

COMPLEXITY OF STOCHASTIC BRANCH AND BOUND METHODS FOR BELIEF TREE SEARCH IN BAYESIAN REINFORCEMENT LEARNING

Christos Dimitrakakis
University of Amsterdam, The Netherlands
christos.dimitrakakis@gmail.com

Keywords: Exploration, Bayesian reinforcement learning, Belief tree search, Complexity, PAC bounds.

Abstract: There has been a lot of recent work on Bayesian methods for reinforcement learning exhibiting near-optimal online performance. The main obstacle facing such methods is that in most problems of interest, the optimal solution involves planning in an infinitely large tree. However, it is possible to obtain stochastic lower and upper bounds on the value of each tree node. This enables us to use stochastic branch and bound algorithms to search the tree efficiently. This paper proposes some algorithms and examines their complexity in this setting.

1 INTRODUCTION

Various Bayesian methods for exploration in Markov decision processes (MDPs) and for solving known partially-observable Markov decision processes (POMDPs), were proposed previously (c.f. (Poupart et al., 2006; Duff, 2002; Ross et al., 2008)). However, such methods often suffer from computational tractability problems. Optimal Bayesian exploration requires the creation of an augmented MDP model in the form of a tree (Duff, 2002), where the root node is the current belief-state pair and children are all possible subsequent belief-state pairs. The size of the belief tree increases exponentially with the horizon, while the branching factor is infinite in the case of continuous observations or actions.

In this work, we examine the complexity of efficient algorithms for expanding the tree. In particular, we propose and analyse stochastic search methods similar to the ones proposed in (Bubeck et al., 2008; Norkin et al., 1998). Related methods have been previously examined experimentally in the context of Bayesian reinforcement learning in (Dimitrakakis, 2008; Wang et al., 2005).

The remainder of this section summarises the Bayesian planning framework. Our main results are presented in Sect. 2. Section 3 concludes with a discussion of related work.

1.1 Markov Decision Processes

Reinforcement learning [c.f. (Puterman, 2005)] is discrete-time sequential decision making problem, where we wish to act so as to maximise the expected sum of discounted future rewards $\mathbf{E} \sum_{k=1}^T \gamma^k r_{t+k}$, where $r_t \in \mathbb{R}$ is a stochastic reward at time t . We are only interested in rewards from time t to $T > 0$, and $\gamma \in [0, 1]$ plays the role of a discount factor. Typically, we assume that γ and T are known (or have known prior distribution) and that the sequence of rewards arises from a Markov decision process μ :

Definition 1 (MDP). A Markov decision process is a discrete-time stochastic process with: A state $s_t \in \mathcal{S}$ at time t and a reward $r_t \in \mathbb{R}$, generated by the process μ , and an action $a_t \in \mathcal{A}$, chosen by the decision maker. We denote the distribution over next states s_{t+1} , which only depends on s_t and a_t , by $\mu(s_{t+1}|s_t, a_t)$. Furthermore $\mu(r_{t+1}|s_t, a_t)$ is a reward distribution conditioned on states and actions. Finally, $\mu(r_{t+1}, s_{t+1}|s_t, a_t) = \mu(r_{t+1}|s_t, a_t)\mu(s_{t+1}|s_t, a_t)$.

In the above, and throughout the text, we usually take $\mu(\cdot)$ to mean $\mathbf{P}_\mu(\cdot)$, the distribution under the process μ , for compactness. Frequently such a notation will imply a marginalisation. For example, we shall write $\mu(s_{t+k}|s_t, a_t)$ to mean $\sum_{s_{t+1}, \dots, s_{t+k-1}} \mu(s_{t+k}, \dots, s_{t+1}|s_t, a_t)$. The decision maker takes actions according to a policy π , which

defines a distribution $\pi(a_t|s_t)$ over \mathcal{A} , conditioned on the state s_t . A policy π is stationary if $\pi(a_t = a|s_t = s) = \pi(a_{t'} = a|s_{t'} = s)$ for all t, t' . The expected utility of a policy π selecting actions in the MDP μ , from time t to T can be written as the *value function*:

