great expense in terms of learning time. The optimal parameter setting achieves a decent slope in a short amount of time. Although not provided, **RiverSwim** yields a similar looking graph.

6. Conclusion

The confidence-interval approach we described is efficient and theoretically motivated. It also performed quite well in the small-scale tests we reported. The basic idea of the approach is (1) use experienced transitions to create a model of the environment, (2) assign values to all states in a way that assigns state–action pairs initially optimistic values but more realistic values with increasing experience, (3) choose actions greedily with respect to these values. Our MBIE scheme uses a particular approach to assigning values, but many other options are also available. In Laplace smoothing, a prior is placed on each state–action pair that assumes it has experienced z transitions to an optimistically initialized state. We have experimented with this approach in bandit problems and found it to match or outperform MBIE. However, we also expect that theoretical results will be easier to obtain for MBIE.

We have also considered Good-Turing smoothing (McAllester & Schapire, 2000), since another way of viewing the problem of exploration is as one of estimating the probability of an unseen event (a better transition than any of those already witnessed).

We believe exploration is the key to efficient reinforcement learning. Future work needs to focus on automatically identifying effective parameter values and on extending exploration approaches to larger and more complex state spaces in which generalization is necessary.

Acknowledgments

Thanks to the National Science Foundation (IIS-0325281) and to DARPA IPTO for some early support. We also thank colleagues Sanjoy Dasgupta, John Langford, Shie Mannor, David McAllester, Rob Schapire,

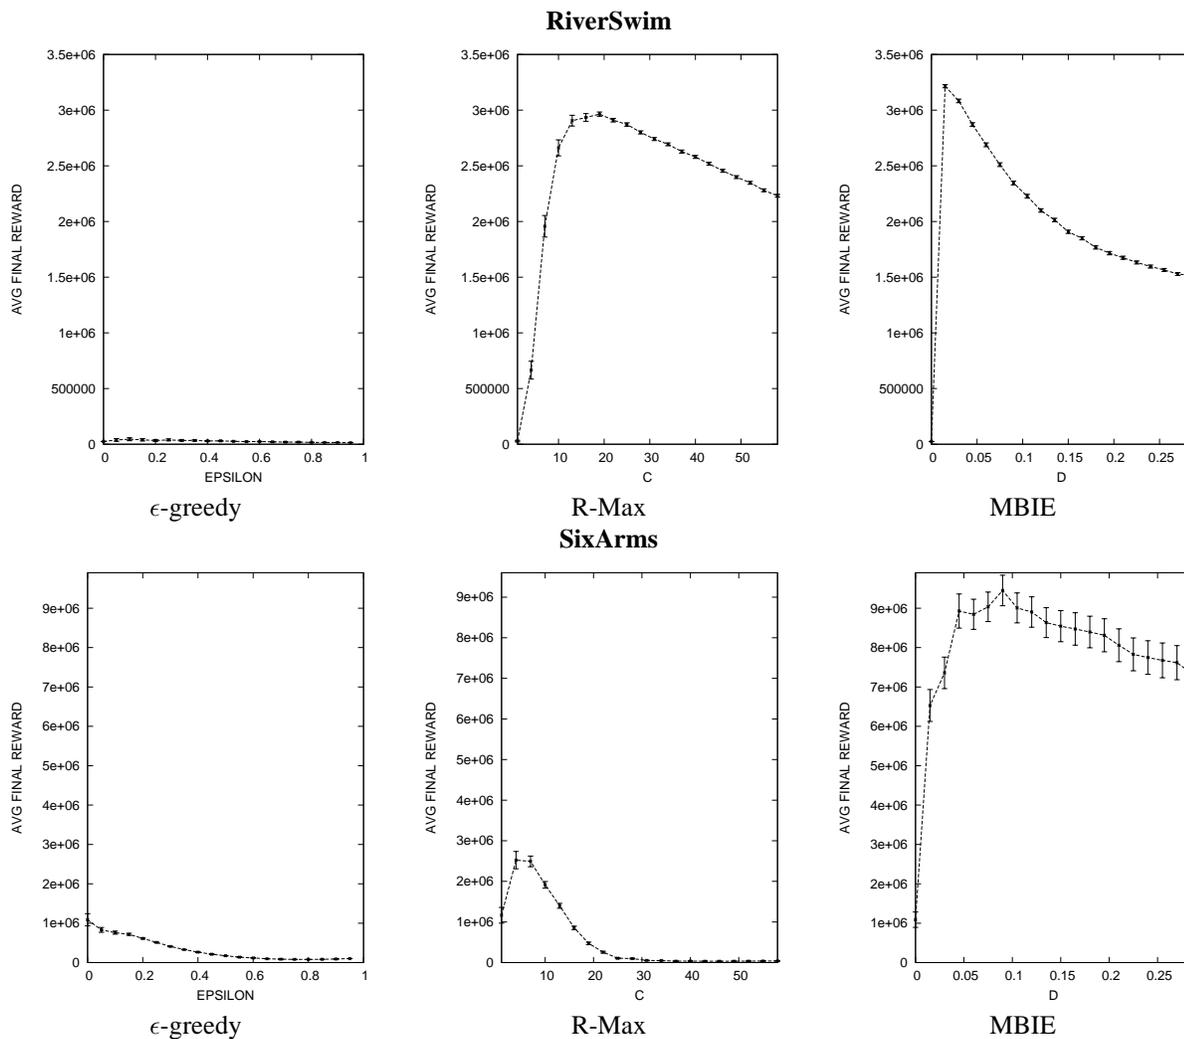


Figure 2. A comparison of three exploration schemes— ϵ -greedy (left), R-Max (middle), and MBIE (right)—on two example MDPs—RiverSwim (top) and SixArms (bottom)—shows how their parameters can be set to best trade off exploration and exploitation to maximize cumulative reward over an initial learning period.

