
Monte-Carlo Tree Search by Best Arm Identification

Emilie Kaufmann

CNRS & Univ. Lille, UMR 9189 (CRIStAL), Inria SequeL
Lille, France
emilie.kaufmann@univ-lille1.fr

Wouter M. Koolen

Centrum Wiskunde & Informatica,
Science Park 123, 1098 XG Amsterdam, The Netherlands
wmkoolen@cwi.nl

Abstract

Recent advances in bandit tools and techniques for sequential learning are steadily enabling new applications and are promising the resolution of a range of challenging related problems. We study the game tree search problem, where the goal is to quickly identify the optimal move in a given game tree by sequentially sampling its stochastic payoffs. We develop new algorithms for trees of arbitrary depth, that operate by summarizing all deeper levels of the tree into confidence intervals at depth one, and applying a best arm identification procedure at the root. We prove new sample complexity guarantees with a refined dependence on the problem instance. We show experimentally that our algorithms outperform existing elimination-based algorithms and match previous special-purpose methods for depth-two trees.

1 Introduction

We consider two-player zero-sum turn-based interactions, in which the sequence of possible successive moves is represented by a maximin game tree \mathcal{T} . This tree models the possible actions sequences by a collection of MAX nodes, that correspond to states in the game in which player A should take action, MIN nodes, for states in the game in which player B should take action, and leaves which specify the payoff for player A. The goal is to determine the best action at the root for player A. For deterministic payoffs this search problem is primarily algorithmic, with several powerful pruning strategies available [20]. We look at problems with stochastic payoffs, which in addition present a major statistical challenge.

Sequential identification questions in game trees with stochastic payoffs arise naturally as robust versions of bandit problems. They are also a core component of Monte Carlo tree search (MCTS) approaches for solving intractably large deterministic tree search problems, where an entire sub-tree is represented by a stochastic leaf in which randomized play-out and/or evaluations are performed [4]. A play-out consists in finishing the game with some simple, typically random, policy and observing the outcome for player A.

For example, MCTS is used within the AlphaGo system [21], and the evaluation of a leaf position combines supervised learning and (smart) play-outs. While MCTS algorithms for Go have now reached expert human level, such algorithms remain very costly, in that many (expensive) leaf evaluations or play-outs are necessary to output the next action to be taken by the player. In this paper, we focus on the *sample complexity* of Monte-Carlo Tree Search methods, about which very little is known. For this purpose, we work under a simplified model for MCTS already studied by [22], and that generalizes the depth-two framework of [10].

1.1 A simple model for Monte-Carlo Tree Search

We start by fixing a game tree \mathcal{T} , in which the root is a MAX node. Letting \mathcal{L} be the set of leaves of this tree, for each $\ell \in \mathcal{L}$ we introduce a stochastic oracle \mathcal{O}_ℓ that represents the leaf evaluation or play-out performed when this leaf is reached by an MCTS algorithm. In this model, we do not try to optimize the evaluation or play-out strategy, but we rather assume that the oracle \mathcal{O}_ℓ produces i.i.d. samples from an unknown distribution whose mean μ_ℓ is the value of the position ℓ . To ease the presentation, we focus on binary oracles (indicating the win or loss of a play-out), in which the oracle \mathcal{O}_ℓ is a Bernoulli distribution with unknown mean μ_ℓ (the probability of player A winning the game in the corresponding state). Our algorithms can be used without modification in case the oracle is a distribution bounded in $[0, 1]$.

For each node s in the tree, we denote by $\mathcal{C}(s)$ the set of its children and by $\mathcal{P}(s)$ its parent. The root is denoted by s_0 . The *value* (for player A) of any node s is recursively defined by $V_\ell = \mu_\ell$ if $\ell \in \mathcal{L}$ and

$$V_s = \begin{cases} \max_{c \in \mathcal{C}(s)} V_c & \text{if } s \text{ is a MAX node,} \\ \min_{c \in \mathcal{C}(s)} V_c & \text{if } s \text{ is a MIN node.} \end{cases}$$

The best move is the action at the root with highest value,

$$s^* = \operatorname{argmax}_{s \in \mathcal{C}(s_0)} V_s.$$

To identify s^* (or an ϵ -close move), an MCTS algorithm sequentially selects paths in the game tree and calls the corresponding leaf oracle. At round t , a leaf $L_t \in \mathcal{L}$ is chosen by this adaptive *sampling rule*, after which a sample $X_t \sim \mathcal{O}_{L_t}$ is collected. We consider here the same PAC learning framework as [22, 10], in which the strategy also requires a *stopping rule*, after which leaves are no longer evaluated, and a *recommendation rule* that outputs upon stopping a guess $\hat{s}_\tau \in \mathcal{C}(s_0)$ for the best move of player A.