$$V_{t,T}^{\pi,\mu}(s) = \mathbf{E}_{\pi,\mu} \left(\sum_{k=1}^T \gamma^k r_{t+k} \mid s_t \right), \quad (1)$$

where $\mathbf{E}_{\pi,\mu}$ denotes the expectation under the Markov chain arising from acting policy π on the MDP μ . Whenever it is clear from context, superscripts and subscripts shall be omitted for brevity. The *optimal* value function will be denoted by $V^* \triangleq \max_{\pi} V^{\pi}$. If the MDP is known, we can evaluate the optimal value function policy in time polynomial to the sizes of the state and action sets (Puterman, 2005) via backwards induction (value iteration).

1.2 Bayesian Reinforcement Learning

If the MDP is unknown, we may use a Bayesian framework to represent our uncertainty (Duff, 2002). This requires maintaining a belief ξ_t , about which MDP $\mu \in \mathcal{M}$ corresponds to reality. More precisely, we define a measurable space $(\mathcal{M}, \mathfrak{M})$, where \mathcal{M} is a (usually uncountable) set of MDPs, and \mathfrak{M} is a suitable σ -algebra. With an appropriate initial density $\xi_0(\mu)$, we can obtain a sequence of densities $\xi_t(\mu)$, representing our subjective belief at time t , by conditioning $\xi_t(\mu)$ on the latest observations:

$$\xi_{t+1}(\mu) \triangleq \frac{\mu(r_{t+1}, s_{t+1} | s_t, a_t) \xi_t(\mu)}{\int_{\mathcal{M}} \mu'(r_{t+1}, s_{t+1} | s_t, a_t) \xi_t(\mu') d\mu'}. \quad (2)$$

In the following, we write \mathbf{E}_{ξ} to denote expectations with respect to any belief ξ .

1.3 Belief-augmented MDPs

In order to optimally select actions in this framework, it is necessary to *explicitly* take into account future changes in the belief when planning (Duff, 2002). The idea is to combine the original MDP's state s_t and our belief state ξ_t into a *hyper-state*.

Definition 2 (BAMDP). A *Belief-Augmented MDP* ν (BAMPD) is an MDP with a set of hyper-states $\Omega = \mathcal{S} \times \mathfrak{B}$, where \mathfrak{B} is an appropriate set of probability measures on \mathcal{M} and \mathcal{S}, \mathcal{A} are the state and action sets of all $\mu \in \mathcal{M}$. At time t , the agent observes the hyper-state $\omega_t = (s_t, \xi_t) \in \Omega$ and takes action $a_t \in \mathcal{A}$. We write the transition distribution as $\nu(\omega_{t+1} | \omega_t, a_t)$ and the reward distribution as $\nu(r_t | \omega_t)$.

The hyper-state ω_t has the Markov property. This allows us to treat the BAMDP as an infinite-state MDP with transitions $\nu(\omega_{t+1} | \omega_t, a_t)$, and rewards $\nu(r_t | \omega_t)$.¹ When the horizon T is finite, we need only require expand the tree to depth $T - t$. Thus, backwards induction starting from the set of terminal hyper-states Ω_T and proceeding backwards to $T - 1, \dots, t$ provides a solution:

$$V_n^*(\omega) = \max_{a \in \mathcal{A}} \mathbf{E}_{\nu}(r | \omega) + \gamma \sum_{\omega' \in \Omega_{n+1}} \nu(\omega' | \omega, a) V_{n+1}^*(\omega'), \quad (3)$$

where Ω_n is the set of hyper-states at time n . We can approximately solve infinite-horizon problems if we expand the tree to some finite depth, if we have bounds on the value of leaf nodes.

