

# Tree Exploration for Bayesian RL Exploration

Christos Dimitrakakis  
Intelligent Systems Laboratory Amsterdam,  
University of Amsterdam  
The Netherlands  
c.dimitrakakis@uva.nl

## Abstract

*Research in reinforcement learning has produced algorithms for optimal decision making under uncertainty that fall within two main types. The first employs a Bayesian framework, where optimality improves with increased computational time. This is because the resulting planning task takes the form of a dynamic programming problem on a belief tree with an infinite number of states. The second type employs relatively simple algorithms which are shown to suffer small regret within a distribution-free framework. This paper presents a lower bound and a high probability upper bound on the optimal value function for the nodes in the Bayesian belief tree, which are analogous to similar bounds in POMDPs. The bounds are then used to create more efficient strategies for exploring the tree. The resulting algorithms are compared with the distribution-free algorithm UCBI, as well as a simpler baseline algorithm on multi-armed bandit problems.*

## 1 Introduction

In recent work [6,7,10,15–17,21], Bayesian methods for exploration in Markov decision processes (MDPs) and for solving known partially-observable Markov decision processes (POMDPs), as well as for exploration in the latter case, have been proposed. All such methods suffer from computational intractability problems for most domains of interest.

The sources of intractability are two-fold. Firstly, there may be no compact representation of the current belief. This is especially true for POMDPs. Secondly, optimally behaving under uncertainty requires that we create an *augmented* MDP model in the form of a tree [6], where the root node is the current belief-state pair and children are all possible subsequent belief-state pairs. This tree grows large very fast, and it is particularly problematic to grow in the case of continuous observations or actions. In this work,

we concentrate on the second problem – and consider algorithms for expanding the tree.

Since the Bayesian exploration methods require a tree expansion to be performed, we can view the whole problem as that of *nested* exploration. For the simplest exploration-exploitation trade-off setting, bandit problems, there already exist nearly optimal, computationally simple methods [1]. Such methods have recently been extended to tree search [12]. This work proposes to take advantage of the special structure of belief trees in order to design nearly-optimal algorithms for expansion of nodes. In a sense, by recognising that the tree expansion problem in Bayesian look-ahead exploration methods is also an optimal exploration problem, we develop tree algorithms that can solve this problem efficiently. Furthermore, we are able to derive interesting upper and lower bounds for the value of branches and leaf nodes which can help limit the amount of search. The ideas developed are tested in the multi-armed bandit setting for which nearly-optimal algorithms already exist.

The remainder of this section introduces the augmented MDP formalism employed within this work and discusses related work. Section 2 discusses tree expansion in exploration problems and introduces some useful bounds. These bounds are used in the algorithms detailed in Section 3, which are then evaluated in Section 4. We conclude with an outlook to further developments.

### 1.1 Preliminaries

We are interested in sequential decision problems where, at each time step  $t$ , the agent seeks to maximise the expected utility

$$\mathbf{E}[u_t \mid \cdot] \triangleq \sum_{k=1}^{\infty} \gamma^k \mathbf{E}[r_{t+k} \mid \cdot],$$

where  $r$  is a stochastic reward and  $u_t$  is simply the discounted sum of future rewards. We shall assume that the sequence of rewards arises from a Markov decision process, defined below.

**Definition 1.1 (Markov decision process)** A Markov decision process (MDP) is defined as the tuple  $\mu = (\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R})$  comprised of a set of states  $\mathcal{S}$ , a set of actions  $\mathcal{A}$ , a transition distribution  $\mathcal{T}$  conditioning the next state on the current state and action,

$$\mathcal{T}(s'|s, a) \triangleq \mu(s_{t+1}=s' | s_t=s, a_t=a) \quad (1)$$

satisfying the Markov property  $\mu(s_{t+1} | s_t, a_t) = \mu(s_{t+1} | s_t, a_t, s_{t-1}, a_{t-1}, \dots)$ , and a reward distribution  $\mathcal{R}$  conditioned on states and actions:

$$\mathcal{R}(r|s, a) \triangleq \mu(r_{t+1}=r | s_t=s, a_t=a), \quad (2)$$

with  $a \in \mathcal{A}$ ,  $s, s' \in \mathcal{S}$ ,  $r \in \mathbb{R}$ . Finally,

$$\mu(r_{t+1}, s_{t+1} | s_t, a_t) = \mu(r_{t+1} | s_t, a_t) \mu(s_{t+1} | s_t, a_t). \quad (3)$$