Given a risk level δ and some accuracy parameter $\epsilon \geq 0$ our goal is have a recommendation $\hat{s}_\tau \in \mathcal{C}(s_0)$ whose value is within ϵ of the value of the best move, with probability larger than $1 - \delta$, that is

$$\mathbb{P}(V(s_0) - V(\hat{s}_\tau) \leq \epsilon) \geq 1 - \delta.$$

An algorithm satisfying this property is called (ϵ, δ) -correct. The main challenge is *to design (ϵ, δ) -correct algorithms that use as few leaf evaluations τ as possible.*

Related work The model we introduce for Monte-Carlo Tree Search is very reminiscent of a stochastic bandit model. In those, an agent repeatedly selects one out of several probability distributions, called arms, and draws a sample from the chosen distribution. Bandits models have been studied since the 1930s [23], mostly with a focus on *regret minimization*, where the agent aims to maximize the sum of the samples collected, which are viewed as rewards [18]. In the context of MCTS, a sample corresponds to a win or a loss in one play-out, and maximizing the number of successful play-outs (that correspond to simulated games) may be at odds with identifying quickly the next best action to take at the root. In that, our best action identification problem is closer to a so-called *Best Arm Identification* (BAI) problem.

The goal in the standard BAI problem is to find quickly and accurately the arm with highest mean. The BAI problem in the fixed-confidence setting [7] is the special case of our simple model for a tree of depth one. For deeper trees, rather than finding the best arm (i.e. leaf), we are interested in finding the best action at the root. As the best root action is a function of the means of all leaves, this is a more structured problem.

Bandit algorithms, and more recently BAI algorithms have been successfully adapted to tree search. Building on the UCB algorithm [2], a regret minimizing algorithm, variants of the UCT algorithm [17] have been used for MCTS in growing trees, leading to successful AIs for games. However, there are only very weak theoretical guarantees for UCT. Moreover, observing that maximizing the number of successful play-outs is not the target, recent work rather tried to leverage tools from the BAI literature. In [19, 6] Sequential Halving [14] is used for exploring game trees. The latter algorithm is a state-of-the-art algorithm for the fixed-budget BAI problem [1], in which the goal is to identify the best arm with the smallest probability of error based on a given budget of draws. The proposed SHOT (Sequential Halving applied to Trees) algorithm [6] is compared empirically to the UCT approach of [17], showing improvements in some cases. A hybrid approach mixing SHOT and UCT is also studied [19], still without sample complexity guarantees.

In the fixed-confidence setting, [22] develop the first sample complexity guarantees in the model we consider. The proposed algorithm, FindTopWinner is based on *uniform sampling and eliminations*, an approach that may be related to the Successive Eliminations algorithm [7] for fixed-confidence BAI in bandit models. FindTopWinner proceeds in rounds, in which the leaves that have not been eliminated are sampled repeatedly until the precision of their estimates doubled. Then the tree is pruned of every node whose estimated value differs significantly from the estimated value of its parent, which leads to the possible elimination of several leaves. For depth-two trees, [10] propose an elimination procedure that is not round-based. In this simpler setting, an algorithm that exploits *confidence intervals* is also developed, inspired by the LUCB algorithm for fixed-confidence BAI [13]. Some variants of the proposed M-LUCB algorithm appear to perform better in simulations than elimination based algorithms. We now investigate this trend further in deeper trees, both in theory and in practice.

Our Contribution. In this paper, we propose a generic architecture, called BAI-MCTS, that builds on a Best Arm Identification (BAI) algorithm and on confidence intervals on the node values in order to solve the best action identification problem in a tree of arbitrary depth. In particular, we study two specific instances, UGapE-MCTS and LUCB-MCTS, that rely on *confidence-based* BAI algorithms [8, 13]. We prove that these are (ϵ, δ) -correct and give a high-probability upper bound on their sample complexity. Both our theoretical and empirical results improve over the elimination-based state-of-the-art algorithm, FindTopWinner [22].

2 BAI-MCTS algorithms

We present a generic class of algorithms, called BAI-MCTS, that combines a BAI algorithm with an exploration of the tree based on confidence intervals on the node values. Before introducing the algorithm and two particular instances, we first explain how to build such confidence intervals, and also introduce the central notion of *representative child* and *representative leaf*.

2.1 Confidence intervals and representative nodes

For each leaf $\ell \in \mathcal{L}$, using the past observations from this leaf we may build a confidence interval

$$\mathcal{I}_\ell(t) = [L_\ell(t), U_\ell(t)],$$

where $U_\ell(t)$ (resp. $L_\ell(t)$) is an Upper Confidence Bound (resp. a Lower Confidence Bound) on the value $V(\ell) = \mu_\ell$. The specific confidence interval we shall use will be discussed later.