1.4 Bounds on the Value Function

We shall relate the optimal value function of the BAMDP, $V^*(\omega)$, for some $\omega(s, \xi)$, to the value functions V_{μ}^{π} of MDPs $\mu \in \mathcal{M}$ for some π . The optimal policy for μ is denoted as $\pi^*(\mu)$. The *mean* MDP resulting from belief ξ is denoted as $\bar{\mu}_{\xi}$ and has the properties: $\bar{\mu}_{\xi}(s_{t+1} | s_t, a_t) = \mathbf{E}_{\xi}[\mu(s_{t+1} | s_t, a_t)]$, $\bar{\mu}_{\xi}(r_{t+1} | s_t, a_t) = \mathbf{E}_{\xi}[\mu(r_{t+1} | s_t, a_t)]$.

Proposition 1 (Dimitrakakis, 2008). For any $\omega = (s, \xi)$, the BAMDP value function V^* obeys:

$$\int_{\mathcal{M}} V_{\mu}^{\pi^*(\mu)}(s) \xi(\mu) d\mu \geq V^*(\omega) \geq \int_{\mathcal{M}} V_{\mu}^{\pi^*(\bar{\mu}_{\xi})}(s) \xi(\mu) d\mu \quad (4)$$

Proof. By definition, $V^*(\omega) \geq V^{\pi}(\omega)$ for all ω , for any policy π . It is easy to see that the lower bound equals $V^{\pi^*(\bar{\mu}_{\xi})}(\omega)$, thus proving the right hand side. The upper bound follows from the fact that for any function f , $\max_x \int f(x, u) du \leq \int \max_x f(x, u) du$. ■

If \mathcal{M} is not finite, then we cannot calculate the upper bound of $V(\omega)$ in closed form. However, we can use Monte Carlo sampling: Given a hyper-state $\omega = (s, \xi)$, we draw m MDPs from its belief ξ : $\mu_1, \dots, \mu_m \sim \xi$,² estimate the value function for each μ_k , $\hat{v}_{U,k}^{\omega} \triangleq V_{\mu_k}^{\pi^*(\mu_k)}(s)$, and average the samples: $\hat{v}_{U,m}^{\omega} \triangleq$

¹Because of the way that the BAMDP ν is constructed from beliefs over \mathcal{M} , the next reward now depends on the next state rather than the current state and action.

²In the discrete case, we sample a multinomial distribution from each of the Dirichlet densities independently for the transitions. For the rewards we draw independent Bernoulli distributions from the Beta of each state-action pair.

$\frac{1}{m} \sum_{k=1}^m \hat{v}_{U,k}^\omega$. Let $v_U^\omega \triangleq \int_{\mathcal{M}} \xi_\omega(\mu) V_\mu^*(s_\omega) d\mu$. Then, $\lim_{m \rightarrow \infty} [\hat{v}_{U,m}^\omega] = v_U^\omega$ almost surely and $\mathbf{E}[\hat{v}_{U,m}^\omega] = v_U^\omega$.

Lower bounds can be calculated via a similar procedure. We begin by calculating the optimal policy $\pi^*(\bar{\mu}_\xi)$ for the mean MDP $\bar{\mu}_\xi$ arising from ξ . We then compute $\hat{v}_{L,k}^\omega \triangleq V_{\mu_k}^{\pi^*(\bar{\mu}_\xi)}$, the value of that policy for each sample μ_k and estimate $\hat{v}_{L,m}^\omega \triangleq \frac{1}{m} \sum_{k=1}^m \hat{v}_{L,k}^\omega$.

2 COMPLEXITY OF BELIEF TREE SEARCH

We now present our main results. Detailed proofs are available in an accompanying technical report (Dimitrakakis, 2009). We search trees which arise in the context of planning under uncertainty in MDPs using the BAMDP framework. We can use value function bounds on the leaf nodes of a partially expanded BAMDP tree to obtain bounds for the inner nodes through backwards induction. The bounds can be used both for action selection and for further tree expansion. However, the bounds are estimated via Monte Carlo sampling, something that necessitates the use of stochastic branch and bound technique to expand the tree.