We shall denote the set of all MDPs as  $\mathcal{M}$ . For any policy  $\pi$  that is an arbitrary distribution on actions, we can define a  $T$ -horizon value function for an MDP  $\mu \in \mathcal{M}$  at time  $t$  as:

$$V_{t,T}^{\pi,\mu}(s, a) = \mathbf{E}[r_{t+1} | s_t=s, a_t=a, \mu] + \gamma \sum_{s'} \mu(s_{t+1}=s' | s_t=s, a_t=a) V_{\mu, t+1, T}^{\pi}(s').$$

Note that for the infinite-horizon case,  $\lim_{T \rightarrow \infty} V_{t,T}^{\pi,\mu} = V^{\pi,\mu}$  for all  $t$ .

In the case where the MDP is unknown, it is possible to use a Bayesian framework to represent our uncertainty (c.f. [6]). This essentially works by maintaining a belief  $\xi_t \in \Xi$ , about which MDP  $\mu \in \mathcal{M}$  corresponds to reality. In a Bayesian setting,  $\xi_t(\mu)$  is our subjective probability measure that  $\mu$  is true.

In order to optimally select actions in this framework, we need to use the approach suggested originally in [3] under the name of Adaptive Control Processes. The approach was investigated more fully in [6, 7]. This creates an *augmented* MDP, with a state comprised of the original MDP's state  $s_t$  and our belief state  $\xi_t$ . We can then solve the exploration *in principle* via standard dynamic programming algorithms such as backwards induction. We shall call such models Belief-Augmented MDPs, analogously to the Bayes-Adaptive MDPs of [6]. This is done by not only considering densities conditioned on the state-action pairs  $(s_t, a_t)$ , i.e.  $p(r_{t+1}, s_{t+1} | s_t, a_t)$ , but taking into account the belief  $\xi_t \in \Xi$ , a probability space over possible MDPs, i.e. augmenting the state space from  $\mathcal{S}$  to  $\mathcal{S} \times \Xi$  and considering the following conditional density:  $p(r_{t+1}, s_{t+1}, \xi_{t+1} | s_t, a_t, \xi_t)$ . More formally, we may give the following definition:

**Definition 1.2 (Belief-Augmented MDP)** A Belief-Augmented MDP  $\nu$  (BAMPD) is an MDP  $\nu =$

$(\Omega, \mathcal{A}, \mathcal{T}', \mathcal{R}')$  where  $\Omega = \mathcal{S} \times \Xi$ , where  $\Xi$  is the set of probability measures on  $\mathcal{M}$ , and  $\mathcal{T}', \mathcal{R}'$  are the transition and reward distributions conditioned jointly on the MDP state  $s_t$ , the belief state  $\xi_t$ , and the action  $a_t$ . Here  $\xi_t(\xi_{t+1} | r_{t+1}, s_{t+1}, s_t, a_t)$  is singular, so that we can define the transition

$$p(\omega_{t+1} | a_t, \omega_t) \equiv p(s_{t+1}, \xi_{t+1} | a_t, s_t, \xi_t).$$

It should be obvious that  $s_t, \xi_t$  jointly form a Markov state in this setting, called the *hyper-state*. In general, we shall denote the components of a future hyper-state  $\omega_t^i$  as  $(s_t^i, \xi_t^i)$ . However, in occasion we will abuse notation by referring to the components of some hyperstate  $\omega$  as  $s_\omega, \xi_\omega$ . We shall use  $\mathcal{M}_B$  to denote the set of BMDPs.