These confidence intervals are then propagated upwards in the tree using the following construction. For each internal node s , we recursively define $\mathcal{I}_s(t) = [L_s(t), U_s(t)]$ with

$$L_s(t) = \begin{cases} \max_{c \in \mathcal{C}(s)} L_c(t) & \text{for a MAX node } s, \\ \min_{c \in \mathcal{C}(s)} L_c(t) & \text{for a MIN node } s, \end{cases} \quad U_s(t) = \begin{cases} \max_{c \in \mathcal{C}(s)} U_c(t) & \text{for a MAX node } s, \\ \min_{c \in \mathcal{C}(s)} U_c(t) & \text{for a MIN node } s. \end{cases}$$

Note that these intervals are the tightest possible on the parent under the sole assumption that the child confidence intervals are all valid. A similar construction was used in the OMS algorithm of [3] in a different context. It is easy to convince oneself (or prove by induction, see Appendix B.1) that the accuracy of the confidence intervals is preserved under this construction, as stated below.

Proposition 1. *Let $t \in \mathbb{N}$. One has $\bigcap_{\ell \in \mathcal{L}} (\mu_\ell \in \mathcal{I}_\ell(t)) \Rightarrow \bigcap_{s \in \mathcal{T}} (V_s \in \mathcal{I}_s(t))$.*

We now define the *representative child* $c_s(t)$ of an internal node s as

$$c_s(t) = \begin{cases} \operatorname{argmax}_{c \in \mathcal{C}(s)} U_c(t) & \text{if } s \text{ is a MAX node,} \\ \operatorname{argmin}_{c \in \mathcal{C}(s)} L_c(t) & \text{if } s \text{ is a MIN node,} \end{cases}$$

and the *representative leaf* $\ell_s(t)$ of a node $s \in \mathcal{T}$, which is the leaf obtained when going down the tree by always selecting the representative child:

$$\ell_s(t) = s \text{ if } s \in \mathcal{L}, \quad \ell_s(t) = \ell_{c_s(t)}(t) \text{ otherwise.}$$

The confidence intervals in the tree represent the statistically plausible values in each node, hence the representative child can be interpreted as an “optimistic move” in a MAX node and a “pessimistic move” in a MIN node (assuming we play against the best possible adversary). This is reminiscent of the behavior of the UCT algorithm [17]. The construction of the confidence intervals and associated representative children are illustrated in Figure 1.




Figure 1: Construction of confidence interval and representative child (in red) for a MAX node.

Input: a BAI algorithm
Initialization: $t = 0$.
while not BAIStop ($\{s \in \mathcal{C}(s_0)\}$) **do**
 $R_{t+1} = \text{BAIStep}(\{s \in \mathcal{C}(s_0)\})$
 Sample the representative leaf
 $L_{t+1} = \ell_{R_{t+1}}(t)$
 Update the information about the arms.
 $t = t + 1$.
end
Output: BAIReco ($\{s \in \mathcal{C}(s_0)\}$)

Figure 2: The BAI-MCTS architecture

2.2 The BAI-MCTS architecture

In this section we present the generic BAI-MCTS algorithm, whose sampling rule combines two ingredients: a best arm identification step which selects an action at the root, followed by a *confidence based exploration step*, that goes down the tree starting from this depth-one node in order to select the *representative leaf* for evaluation.

The structure of a BAI-MCTS algorithm is presented in Figure 2. The algorithm depends on a Best Arm Identification (BAI) algorithm, and uses the three components of this algorithm:

- the sampling rule $\text{BAIStep}(\mathcal{S})$ selects an arm in the set \mathcal{S}
- the stopping rule $\text{BAIStop}(\mathcal{S})$ returns `True` if the algorithm decides to stop
- the recommendation rule $\text{BAIReco}(\mathcal{S})$ selects an arm as a candidate for the best arm

In BAI-MCTS, the arms are the depth-one nodes, hence the information needed by the BAI algorithm to make a decision (e.g. BAIStep for choosing an arm, or BAIStop for stopping) is information about depth-one nodes, that has to be updated at the end of each round (last line in the `while` loop). Different BAI algorithms may require different information, and we now present two instances that rely on confidence intervals (and empirical estimates) for the value of the depth-one nodes.

2.3 UGapE-MCTS and LUCB-MCTS

Several Best Arm Identification algorithms may be used within BAI-MCTS, and we now present two variants, that are respectively based on the UGapE [8] and the LUCB [13] algorithms. These two algorithms are very similar in that they exploit confidence intervals and use the same stopping rule, however the LUCB algorithm additionally uses the empirical means of the arms, which within BAI-MCTS requires defining an estimate $\hat{V}_s(t)$ of the value of the depth-one nodes.