We analyse a set of such algorithms. The first is a search to a fixed depth that employs exact lower bounds. We then show that if only stochastic bounds are available, the complexity of fixed depth search only increases logarithmically. We then present two stochastic branch and bound algorithms, whose complexity is dependent on the number of near-optimal branches. The first of these uses bound samples on leaf nodes only, while the second uses samples obtained in the last half of the parents of leaf nodes, thus using the collected samples more efficiently.

2.1 Assumptions and Notation

We present the main assumptions concerning the tree search, pointing out the relations to Bayesian RL. The symbols V and v have been overloaded to make this correspondence more apparent. The tree that has a branching factor at most ϕ . The branching is due to both action choices and random outcomes (see Fig. 1). Thus, the nodes at depth k correspond to the set of hyper-states $\{\omega_{t+k}\}$ in the BAMDP. By abusing notation, we may also refer to the components of each node $\omega = (s, \xi)$ as $s(\omega), \xi(\omega)$.

We define a branch b as a *set of policies* (i.e. the set of all policies starting with a particular action). The value of a branch b is $V^b \triangleq \max_{\pi \in b} V^\pi$. The root branch is the set of all policies, with value V^* . A




Figure 1: A belief tree, where the rewards are ignored for simplicity, with actions $\mathcal{A} = \{a^1, a^2\}$ and states $\mathcal{S} = \{s^1, s^2\}$.

hyper-state ω is b -reachable if $\exists \pi \in b$ s.t. $\mathbf{P}_{\pi, v}(\omega | \omega_t) > 0$. Any branch b can be partitioned at any b -reachable ω into a set of branches $B(b, \omega)$. A possible partition is any $b_i = \{\pi \in b : i = \arg \max_a \pi(a | \omega)\}$ for any $b_i \in B(b, \omega)$. We simplify this by considering only deterministic policies. We denote the k -horizon value function by $V^b(k) \triangleq \max_{\pi \in b} V_{t,k}^\pi(\omega_t)$. For each tree node $\omega = (s, \xi)$, we define upper and lower bounds $v_U(\omega) \triangleq \mathbf{E}_\xi[V_\mu^*(s)]$, $v_L(\omega) \triangleq \mathbf{E}_\xi[V^{\pi^*(\bar{\mu}_\xi)}(s)]$, from (4). By fully expanding the tree to depth k and performing backwards induction (3), using either v_U or v_L as the value of leaf nodes, we obtain respectively upper and lower bounds $V_U^b(k), V_L^b(k)$ on the value of any branch. Finally, we use $\mathcal{C}(\omega)$ for the set of immediate children of a node ω and the short-hand Ω_k for $\mathcal{C}^k(\omega)$, the set of all children of ω at depth k . We assume the following:

Assumption 1 (Uniform linear convergence). There exists $\gamma \in (0, 1)$ and $\beta > 0$ s.t. for any branch b , and depth k , $V^b - V_L^b(k) \leq \beta \gamma^k$, $V_U^b(k) - V^b \leq \beta \gamma^k$.

Remark 1. For BAMDPs with $r_t \in [0, 1]$ and $\gamma < 1$, Ass. 1 holds, from boundedness and the geometric series, with $\beta = 1/(1 - \gamma)$, since $V_L^b(k)$ and $V_U^b(k)$ are the k -horizon value functions with the value of leaf nodes bounded in $1/(1 - \gamma)$.

We analyse algorithms which search the tree and then select an (action) branch \hat{b}^* . For each algorithm, we examine the number of leaf node evaluations required to bound the regret $V^* - V^{\hat{b}^*}$.

2.2 Flat Search

With exact bounds, we can expand all branches to a fixed depth and then select the branch \hat{b}^* , with the highest lower bound. This is Alg. 1, with complexity given by the following lemma.

Algorithm 1. Flat oracle search.