As in the MDP case, finite horizon problems only require sampling all future actions until the horizon  $T$ .

$$V_{t,T}^{\pi*}(\omega_t, a_t) = \mathbf{E}[r_{t+1} | \omega_t, a_t] + \gamma \int_{\Omega} V_{t+1,T}^{\pi*}(\omega_{t+1}) \nu(\omega_{t+1} | \omega_t, a_t) d\omega_{t+1}. \quad (4)$$

However, because the set of hyper-states available at each time-step is necessarily different from those at other time-steps, the value function cannot be easily calculated for the infinite horizon case.

In fact, the only clear solution is to continue expanding a belief tree until we are certain of the optimality of an action. As has previously been observed [4,5], this is possible since we can always obtain upper and lower bounds on the utility of any policy from the current hyper-state. We can apply such bounds on future hyper-states in order to efficiently expand the tree.

## 1.2 Related work

Up to date, most work had only used full expansion of the belief tree up to a certain depth. A notable exception is [22], which uses Thompson sampling [20] to expand the tree. In very recent work [18], the importance of tree expansion in the closely related POMDP setting<sup>1</sup> has been recognised. Therein, the authors contrast and compare many different methods for tree expansion, including branch-and-bound [13] methods and Monte Carlo sampling.

Monte Carlo sampling methods have also been recently explored in the upper confidence bounds on trees (UCT) algorithms, proposed in [8, 12] in the context of planning in games. Our case is similar, however we can take advantage of the special structure of the belief tree. In particular, for each node we can obtain a lower bound and a high-probability upper bound on the value of the optimal policy.

<sup>1</sup>The BAMDP setting is equivalent to a POMDP where the unobservable part of the state is stationary, but continuous (chap. 5 [6])

This paper’s contribution is to recognise that tree expansion in Bayesian exploration is itself an exploration problem with very special properties. Based on this insight, it proposes to combine sampling with lower bounds and upper bound estimates at the leaves. This allows us to obtain high-probability bounds for expansion of the tree. While the proposed methods are similar to the ones used in the discrete-state POMDP setting [18], the BAMDP requires the evaluation of different bounds at leaf nodes. On the experimental side, we present first results on bandit problems, for which nearly-optimal distribution-free algorithms are known. We believe that this is a very important step towards extending the applicability of Bayesian look-ahead methods in exploration.

## 2 Belief tree expansion

Let the current belief be  $\xi_t$  and suppose we observe  $x_t^i \triangleq (s_{t+1}^i, r_{t+1}^i, a_t^i)$ . This observation defines a unique subsequent belief  $\xi_{t+1}^i$ . Together with the MDP state  $s$ , this creates a hyper-state transition from  $\omega_t$  to  $\omega_{t+1}^i$ . By recursively obtaining observations for future beliefs, we can obtain an unbalanced tree with nodes  $\{\omega_{t+k}^i : k = 1, \dots, T; i = 1, \dots\}$ . However, we cannot hope to be able to fully expand the tree. This is especially true in the case where observations (i.e. states, rewards, or actions) are continuous, where we cannot perform even a full single-step expansion. Even in the discrete case the problem is intractable for infinite horizons – and far too complex computationally for the finite horizon case. However, had there been efficient tree expansion methods, this problem would be largely alleviated. The remainder of this section details bounds and algorithms that can be used to reduce the computational complexity of the Bayesian lookahead approach.

### 2.1 Expanding a given node

All tree search methods require the expansion of leaf nodes. However, in general, a leaf node may have an infinite number of children. We thus need some strategies to limit the number of children.

More formally, let us assume that we wish to expand in node  $\omega_t^i = (\xi_t^i, s_t^i)$ , with  $\xi_t^i$  defining a density over  $\mathcal{M}$ . For discrete state/action/reward spaces, we can simply enumerate all the possible outcomes  $\{\omega_{t+1}^j\}_{j=1}^{|\mathcal{S} \times \mathcal{A} \times \mathcal{R}|}$ , where  $\mathcal{R}$  is the set of possible reward outcomes. Note that if the reward is deterministic, there is only one possible outcome per state-action pair. The same holds if  $\mathcal{T}$  is deterministic, in both cases making an enumeration possible. While in general this may not be the case, since rewards, states, or actions can be continuous, in this paper we shall only examine the discrete case.