The generic structure of the two algorithms is similar. At round $t + 1$ two promising depth-one nodes are computed, that we denote by \underline{b}_t and \underline{c}_t . Among these two candidates, the node whose confidence interval is the largest (that is, the most uncertain node) is selected:

$$R_{t+1} = \underset{i \in \{\underline{b}_t, \underline{c}_t\}}{\text{argmax}} [U_i(t) - L_i(t)].$$

Then, following the BAI-MCTS architecture, the representative leaf of R_{t+1} (computed by going down the tree) is sampled: $L_{t+1} = \ell_{R_{t+1}}(t)$. The algorithm stops whenever the confidence intervals of the two promising arms overlap by less than ϵ :

$$\tau = \inf \{t \in \mathbb{N} : U_{\underline{c}_t}(t) - L_{\underline{b}_t}(t) < \epsilon\},$$

and it recommends $\hat{s}_\tau = \underline{b}_\tau$.

In both algorithms that we detail below \underline{b}_t represents a guess for the best depth-one node, while \underline{c}_t is an “optimistic” challenger, that has the maximal possible value among the other depth-one nodes. Both nodes need to be explored enough in order to discover the best depth-one action quickly.

UGapE-MCTS. In UGapE-MCTS, introducing for each depth-one node the index

$$B_s(t) = \max_{s' \in \mathcal{C}(s_0) \setminus \{s\}} U_{s'}(t) - L_s(t),$$

the promising depth-one nodes are defined as

$$\underline{b}_t = \operatorname{argmin}_{a \in \mathcal{C}(s_0)} B_a(t) \quad \text{and} \quad \underline{c}_t = \operatorname{argmax}_{b \in \mathcal{C}(s_0) \setminus \{\underline{b}_t\}} U_b(t).$$

LUCB-MCTS. In LUCB-MCTS, the promising depth-one nodes are defined as

$$\underline{b}_t = \operatorname{argmax}_{a \in \mathcal{C}(s_0)} \hat{V}_a(t) \quad \text{and} \quad \underline{c}_t = \operatorname{argmax}_{b \in \mathcal{C}(s_0) \setminus \{\underline{b}_t\}} U_b(t),$$

where $\hat{V}_s(t) = \hat{\mu}_{\ell_s(t)}(t)$ is the empirical mean of the representative leaf of node s . Note that several alternative definitions of $\hat{V}_s(t)$ may be proposed (such as the middle of the confidence interval $\mathcal{I}_s(t)$, or $\max_{a \in \mathcal{C}(s)} \hat{V}_a(t)$), but our choice is crucial for the analysis of LUCB-MCTS, given in Appendix C.

3 Analysis of UGapE-MCTS

In this section we first prove that UGapE-MCTS and LUCB-MCTS are both (ϵ, δ) -correct. Then we give in Theorem 3 a high-probability upper bound on the number of samples used by UGapE-MCTS. A similar upper bound is obtained for LUCB-MCTS in Theorem 9, stated in Appendix C.

3.1 Choosing the Confidence Intervals

From now on, we assume that the confidence intervals on the leaves are of the form

$$L_\ell(t) = \hat{\mu}_\ell(t) - \sqrt{\frac{\beta(N_\ell(t), \delta)}{2N_\ell(t)}} \quad \text{and} \quad U_\ell(t) = \hat{\mu}_\ell(t) + \sqrt{\frac{\beta(N_\ell(t), \delta)}{2N_\ell(t)}}. \quad (1)$$

$\beta(s, \delta)$ is some *exploration function*, that can be tuned to have a δ -PAC algorithm, as expressed in the following lemma, whose proof can be found in Appendix B.2

Lemma 2. *If $\delta \leq \max(0.1|\mathcal{L}|, 1)$, for the choice*

$$\beta(s, \delta) = \ln(|\mathcal{L}|/\delta) + 3 \ln \ln(|\mathcal{L}|/\delta) + (3/2) \ln(\ln s + 1) \quad (2)$$

both UGapE-MCTS and LUCB-MCTS satisfy $\mathbb{P}(V(s^) - V(\hat{s}_\tau) \leq \epsilon) \geq 1 - \delta$.*

An interesting practical feature of these confidence intervals is that they only depend on the local number of draws $N_\ell(t)$, whereas most of the BAI algorithms use exploration functions that depend on the number of rounds t . Hence the only confidence intervals that need to be updated at round t are those of the ancestors of the selected leaf, which can be done recursively.

Moreover, $\beta(s, \delta)$ scales with $\ln(\ln(s))$, and not $\ln(s)$, leveraging some tools recently introduced to obtain tighter confidence intervals [12, 15]. The union bound over \mathcal{L} (that may be an artifact of our current analysis) however makes the exploration function of Lemma 2 still a bit over-conservative and in practice, we recommend the use of $\beta(s, \delta) = \ln(\ln(es)/\delta)$.

Finally, similar correctness results (with slightly larger exploration functions) may be obtained for confidence intervals based on the Kullback-Leibler divergence (see [5]), which are known to lead to better performance in standard best arm identification problems [16] and also depth-two tree search problems [10]. However, the sample complexity analysis is much more intricate, hence we stick to the above Hoeffding-based confidence intervals for the next section.