-
- 1: Expand all branches until depth $k = \log_\gamma \varepsilon / \beta$ or $\hat{\Delta}_L > \beta \gamma^k - \varepsilon$.
 - 2: Select the root branch $\hat{b}^* = \arg \max_b V_L^b(k)$.
-

Algorithm 2. Flat stochastic search.

-
- 1: **procedure** FSSEARCH(ω_r, k, m)
 - 2: Let $\Omega_k = \{\omega_{r+k}^i : i = 1, \dots, \phi^k\}$ be the set of all k -step children of ω
 - 3: **for** $\omega \in \Omega_k$ **do**
 - 4: Draw m samples $\tilde{v}_{L,j}^\omega = V_\mu^\pi, \mu \sim \xi(\omega)$
 - 5: $\hat{v}_L^\omega = \frac{1}{m} \sum_{j=1}^m \tilde{v}_{L,j}^\omega$,
 - 6: **end for**
 - 7: Calculate \hat{V}^b , **return** $\hat{b}^* = \arg \max \hat{V}^b$.
 - 8: **end procedure**
-

Lemma 1. Alg. 1 on a tree with branching factor ϕ , $\gamma \in (0, 1)$, samples $O(\phi^{1+\log_\gamma \varepsilon / \beta})$ times to bound the regret by ε .

Proof. Bound the k -horizon value function error with Ass. 1 and note that there are ϕ^{k+1} leaves. ■

In our case, we only have a stochastic lower bound on the value of each node. Algorithm 2 expands the tree to a fixed depth and then takes multiple samples from each leaf node.

Lemma 2. Calling Alg. 2 with $k = \lceil \log_\gamma \varepsilon / 2\beta \rceil$, $m = 2 \lceil \log_\gamma (\varepsilon / 2\beta) \rceil \cdot \log \phi$, we bound the regret by ε using $O(\phi^{1+\log_\gamma \varepsilon / 2\beta} \log_\gamma (\varepsilon / 2\beta) \cdot \log \phi)$ samples.

Proof. The regret now is due to both limited depth and stochasticity. We bound each by $\varepsilon/2$, the first via Lem. 1 and the second via Hoeffding's inequality. ■

Thus, stochasticity mainly adds a logarithmic factor to the oracle search. We now consider two algorithms which do not search to a fixed depth, but select branches to deepen adaptively.

2.3 Stochastic Branch and Bound 1

A stochastic branch and bound algorithm similar to those examined here was originally developed by (Norkin et al., 1998) for optimisation problems. At each stage, it takes an additional sample at each leaf node, to improve their upper bound estimates, then expands the node with the highest mean upper bound. Algorithm 3 uses the same basic idea, averaging the value function samples at every leaf node.

In order to bound complexity, we need to bound the time required until we discover a nearly optimal

branch. We calculate the number of times a suboptimal branch is expanded before its suboptimality is discovered. Similarly, we calculate the number of times we shall sample the optimal node until its mean upper bound becomes dominant. These two results cover the time spent sampling upper bounds of nodes in the optimal branch without expanding them and the time spent expanding nodes in a sub-optimal branch.

Algorithm 3. Stochastic branch and bound 1.

-
- Let \mathcal{L}_0 be the root.
- for** $n = 1, 2, \dots$ **do**
- for** $\omega \in \mathcal{L}_n$ **do**
- $m_{\omega++}, \mu \sim \xi(\omega), \tilde{v}_{m_\omega}^\omega = V_\mu^*(s(\omega))$.
- $\hat{v}_U^\omega = \frac{1}{m_\omega} \sum_{i=1}^{m_\omega} \tilde{v}_i^\omega$
- end for**
- $\hat{\omega}_n^* = \arg \max_{\omega} \hat{v}_U^\omega$.
- $\mathcal{L}_{n+1} = \mathcal{C}(\hat{\omega}_n^*) \cup \mathcal{L}_n \setminus \hat{\omega}_n^*$
- end for**
-