### 2.2 Bounds on the optimal value function

At each point in the process, the next node  $\omega_{t+k}^i$  to be expanded is the one maximising a utility  $U(\omega_{t+k}^i)$ . Let  $\Omega_T$  be the set of leaf nodes. If their values were known, then we could easily perform the backwards induction procedure shown in Algorithm 1. The main problem is obtaining a

---

#### Algorithm 1 Backwards induction action selection

---

```

1: procedure BACKWARDSINDUCTION( $t, \nu, \Omega_T, V_T^*$ )
2:   for  $n = T - 1, T - 2, \dots, t$  do
3:     for  $\omega \in \Omega_n$  do
4:        $a(\omega) = \arg \max_a [\mathbf{E}(r|\omega', \omega, \nu) +$ 
5:          $\gamma V_{n+1}^*(\omega')] \nu(\omega'|\omega, a)$ 
6:        $V_n^*(\omega) = \sum_{\omega' \in \Omega_{n+1}} \nu(\omega'|\omega, a_n^*) [\mathbf{E}(r|\omega', \omega, \nu) +$ 
7:          $\gamma V_{n+1}^*(\omega')]$ 
8:     return  $a_t^*$ 
9: end procedure

```

---

good estimate for  $V_T^*$ , i.e. the value of leaf nodes. Let  $\pi^*(\mu)$  denote the policy such that, for any  $\pi$ ,

$$V_\mu^{\pi^*(\mu)}(s) \geq V_\mu^\pi(s) \quad \forall s \in \mathcal{S}.$$

Furthermore, let the maximum probability MDP arising from the belief at hyper-state  $\omega$  be  $\hat{\mu}_\omega \triangleq \arg \max_\mu \hat{\mu}$ . Similarly, we denote the mean MDP with  $\bar{\mu}_\omega \triangleq \mathbf{E}[\mu|\xi_\omega]$ .

**Proposition 2.1** *The optimal value function at any leaf node  $\omega$  is bounded by the following inequalities*

$$\int V_\mu^{\pi^*(\mu)}(s_\omega) \xi_\omega(\mu) d\mu \geq V^*(\omega) \geq \int V_\mu^{\pi^*(\bar{\mu}_\omega)}(s_\omega) \xi_\omega(\mu) d\mu. \quad (5)$$

**Proof** By definition,  $V^*(\omega) \geq V^\pi(\omega)$  for all  $\omega$ , for any policy  $\pi$ . The lower bound follows trivially, since

$$V^{\pi^*(\bar{\mu}_\omega)}(\omega) \triangleq \int V_\mu^{\pi^*(\bar{\mu}_\omega)}(s_\omega) \xi_\omega(\mu) d\mu. \quad (6)$$

The upper bound is derived as follows. First note that for any function  $f$ ,  $\sup_x \int f(x, u) du \leq \int \sup_x f(x, u) du$ . Then, we remark that:

$$V^*(\omega) = \sup_\pi \int V_\mu^\pi(s_\omega) \xi_\omega(\mu) d\mu \quad (7a)$$

$$\leq \int \sup_\pi V_\mu^\pi(s_\omega) \xi_\omega(\mu) d\mu \quad (7b)$$

$$= \int V_\mu^{\pi^*(\mu)}(s_\omega) \xi_\omega(\mu) d\mu. \quad (7c)$$

■

In POMDPs, a trivial lower bound can be obtained by calculating the value of the blind policy [9, 19], which always takes the same action. Our lower bound is in fact the BAMDP analogue of the value of the blind policy in POMDPs. This is because for any *fixed policy*  $\pi$ , it holds trivially that  $V^\pi(\omega) \leq V^*(\omega)$ . In our case, we have made this lower bound tighter by considering  $\pi^*(\bar{\mu}_\omega)$ , the policy that is greedy with respect to the current mean estimate.

The upper bound itself is analogous to the POMDP value function bound given in Theorem 9 of [9]. However, while the lower bound is easy to compute in our case, the upper bound can only be approximated via Monte Carlo sampling with some probability.