3.2 Complexity term and sample complexity guarantees

We first introduce some notation. Recall that s^* is the optimal action at the root, identified with the depth-one node satisfying $V(s^*) = V(s_0)$, and define the second-best depth-one node as $s_2^* =$

$\operatorname{argmax}_{s \in \mathcal{C}(s_0) \setminus \{s^*\}} V_s$. Recall $\mathcal{P}(s)$ denotes the parent of a node s different from the root. Introducing furthermore the set $\operatorname{Anc}(s)$ of all the ancestors of a node s , we define the complexity term by

$$H_\epsilon^*(\boldsymbol{\mu}) := \sum_{\ell \in \mathcal{L}} \frac{1}{\Delta_\ell^2 \vee \Delta_*^2 \vee \epsilon^2}, \quad \text{where} \quad \begin{aligned} \Delta_* &:= V(s^*) - V(s_2^*) \\ \Delta_\ell &:= \max_{s \in \operatorname{Anc}(\ell) \setminus \{s_0\}} |V_s - V(\mathcal{P}(s))| \end{aligned} \quad (3)$$

The intuition behind these squared terms in the denominator is the following. We will sample a leaf ℓ until we either prune it (by determining that it or one of its ancestors is a bad move), prune everyone else (this happens for leaves below the optimal arm) or reach the required precision ϵ .

Theorem 3. *Let $\delta \leq \min(1, 0.1|\mathcal{L}|)$. UGapE-MCTS using the exploration function (2) is such that, with probability larger than $1 - \delta$, $(V(s^*) - V(\hat{s}_\tau) < \epsilon)$ and, letting $\bar{\Delta}_{\ell, \epsilon} = \Delta_\ell \vee \Delta_* \vee \epsilon$,*

$$\begin{aligned} \tau &\leq 8H_\epsilon^*(\boldsymbol{\mu}) \ln \frac{|\mathcal{L}|}{\delta} + \sum_{\ell} \frac{16}{\bar{\Delta}_{\ell, \epsilon}^2} \ln \ln \frac{1}{\bar{\Delta}_{\ell, \epsilon}} \\ &\quad + 8H_\epsilon^*(\boldsymbol{\mu}) \left[3 \ln \ln \frac{|\mathcal{L}|}{\delta} + 2 \ln \ln \left(8e \ln \frac{|\mathcal{L}|}{\delta} + 24e \ln \ln \frac{|\mathcal{L}|}{\delta} \right) \right] + 1. \end{aligned}$$

Remark 4. *If $\beta(N_a(t), \delta)$ is changed to $\beta(t, \delta)$, one can still prove (ϵ, δ) correctness and furthermore upper bound the expectation of τ . However the algorithm becomes less efficient to implement, since after each leaf observation, ALL the confidence intervals have to be updated. In practice, this change lowers the probability of error but does not effect significantly the number of play-outs used.*

3.3 Comparison with previous work

To the best of our knowledge¹, the FindTopWinner algorithm [22] is the only algorithm from the literature designed to solve the best action identification problem in any-depth trees. The number of play-outs of this algorithm is upper bounded with high probability by

$$\sum_{\ell: \Delta_\ell > 2\epsilon} \left(\frac{32}{\Delta_\ell^2} \ln \frac{16|\mathcal{L}|}{\Delta_\ell \delta} + 1 \right) + \sum_{\ell: \Delta_\ell \leq 2\epsilon} \left(\frac{8}{\epsilon^2} \ln \frac{8|\mathcal{L}|}{\epsilon \delta} + 1 \right)$$

One can first note the improvement in the constant in front of the leading term in $\ln(1/\delta)$, as well as the presence of the $\ln \ln(1/\Delta_{\ell, \epsilon^2})$ second order, that is unavoidable in a regime in which the gaps are small [12]. The most interesting improvement is in the control of the number of draws of 2ϵ -optimal leaves (such that $\Delta_\ell \leq 2\epsilon$). In UGapE-MCTS, the number of draws of such leaves is at most of order $(\epsilon \vee \Delta_*^2)^{-1} \ln(1/\delta)$, which may be significantly smaller than $\epsilon^{-1} \ln(1/\delta)$ if there is a gap in the best and second best value. Moreover, unlike FindTopWinner and M-LUCB [10] in the depth two case, UGapE-MCTS can also be used when $\epsilon = 0$, with provable guarantees.