Lemma 3. If N is the (random) number of samples \tilde{v}_i from random variable $V \in [0, \beta]$ we must take until its empirical mean $\hat{V}_k \triangleq \sum_{i=1}^k \tilde{v}_i > \mathbf{E}V - \Delta$, then:

$$\mathbf{E}[N] \leq 1 + \beta^2 \Delta^{-2} \quad (5)$$

$$\mathbf{P}[N > n] \leq \exp(-2\beta^{-2} n^2 \Delta^2). \quad (6)$$

Proof. The first inequality follows from the Hoeffding inequality and an integral bound on the resulting sum, while the second inequality is proven directly via a Hoeffding bound. ■

By setting Δ to be the difference between the optimal and second optimal branch, we can use the above lemma to bound the number of times N the leaf nodes in the optimal branch will be sampled without being expanded. The converse problem is bounding the number of times that a suboptimal branch will be expanded.

Lemma 4. If b is a branch with $V^b = V^* - \Delta$, then it will be expanded at least to depth $k_0 = \log_\gamma \Delta / \beta$. Subsequently,

$$\mathbf{P}(K > k) < O(\exp\{-2\beta^{-2} [(k - k_0)\Delta^2]\}). \quad (7)$$

Proof. In the worst case, the branch is degenerate and only one leaf has non-zero probability. We then apply a Hoeffding bound to obtain the desired result. ■

Algorithm 4. Stochastic branch and bound 2.

for $\omega \in \mathcal{L}_n$ **do**
 $\hat{V}_U^\omega = \frac{1}{\sum_{\omega' \in C(\omega)} m_{\omega'}} \sum_{\omega' \in C(\omega)} \sum_{i=1}^{m_{\omega'}} \hat{V}_i^{\omega'}$
end for
 Use (3) to obtain \hat{V}_U for all nodes.
 Set ω_0 to root.
for $d = 1, \dots$ **do**
 $a_d^* = \arg \max_a \sum_{\omega \in \Omega_d} \omega_{d-1}(j|a) \hat{V}_U(\omega)$
 $\omega_d \sim \omega_{d-1}(j|a_d^*)$.
if $\omega_d \in \mathcal{L}_n$ **then**,
 $\mathcal{L}_{n+1} = C(\omega_d) \cup \mathcal{L}_n \setminus \omega_d$
Break
end if
end for

2.4 Stochastic Branch and Bound 2

The degeneracy is the main problem of Alg. 3. Alg. 4 not only propagates upper bounds from multiple leaf nodes to the root, but also re-uses upper bound samples from inner nodes, in order to handle the degenerate case where only one path has non-zero probability. (Nevertheless, Lemma 3 applies without modification to Alg. 4). Because we are no longer operating on leaf nodes, we can take advantage of the upper bound samples collected along a given trajectory. However, if we use all of the upper bounds along a branch, then the early samples may bias our estimates a lot. For this reason, if a leaf is at depth k , we only average the upper bounds along the branch to depth $k/2$. The complexity of this approach is given by the following lemma:

Lemma 5. If b is s.t. $V^b = V^* - \Delta$, it will be expanded to depth $k_0 > \log_\gamma \Delta / \beta$ and

$$\mathbf{P}(K > k) \lesssim \exp(-2(k - k_0)^2(1 - \gamma^2)), \quad k > k_0$$

Proof. There is a degenerate case where only one sub-branch has non-zero probability. However we now re-use the samples that were obtained at previous expansions, thus allowing us to upper bound the bias by $\frac{\Delta(1 - \gamma^{k+1})}{(k - k_0)(1 - \gamma)}$. This allows to use a tighter Hoeffding bound and so obtain the desired outcome. ■

This bound decreases faster with k . Furthermore, there is no dependence on Δ after the initial transitory period, which may however be very long. The gain is due to the fact that we are re-using the upper bounds previously obtained in inner nodes. Thus, this algorithm should be particularly suitable for stochastic problems.