### 2.3 Calculating the lower bound

The lower bound can be calculated by performing value iteration in the mean MDP. This is because, for any policy  $\pi$  and belief  $\xi$ ,  $\int V_\mu^\pi(s)\xi(s) d\mu$  can be written as

$$\begin{aligned} & \int \left\{ \mathbf{E}[r|s, \mu, \pi] + \gamma \sum_{s'} \mu(s'|s, \pi(s)) V_\mu^\pi(s') \right\} \xi(\mu) d\mu = \\ & = \sum_a \pi(a|s) \left\{ \int \mathbf{E}[r|s, a, \mu] \xi(\mu) d\mu \right. \\ & \quad \left. + \gamma \sum_{s'} \int \mu(s'|s, a) V_\mu^\pi(s') \xi(\mu) d\mu \right\} \\ & = \sum_a \pi(a|s) \left\{ \mathbf{E}[r|s, a, \bar{\mu}_\xi] + \gamma \sum_{s'} \bar{\mu}_\xi(s'|s, a) V_{\bar{\mu}_\xi}^\pi(s') \right\}. \end{aligned}$$

where  $\bar{\mu}_\xi$  is the *mean MDP* for belief  $\xi$ . If the beliefs  $\xi$  can be expressed in closed form, it is easy to calculate the mean transition distribution and the mean reward from  $\xi$ . For discrete state spaces, transitions can be expressed as multinomial distributions, to which the Dirichlet density is a conjugate prior. In that case, for Dirichlet parameters  $\{\psi_i^{j,a}(\xi) : i, j \in \mathcal{S}, a \in \mathcal{A}\}$ , we have  $\bar{\mu}_\xi(s'|s, a) = \psi_{s'}^{s,a}(\xi) / \sum_{i \in \mathcal{S}} \psi_i^{s,a}(\xi)$ . Similarly, for Bernoulli rewards, the corresponding mean model arising from the beta prior with parameters  $\{\alpha^{s,a}(\xi), \beta^{s,a}(\xi) : s \in \mathcal{S}, a \in \mathcal{A}\}$  is  $\mathbf{E}[r|s, a, \bar{\mu}_\xi] = \alpha^{s,a}(\xi) / (\alpha^{s,a}(\xi) + \beta^{s,a}(\xi))$ . Then the value function of the mean model, and consequently, a lower bound on the optimal value function, can be found with standard value iteration.

### 2.4 Upper bound with high probability

In general, (7b) cannot be expressed in closed form. However, the integral can be approximated via Monte Carlo sampling.

Let the leaf node which we wish to expand be  $\omega$ . Then, we can obtain  $c$  MDP samples from the belief at

$\omega$ :  $\mu_1, \dots, \mu_c \sim \xi_\omega(\mu)$ . For each  $\mu_k$  we can derive the optimal policy  $\pi^*(\mu_k)$  and estimate its value function  $\tilde{v}_k^* \triangleq V_{\mu_k}^{\pi^*(\mu_k)} \equiv V_{\mu_k}^*$ . We may then average these samples to obtain

$$\hat{v}_c^*(\omega) \triangleq \frac{1}{c} \sum_{k=1}^c \tilde{v}_k^*(s_\omega). \quad (8)$$

Let  $\bar{v}^*(\omega) = \int_{\mathcal{M}} \xi_\omega(\mu) V_\mu^*(s_\omega) d\mu$ . It holds that  $\lim_{c \rightarrow \infty} [\hat{v}_c^*] = \bar{v}^*(\omega)$  and that  $\mathbf{E}[\hat{v}_c^*] = \bar{v}^*(\omega)$ . Due to the latter, we can apply a Hoeffding inequality

$$\mathbf{P}(|\hat{v}_c^*(\omega) - \bar{v}^*(\omega)| > \epsilon) < 2 \exp\left(-\frac{2c\epsilon^2}{(V_{\max} - V_{\min})^2}\right), \quad (9)$$

thus bounding the error within which we estimate the upper bound. For  $r_t \in [0, 1]$  and discount factor  $\gamma$ , note that  $V_{\max} - V_{\min} \leq 1/(1 - \gamma)$ .