Regarding the algorithms themselves, one can note that M-LUCB, an extension of LUCB suited for depth-two tree, does *not* belong to the class of BAI-MCTS algorithms. Indeed, it has a “reversed” structure, first computing the representative leaf for each depth-one node: $\forall s \in \mathcal{C}(s_0), R_{s,t} = \ell_s(t)$ and then performing a BAI step over the representative leaves: $\tilde{L}_{t+1} = \text{BAIStep}(R_{s,t}, s \in \mathcal{C}(s_0))$. This alternative architecture can also be generalized to deeper trees, and was found to have empirical performance similar to BAI-MCTS. M-LUCB, which will be used as a benchmark in Section 4, also distinguish itself from LUCB-MCTS by the fact that it uses an exploration rate that depends on the global time $\beta(t, \delta)$ and that b_t is the empirical maximin arm (which can be different from the arm maximizing \hat{V}_s). This alternative choice is not yet supported by theoretical guarantees in deeper trees.

Finally, the exploration step of BAI-MCTS algorithm bears some similarity with the UCT algorithm [17], as it goes down the tree choosing alternatively the move that yields the highest UCB or the lowest LCB. However, the behavior of BAI-MCTS is very different at the root, where the first move is selected using a BAI algorithm. Another key difference is that BAI-MCTS relies on *exact* confidence

¹In a recent paper, [11] independently proposed the LUCBMinMax algorithm, that differs from UGapE-MCTS and LUCB-MCTS only by the way the best guess b_t is picked. The analysis is very similar to ours, but features some refined complexity measure, in which Δ_ℓ (that is the maximal distance between *consecutive* ancestors of the leaf, see (3)) is replaced by the maximal distance between *any* ancestors of that leaf. Similar results could be obtained for our two algorithms following the same lines.

intervals: each interval $\mathcal{I}_s(t)$ is shown to contain with high probability the corresponding value V_s , whereas UCT uses more heuristic confidence intervals, based on the number of visits of the parent node, and aggregating all the samples from descendant nodes. Using UCT in our setting is not obvious as it would require to define a suitable stopping rule, hence we don't include a comparison with this algorithm in Section 4. A hybrid comparison between UCT and FindTopWinner is proposed in [22], providing UCT with the random number of samples used by the fixed-confidence algorithm. It is shown that FindTopWinner has the advantage for hard trees that require many samples. Our experiments show that our algorithms in turn always dominate FindTopWinner.

3.4 Proof of Theorem 3.

Letting $\mathcal{E}_t = \bigcap_{\ell \in \mathcal{L}} (\mu_\ell \in \mathcal{I}_\ell(t))$ and $\mathcal{E} = \bigcap_{t \in \mathbb{N}} \mathcal{E}_t$, we upper bound τ assuming the event \mathcal{E} holds, using the following key result, which is proved in Appendix D.

Lemma 5. *Let $t \in \mathbb{N}$. $\mathcal{E}_t \cap (\tau > t) \cap (L_{t+1} = \ell) \Rightarrow N_\ell(t) \leq \frac{8\beta(N_\ell(t), \delta)}{\Delta_\ell^2 \vee \Delta_*^2 \vee \epsilon^2}$.*

An intuition behind this result is the following. First, using that the selected leaf ℓ is a representative leaf, it can be seen that the confidence intervals from $s_D = \ell$ to s_0 are nested (Lemma 11). Hence if \mathcal{E}_t holds, $V(s_k) \in \mathcal{I}_\ell(t)$ for all $k = 1, \dots, D$, which permits to lower bound the width of this interval (and thus upper bound $N_\ell(t)$) as a function of the $V(s_k)$ (Lemma 12). Then Lemma 13 exploits the mechanism of UGapE to further relate this width to Δ_* and ϵ .

Another useful tool is the following lemma, that will allow to leverage the particular form of the exploration function β to obtain an explicit upper bound on $N_\ell(\tau)$.

Lemma 6. *Let $\beta(s) = C + \frac{3}{2} \ln(1 + \ln(s))$ and define $S = \sup\{s \geq 1 : a\beta(s) \geq s\}$. Then*

$$S \leq aC + 2a \ln(1 + \ln(aC)).$$

This result is a consequence of Theorem 16 stated in Appendix F, that uses the fact that for $C \geq -\ln(0.1)$ and $a \geq 8$, it holds that

$$\frac{3}{2} \frac{C(1 + \ln(aC))}{C(1 + \ln(aC)) - \frac{3}{2}} \leq 1.7995564 \leq 2.$$

On the event \mathcal{E} , letting τ_ℓ be the last instant before τ at which the leaf ℓ has been played before stopping, one has $N_\ell(\tau - 1) = N_\ell(\tau_\ell)$ that satisfies by Lemma 5

$$N_\ell(\tau_\ell) \leq \frac{8\beta(N_\ell(\tau_\ell), \delta)}{\Delta_\ell^2 \vee \Delta_*^2 \vee \epsilon^2}.$$