2.5 Lower Bounds for Bayesian RL

We can reduce the branching factor ϕ , (which is $|\mathcal{A} \times \mathcal{S} \times \mathcal{R}|$ for a full search) by employing sparse sampling methods (Kearns et al., 1999) to $O(|\mathcal{A}| \exp[1/(1 - \gamma)])$. This was essentially the approach employed by (Wang et al., 2005). However, our main focus here is to reduce the depth to which each branch is searched.

The main problem with the above algorithms is the fact that we must reach $k_0 = \lceil \log_\gamma \Delta \rceil$ to discard Δ -optimal branches. However, since the hyper-state ω_t arises from a Bayesian belief, we can use an additional smoothness property:

Lemma 6. The Dirichlet parameter sequence ψ_t/n_t , with $n_t \triangleq \sum_{i=1}^K \psi_t^i$, is a c -Lipschitz martingale with $c_t = 1/2(n_t + 1)$.

Proof. Simple calculations show that, no matter what is observed, $\mathbf{E}_{\xi_t}(\psi_{t+1}/n_{t+1}) = \psi_t/n_t$. Then, we bound the difference $|\psi_{t+k}/n_{t+k} - \psi_t/n_t|$ by two different bounds, which we equate to obtain c_t . ■

Lemma 7. If $\mu, \hat{\mu}$ are such that $\|\mathcal{T} - \hat{\mathcal{T}}\|_\infty \leq \varepsilon$ and $\|r - \hat{r}\|_\infty \leq \varepsilon$, for some $\varepsilon > 0$, then $\|V^\pi - \hat{V}^\pi\|_\infty \leq \frac{\varepsilon}{(1 - \gamma)^2}$, for any policy π .

Proof. By subtracting the Bellman equations for V, \hat{V} and taking the norm, we can repeatedly apply Cauchy-Schwarz and triangle inequalities to obtain the desired result. ■

The above results help us obtain better lower bounds in two ways. First we note that initially $1/k$ converges faster than γ^k , for large γ , thus we should be able to expand less deeply. Later, n_t is large so we can sample even more sparsely.

If we search to depth k , and the rewards are in $[0, 1]$, then, naively, our error is bounded by $\sum_{n=k}^\infty \gamma^n = \gamma^k / (1 - \gamma)$. However, the mean MDPs for $n > k$ are close to the mean MDP at k due to Lem. 6. This means that β can be significantly smaller than $1/(1 - \gamma)$. In fact, the total error is bounded by $\sum_{n=k}^\infty \gamma^n (n - k)/n$. For undiscounted problems, our error is bounded by $T - k$ in the original case and by $T - k[1 + \log(T/k)]$ when taking into account the smoothness.

3 CONCLUSIONS AND RELATED WORK

Much recent work on Bayesian RL focused on myopic estimates or full expansion of the belief tree

up to a certain depth. Exceptions include (Poupart et al., 2006), which uses an analytical bound based on sampling a small set of beliefs and (Wang et al., 2005), which uses Kearns's sparse sampling algorithm (Kearns et al., 1999) to expand the tree. Both methods have complexity exponential in the horizon, something which we improve via the use of smoothness properties induced by the Bayesian updating.

There are also connections with work on POMDPs problems (Ross et al., 2008). However this setting, though equivalent in an abstract sense, is not sufficiently close to the one we consider. Results on bandit problems, employing the same value function bounds used herein were reported in (Dimitrakakis, 2008), which experimentally compared algorithms operating on leaf nodes only.

Related results on the online sample complexity of Bayesian RL were developed by (Kolter and Ng, 2009), who employs a different upper bound to ours and (Asmuth et al., 2009), who employs MDP samples to plan in an augmented MDP space, similarly to (Auer et al., 2008) (who consider the set of plausible MDPs) and uses Bayesian concentration of measure results (Zhang, 2006) to prove mistake bounds on the online performance of the algorithm.