### 2.5 Bounds on parent nodes

We can obtain upper and lower bounds on the value of every action  $a \in \mathcal{A}$ , at any part of the tree, by iterating over  $\Omega_t$ , the set of possible outcomes following  $\omega_t$ :

$$\bar{v}(\omega_t, a) = \sum_{i=1}^{|\Omega_t|} \mathbf{P}(\omega_t^i | \omega_t, a) [r_t^i + \gamma \bar{v}(\omega_t^i)] \quad (10)$$

$$\hat{v}^*(\omega_t, a) = \sum_{i=1}^{|\Omega_t|} \mathbf{P}(\omega_t^i | \omega_t, a) [r_t^i + \gamma \hat{v}^*(\omega_t^i)], \quad (11)$$

where the probabilities are implicitly conditional on the beliefs at each  $\omega_t$ . For every node, we can calculate an upper and lower bound on the value of all actions. Obviously, if at the root node  $\omega_t$ , there exists some  $\hat{a}_t^*$  such that  $\bar{v}(\omega_t, \hat{a}_t^*) \geq \bar{v}^*(\omega_t, a)$  for all  $a$ , then  $\hat{a}_t^*$  is unambiguously the optimal action.

However, in general, there may be some other action,  $a'$ , whose upper bound is higher than the lower bound of  $\hat{a}_t^*$ . In that case, we should expand either one of the two trees.

It is easy to see that the upper and lower bounds at any node  $\omega_t$  can be expressed as a function of the respective bounds at the leaf nodes. Let  $B(\omega_t, a)$  be the set of all branches from  $\omega_t$  when action  $a$  is taken. For each branch  $b \in B(\omega_t, a)$ , let  $\xi_t(b)$  be the probability of the branch from  $\omega_t$  and  $u_t^b$  be the discounted cumulative reward along the branch. Finally, let  $L(\omega_t, a)$  be the set of leaf nodes reachable from  $\omega_t$  and  $\omega_b$  be the specific node reachable from branch  $b$ . Then, upper or lower bounds on the value function can simply be expressed as  $v(\omega_t, a) = \sum_{b \in B(\omega_t, a)} u_t^b + \gamma^{t_b} \xi_{\omega_t}(b) v(\omega_b)$ . This would allow us to use a heuristic for greedily minimising the uncertainty at any branch.<sup>2</sup> However, the algorithms we shall consider here will only employ evaluation of upper and lower bounds.

<sup>2</sup>As an example, let  $a, a'$  such that  $\bar{v}(\omega_t, a') > \hat{v}^*(\omega_t, a)$ . The

### 3 Algorithms

All the expansion algorithms employed herein calculate a utility function  $U$  for every node  $\omega$  in the set of leaf nodes  $L(\omega_t)$ . The main difference among the algorithms is the way  $U$  is calculated.

1. *Serial*. This results in a nearly balanced tree, as the oldest leaf node is expanded next, i.e.  $U(\omega_i) = -i$ , the negative node index.
2. *Random*. In this case, we expand any of the current leaf nodes with equal probability, i.e.  $U(\omega_i) = U(\omega_j)$  for all  $i, j$ . This can of course lead to unbalanced trees.
3. *Highest lower bound*. We expand the node maximising a lower bound i.e.  $U(\omega_i) = \gamma^{t_i} \bar{V}(\omega_i)$ .
4. *Thompson sampling*. We expand the node for which the currently sampled upper bound is highest, i.e.  $U(\omega_i) = \gamma^{t_i} \tilde{v}(\omega_i)$ .
5. *Extreme upper bound*. We expand the node with the highest overall upper bound  $U(\omega_i) = \gamma^{t_i} \max\{\max_{k=1, \dots, c(i)} \tilde{v}_k^*(\omega_i), \bar{v}(\omega_i)\}$ .
6. *High probability upper bound*. We expand the node with the highest mean upper bound  $U(\omega_i) = \gamma^{t_i} \max_k\{\hat{v}_{c(i)}^*(\omega_i), \bar{v}(\omega_i)\}$ .

While methods 3 and 4 only use one sample from the upper bound calculation at every iteration. The last two methods retain the samples obtained in the previous iterations and use them to calculate the mean estimate.