Applying Lemma 6 with $a = a_\ell = \frac{8}{\Delta_\ell^2 \vee \Delta_*^2 \vee \epsilon^2}$ and $C = \ln \frac{|\mathcal{L}|}{\delta} + 3 \ln \ln \frac{|\mathcal{L}|}{\delta}$ leads to

$$N_\ell(\tau - 1) \leq a_\ell (C + 2 \ln(1 + \ln(a_\ell C))).$$

Letting $\bar{\Delta}_{\ell, \epsilon} = \Delta_\ell \vee \Delta_* \vee \epsilon$ and summing over arms, we find

$$\begin{aligned} \tau &= 1 + \sum_{\ell} N_\ell(\tau - 1) \\ &\leq 1 + \sum_{\ell} \frac{8}{\bar{\Delta}_{\ell, \epsilon}^2} \left(\ln \frac{|\mathcal{L}|}{\delta} + 3 \ln \ln \frac{|\mathcal{L}|}{\delta} + 2 \ln \ln \left(8e \frac{\ln \frac{|\mathcal{L}|}{\delta} + 3 \ln \ln \frac{|\mathcal{L}|}{\delta}}{\bar{\Delta}_{\ell, \epsilon}^2} \right) \right) \\ &= 1 + \sum_{\ell} \frac{8}{\bar{\Delta}_{\ell, \epsilon}^2} \left(\ln \frac{|\mathcal{L}|}{\delta} + 2 \ln \ln \frac{1}{\bar{\Delta}_{\ell, \epsilon}^2} \right) + 8H_\epsilon^*(\mu) \left[3 \ln \ln \frac{|\mathcal{L}|}{\delta} + 2 \ln \ln \left(8e \ln \frac{|\mathcal{L}|}{\delta} + 24e \ln \ln \frac{|\mathcal{L}|}{\delta} \right) \right]. \end{aligned}$$

To conclude the proof, we remark that from the proof of Lemma 2 (see Appendix B.2) it follows that on \mathcal{E} , $V(s^*) - V(\hat{s}_\tau) < \epsilon$ and that \mathcal{E} holds with probability larger than $1 - \delta$.

4 Experimental Validation

In this section we evaluate the performance of our algorithms in three experiments. We evaluate on the depth-two benchmark tree from [10], a new depth-three tree and the random tree ensemble from [22]. We compare to the FindTopWinner algorithm from [22] in all experiments, and in the depth-two experiment we include the M-LUCB algorithm from [10]. Its relation to BAI-MCTS is discussed in Section 3.3. For our BAI-MCTS algorithms and for M-LUCB we use the exploration rate $\beta(s, \delta) = \ln \frac{|\mathcal{L}|}{\delta} + \ln(\ln(s) + 1)$ (a stylized version of Lemma 2 that works well in practice), and we use the KL refinement of the confidence intervals (1). To replicate the experiment from [22], we supply all algorithms with $\delta = 0.1$ and $\epsilon = 0.01$. For comparing with [10] we run all algorithms with $\epsilon = 0$ and $\delta = 0.1|\mathcal{L}|$ (undoing the conservative union bound over leaves. This excessive choice, which might even exceed one, does not cause a problem, as the algorithms depend on $\frac{\delta}{|\mathcal{L}|} = 0.1$). In none of our experiments the observed error rate exceeds 0.1.

Figure 3 shows the benchmark tree from [10, Section 5] and the performance of four algorithms on it. We see that the special-purpose depth-two M-LUCB performs best, very closely followed by both our new arbitrary-depth LUCB-MCTS and UGapE-MCTS methods. All three use significantly fewer samples than FindTopWinner. Figure 4 (displayed in Appendix A for the sake of readability) shows a full 3-way tree of depth 3 with leaves drawn uniformly from $[0, 1]$. Again our algorithms outperform the previous state of the art by an order of magnitude. Finally, we replicate the experiment from [22, Section 4]. To make the comparison as fair as possible, we use the proven exploration rate from (2). On 10K full 10-ary trees of depth 3 with Bernoulli leaf parameters drawn uniformly at random from $[0, 1]$ the average numbers of samples are: LUCB-MCTS 141811, UGapE-MCTS 142953 and FindTopWinner 2254560. To closely follow the original experiment, we do apply the union bound over leaves to all algorithms, which are run with $\epsilon = 0.01$ and $\delta = 0.1$. We did not observe any error from any algorithm (even though we allow 10%). Our BAI-MCTS algorithms deliver an impressive 15-fold reduction in samples.




Figure 3: The 3×3 tree of depth 2 that is the benchmark in [10]. Shown below the leaves are the average numbers of pulls for 4 algorithms: LUCB-MCTS (0.89% errors, 2460 samples), UGapE-MCTS (0.94%, 2419), FindTopWinner (0%, 17097) and M-LUCB (0.14%, 2399). All counts are averages over 10K repetitions with $\epsilon = 0$ and $\delta = 0.1 \cdot 9$.