Rich Sutton, Sergio Verdu, Tsachy Weissman, and Martin Zinkovich for suggestions.

References

- Berry, D. A., & Fristedt, B. (1985). *Bandit problems: Sequential allocation of experiments*. London, UK: Chapman and Hall.
- Brafman, R. I., & Tennenholtz, M. (2002). R-MAX—a general polynomial time algorithm for near-optimal reinforcement learning. *Journal of Machine Learning Research*, 3, 213–231.
- Dayan, P., & Sejnowski, T. J. (1996). Exploration bonuses and dual control. *Machine Learning*, 25, 5–22.
- Dearden, R., Friedman, N., & Andre, D. (1999). Model-based bayesian exploration. *Proceedings of the 15th Annual Conference on Uncertainty in Artificial Intelligence (UAI-99)* (pp. 150–159).
- Duff, M. O., & Barto, A. G. (1997). Local bandit approximation for optimal learning problems. *Advances in Neural Information Processing Systems* (pp. 1019–1025). The MIT Press.
- Even-Dar, E., Mannor, S., & Mansour, Y. (2003). Action elimination and stopping conditions for reinforcement

Cumulative Reward Plot for SixArms

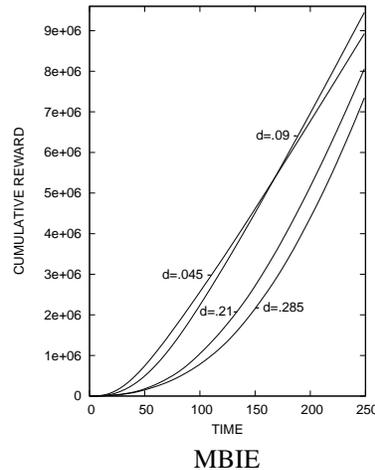


Figure 3. A plot of the average cumulative reward verses time for MBIE with a variety of parameter settings.

- learning. *The Twentieth International Conference on Machine Learning (ICML 2003)* (pp. 162–169).
- Fong, P. W. L. (1995). A quantitative study of hypothesis selection. *Proceedings of the Twelfth International Conference on Machine Learning (ICML-95)* (pp. 226–234).
- Gittins, J. C. (1989). *Multi-armed bandit allocation indices*. Wiley-Interscience series in systems and optimization. Chichester, NY: Wiley.
- Givan, R., Leach, S., & Dean, T. (2000). Bounded-parameter Markov decision processes. *Artificial Intelligence, 122*, 71–109.
- Kaelbling, L. P. (1993). *Learning in embedded systems*. Cambridge, MA: The MIT Press.
- Kearns, M. J., & Singh, S. P. (2002). Near-optimal reinforcement learning in polynomial time. *Machine Learning, 49*, 209–232.
- Koenig, S., & Simmons, R. G. (1993). Complexity analysis of real-time reinforcement learning. *Proceedings of the Eleventh National Conference on Artificial Intelligence* (pp. 99–105). Menlo Park, CA: AAAI Press/MIT Press.
- McAllester, D., & Schapire, R. E. (2000). On the convergence rate of Good-Turing estimators. *Proceedings of the 13th Annual Conference on Computational Learning Theory* (pp. 1–6).
- Meuleau, N., & Bourguine, P. (1999). Exploration of multi-state environments: Local measure and back-propagation of uncertainty. *Machine Learning, 35*, 117–154.
- Moore, A. W., & Atkeson, C. G. (1993). Prioritized sweeping: Reinforcement learning with less data and less real time. *Machine Learning, 13*, 103–130.
- Puterman, M. L. (1994). *Markov decision processes—discrete stochastic dynamic programming*. New York, NY: John Wiley & Sons, Inc.
- Strehl, A. L., & Littman, M. L. (2004). Efficient exploration via model-based interval estimation. Working paper.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. The MIT Press.
- Thrun, S. B. (1992). The role of exploration in learning control. In D. A. White and D. A. Sofge (Eds.), *Handbook of intelligent control: Neural, fuzzy, and adaptive approaches*, 527–559. New York, NY: Van Nostrand Reinhold.
- Weissman, T., Ordentlich, E., Seroussi, G., Verdu, S., & Weinberger, M. J. (2003). *Inequalities for the L1 deviation of the empirical distribution* (Technical Report HPL-2003-97R1). Hewlett-Packard Labs.
- Wiering, M. (1999). *Explorations in efficient reinforcement learning*. Doctoral dissertation, University of Amsterdam.
- Wiering, M., & Schmidhuber, J. (1998). Efficient model-based exploration. *Proceedings of the Fifth International Conference on Simulation of Adaptive Behavior (SAB'98)* (pp. 223–228).
- Wyatt, J. L. (2001). Exploration control in reinforcement learning using optimistic model selection. *Proceedings of the Eighteenth International Conference on Machine Learning (ICML 2001)* (pp. 593–600).