Interestingly, Alg. 4 resembles HOO (Bubeck et al., 2008) in the way that it traverses the tree, with two major differences. (a) The search is adapted to *stochastic* trees. (b) We use means of samples of upper bounds, rather than upper bounds on sample means. For these reasons, we are unable to simply restate the arguments in (Bubeck et al., 2008).

We presented complexity results and counting arguments for a number of tree search algorithms on trees where stochastic upper and lower bounds satisfying a smoothness property exist. These are the first results of this type and partially extend the results of (Norkin et al., 1998), which provided an asymptotic convergence proof, under similar smoothness conditions, for a stochastic branch and bound algorithm. In addition, we introduce a mechanism to utilise samples obtained at inner nodes when calculating mean upper bounds at leaf nodes. Finally, we relate our complexity results to those of (Kearns et al., 1999), for whose lower bound we provide a small improvement. We plan to address the online sample complexity of the proposed algorithms, as well as their practical performance, in future work.

ACKNOWLEDGEMENTS

This work was part of the ICIS project, supported by the Dutch Ministry of Economic Affairs, grant nr:

BSIK03024. I would like to thank the anonymous reviewers, as well as colleagues at the university of Amsterdam, Leoben and TU Crete for their comments on earlier versions of this paper.

REFERENCES

- Asmuth, J., Li, L., Littman, M. L., Nouri, A., and Wingate, D. (2009). A Bayesian sampling approach to exploration in reinforcement learning. In *UAI 2009*.
- Auer, P., Jaksch, T., and Ortner, R. (2008). Near-optimal regret bounds for reinforcement learning. In *Proceedings of NIPS 2008*.
- Bubeck, S., Munos, R., Stoltz, G., and Szepesvári, C. (2008). Online optimization in X-armed bandits. In *NIPS*, pages 201–208.
- Dimitrakakis, C. (2008). Tree exploration for Bayesian RL exploration. In *CIMCA'08*, pages 1029–1034, Los Alamitos, CA, USA. IEEE Computer Society.
- Dimitrakakis, C. (2009). Complexity of stochastic branch and bound for belief tree search in Bayesian reinforcement learning. Technical Report IAS-UVA-09-01, University of Amsterdam.
- Duff, M. O. (2002). *Optimal Learning Computational Procedures for Bayes-adaptive Markov Decision Processes*. PhD thesis, University of Massachusetts at Amherst.
- Kearns, M. J., Mansour, Y., and Ng, A. Y. (1999). A sparse sampling algorithm for near-optimal planning in large Markov decision processes. In Dean, T., editor, *IJCAI*, pages 1324–1331. Morgan Kaufmann.
- Kolter, J. Z. and Ng, A. Y. (2009). Near-Bayesian exploration in polynomial time. In *ICML 2009*.
- Norkin, V. I., Pflug, G. C., and Ruszczyski, A. (1998). A branch and bound method for stochastic global optimization. *Mathematical Programming*, 83(1):425–450.
- Poupart, P., Vlassis, N., Hoey, J., and Regan, K. (2006). An analytic solution to discrete Bayesian reinforcement learning. In *ICML 2006*, pages 697–704. ACM Press New York, NY, USA.
- Puterman, M. L. (1994,2005). *Markov Decision Processes : Discrete Stochastic Dynamic Programming*. John Wiley & Sons, New Jersey, US.
- Ross, S., Pineau, J., Paquet, S., and Chaib-draa, B. (2008). Online planning algorithms for POMDPs. *Journal of Artificial Intelligence Resesarch*, 32:663–704.
- Wang, T., Lizotte, D., Bowling, M., and Schuurmans, D. (2005). Bayesian sparse sampling for on-line reward optimization. In *ICML '05*, pages 956–963, New York, NY, USA. ACM.
- Zhang, T. (2006). From ϵ -entropy to KL-entropy: Analysis of minimum information complexity density estimation. *Annals of Statistics*, 34(5):2180–2210.