### 4 Experiments

We compared the regret of the tree expansion to the optimal policy in bandit problems with Bernoulli rewards with two benchmarks: the UCB1 algorithm [2], which suffers only logarithmic regret, and secondly a policy that is greedy with respect to a mean Bayesian estimate with a prior density  $\text{Beta}(0, 1)$ , which is a good optimistic heuristic for such problems.

Figure 1 shows the cumulative undiscounted regret for horizon  $T = 2/(1 - \gamma)$ , with  $\gamma = 0.9999$  and  $|\mathcal{A}| = 2$ , averaged over 1000 runs. We compare the UCB1 algorithm

amount by which the lower bound can be increased for any branch  $b \in B(\omega_t, a)$  following action  $a$  is bounded by

$$\Delta V(b) = \gamma^{t_b} \xi_{\omega_t}(b) (\hat{v}^*(\omega_b) - \bar{v}(\omega_b)),$$

while the amount by which the upper bound can be increased for any branch  $b' \in B(\omega_t, a')$  is similarly bounded by  $\Delta V(b')$ . A greedy way to expand the tree would simply select the node with the highest difference in the two bounds' values, and thus the highest potential for removing the ambiguity about the optimality of two actions.




Figure 1. Cumulative undiscounted regret accrued over  $2/(1 - \gamma)$  time-steps.

(*ucb*), and the Bayesian baseline (*base*) with the BAMDP approach. The figure shows the cumulative undiscounted regret as a function of the number of look-aheads, for the following expansion algorithms: *serial*, *random*, highest lower bound (*LB*), Thompson sampling (*Thompson*) and extreme and high probability upper bound (*x UB* and *h.p. UB*). The latter three algorithms use  $\gamma$ -rate discounting for future node expansion.

What is perhaps surprising in these results is how much faster the lower bound and serial expansions converge to a good solution compared to the other expansion methods. For problems with more arms, these differences are amplified. In fact, the high probability upper bound expansion performs worse than random expansion, which is particularly interesting because in the POMDP setting, upper bound expansions appear to be best [18]. This is despite the fact that the high probability upper bound is tighter than the extreme upper bound. In fact, the latter bound is performing slightly better than random expansion and about as well as Thompson sampling. On the other hand, perhaps this problem could be alleviated if the node selection strategy was not flat.

### 5 Conclusion

One of this paper's aims was to draw attention to the interesting problem of tree expansion in Bayesian RL exploration. To this end, bounds on the optimal value function at belief tree leaf nodes have been derived and then utilised as heuristics for tree expansion. It is shown experimentally that the resulting expansion methods have very significant differences in computational complexity for bandit problems. While the results are preliminary in the sense that

no experiments on more complex problems are presented and that only very simple expansion algorithms have been tried, they are nevertheless significant in the sense that the effect of the tree exploration method used is very large.

Apart from performing further experiments, especially with more sophisticated expansion algorithms, future work should include deriving bounds on the minimum and maximum depth reached for each algorithm, as well as more general regret bounds if possible. The regret could be measured either as simply the  $\epsilon, \delta$  optimality of  $\hat{a}^*$ , or, more interestingly, bounds on the cumulative online regret suffered by each algorithm. More importantly, problems with infinite observation spaces (i.e. with continuous rewards) should also be examined.

Looking further ahead, it would also be interesting to use some type of *stochastic* branch-and-bound algorithm such as the ones described in [11, 14]. Another interesting approach would be to develop a new expansion algorithm that achieves a small anytime regret, perhaps in the lines of UCT [12]. Such algorithms have been very successful in solving problems with large spaces and may be useful in this problem as well, especially when the space of observations becomes larger.

## Acknowledgments

This work was supported by the ICIS-IAS project. Thanks to Carsten Cibura and Frans Groen for useful discussions and to Aikaterini Mitrokotsa for proofreading.

## References

- [1] P. Auer. Using confidence bounds for exploitation-exploration trade-offs. *J. Machine Learning Research*, 3(Nov):397–422, 2002. A preliminary version has appeared in *Proc. of the 41th Annual Symposium on Foundations of Computer Science*.
- [2] P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2/3):235–256, 2002. A preliminary version has appeared in *Proc. of the 15th International Conference on Machine Learning*.
- [3] R. Bellman and R. Kalaba. A mathematical theory of adaptive control processes. *Proceedings of the National Academy of Sciences of the United States of America*, 45(8):1288–1290, 1959.
- [4] R. Dearden, N. Friedman, and S. J. Russell. Bayesian Q-learning. In *AAAI/IAAI*, pages 761–768, 1998.
- [5] C. Dimitrakakis. Nearly optimal exploration-exploitation decision thresholds. In *Int. Conf. on Artificial Neural Networks (ICANN)*, 2006. IDIAP-RR 06-12.
- [6] M. O. Duff. *Optimal Learning Computational Procedures for Bayes-adaptive Markov Decision Processes*. PhD thesis, University of Massachusetts at Amherst, 2002.
- [7] M. O. Duff and A. G. Barto. Local bandit approximation for optimal learning problems. In M. C. Mozer, M. I. Jordan, and T. Petsche, editors, *Advances in Neural Information Processing Systems*, volume 9, page 1019. The MIT Press, 1997.
- [8] S. Gelly and D. Silver. Combining online and offline knowledge in UCT. In *ICML '07: Proceedings of the 24th international conference on Machine learning*, pages 273–280, New York, NY, USA, 2007. ACM Press.
- [9] M. Hauskrecht. Value-function approximations for partially observable markov decision processes. *Journal of Artificial Intelligence Research*, pages 33–94, Aug 2000.
- [10] M. Hoffman, A. Doucet, N. D. Freitas, and A. Jasra. Bayesian policy learning with trans-dimensional mcmc. In J. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in Neural Information Processing Systems 20*. MIT Press, Cambridge, MA, 2008.
- [11] A. Kleywegt, A. Shapiro, and T. Homem-de Mello. The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization*, 12(2):479–502, 2001.
- [12] L. Kocsis and C. Szepesvári. Bandit based Monte-Carlo planning. In *Proceedings of ECML-2006*, 2006.
- [13] L. Mitten. Branch-and-bound methods: General formulation and properties. *Operations Research*, 18(1):24–34, 1970.
- [14] V. I. Norkin, G. C. Pflug, and A. Ruszczyński. A branch and bound method for stochastic global optimization. *Mathematical Programming*, 83(1):425–450, January 1998.
- [15] P. Poupart, N. Vlassis, J. Hoey, and K. Regan. An analytic solution to discrete Bayesian reinforcement learning. *Proceedings of the 23rd international conference on Machine learning*, pages 697–704, 2006.
- [16] S. Ross, B. Chaib-draa, and J. Pineau. Bayes-adaptive POMDPs. In J. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in Neural Information Processing Systems 20*, Cambridge, MA, 2008. MIT Press.
- [17] S. Ross, J. Pineau, and B. Chaib-draa. Theoretical analysis of heuristic search methods for online POMDPs. In J. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in Neural Information Processing Systems 20*. MIT Press, Cambridge, MA, 2008.
- [18] S. Ross, J. Pineau, S. Paquet, and B. Chaib-draa. Online planning algorithms for POMDPs. *Journal of Artificial Intelligence Research*, 32:663–704, July 2008.
- [19] T. Smith and R. Simmons. Point-based POMDP algorithms: Improved analysis and implementation. In *Proceedings of the 21st Conference on Uncertainty in Artificial Intelligence (UAI-05)*, pages 542–547, 2005.
- [20] W. Thompson. On the Likelihood that One Unknown Probability Exceeds Another in View of the Evidence of two Samples. *Biometrika*, 25(3-4):285–294, 1933.
- [21] M. Toussaint, S. Harmelign, and A. Storkey. Probabilistic inference for solving (PO)MDPs, 2006.
- [22] T. Wang, D. Lizotte, M. Bowling, and D. Schuurmans. Bayesian sparse sampling for on-line reward optimization. In *ICML '05: Proceedings of the 22nd international conference on Machine learning*, pages 956–963, New York, NY, USA, 2005. ACM.