5 Lower bounds and discussion

Given a tree \mathcal{T} , a MCTS model is parameterized by the leaf values, $\boldsymbol{\mu} := (\mu_\ell)_{\ell \in \mathcal{L}}$, which determine the best root action: $s^* = s^*(\boldsymbol{\mu})$. For $\boldsymbol{\mu} \in [0, 1]^{|\mathcal{L}|}$, We define $\text{Alt}(\boldsymbol{\mu}) = \{\boldsymbol{\lambda} \in [0, 1]^{|\mathcal{L}|} : s^*(\boldsymbol{\lambda}) \neq s^*(\boldsymbol{\mu})\}$. Using the same technique as [9] for the classic best arm identification problem, one can establish the following (non explicit) lower bound. The proof is given in Appendix E.

Theorem 7. Assume $\epsilon = 0$. Any δ -correct algorithm satisfies

$$\mathbb{E}_{\boldsymbol{\mu}}[\tau] \geq T^*(\boldsymbol{\mu})d(\delta, 1 - \delta), \text{ where } T^*(\boldsymbol{\mu})^{-1} := \sup_{\boldsymbol{w} \in \Sigma_{|\mathcal{L}|}} \inf_{\boldsymbol{\lambda} \in \text{Alt}(\boldsymbol{\mu})} \sum_{\ell \in \mathcal{L}} w_{\ell} d(\mu_{\ell}, \lambda_{\ell}) \quad (4)$$

with $\Sigma_k = \{\boldsymbol{w} \in [0, 1]^k : \sum_{i=1}^k w_i = 1\}$ and $d(x, y) = x \ln(x/y) + (1 - x) \ln((1 - x)/(1 - y))$ is the binary Kullback-Leibler divergence.

This result is however not directly amenable for comparison with our upper bounds, as the optimization problem defined in Lemma 7 is not easy to solve. Note that $d(\delta, 1 - \delta) \geq \ln(1/(2.4\delta))$ [15], thus our upper bounds have the right dependency in δ . For depth-two trees with K (resp. M) actions for player A (resp. B), we can moreover prove the following result, that suggests an intriguing behavior.

Lemma 8. Assume $\epsilon = 0$ and consider a tree of depth two with $\boldsymbol{\mu} = (\mu_{i,j})_{1 \leq i \leq K, 1 \leq j \leq M}$ such that $\forall (i, j), \mu_{1,1} > \mu_{i,1}, \mu_{i,1} < \mu_{i,j}$. The supremum in the definition of $T^*(\boldsymbol{\mu})^{-1}$ can be restricted to

$$\tilde{\Sigma}_{K,M} := \{\boldsymbol{w} \in \Sigma_{K \times M} : w_{i,j} = 0 \text{ if } i \geq 2 \text{ and } j \geq 2\}$$

and

$$T^*(\boldsymbol{\mu})^{-1} = \max_{\boldsymbol{w} \in \tilde{\Sigma}_{K,M}} \min_{\substack{i=2, \dots, K \\ a=1, \dots, M}} \left[w_{1,a} d\left(\mu_{1,a}, \frac{w_{1,a}\mu_{1,a} + w_{i,1}\mu_{i,1}}{w_{1,a} + w_{i,1}}\right) + w_{i,1} d\left(\mu_{i,1}, \frac{w_{1,a}\mu_{1,a} + w_{i,1}\mu_{i,1}}{w_{1,a} + w_{i,1}}\right) \right].$$

It can be extracted from the proof of Theorem 7 (see Appendix E) that the vector $\boldsymbol{w}^*(\boldsymbol{\mu})$ that attains the supremum in (4) represents the average proportions of selections of leaves by any algorithm matching the lower bound. Hence, the sparsity pattern of Lemma 8 suggests that matching algorithms should draw many of the leaves *much less than* $O(\ln(1/\delta))$ times. This hints at the exciting prospect of optimal stochastic pruning, at least in the asymptotic regime $\delta \rightarrow 0$.

As an example, we numerically solve the lower bound optimization problem (which is a concave maximization problem) for $\boldsymbol{\mu}$ corresponding to the benchmark tree displayed in Figure 3 to obtain

$$T^*(\boldsymbol{\mu}) = 259.9 \quad \text{and} \quad \boldsymbol{w}^* = (0.3633, 0.1057, 0.0532), (0.3738, \mathbf{0}, \mathbf{0}), (0.1040, \mathbf{0}, \mathbf{0}).$$

With $\delta = 0.1$ we find $\text{kl}(\delta, 1 - \delta) = 1.76$ and the lower bound is $\mathbb{E}_{\boldsymbol{\mu}}[\tau] \geq 456.9$. We see that there is a potential improvement of at least a factor 4.

Future directions An (asymptotically) optimal algorithm for BAI called Track-and-Stop was developed by [9]. It maintains the empirical proportions of draws close to $\boldsymbol{w}^*(\hat{\boldsymbol{\mu}})$, adding forced exploration to ensure $\hat{\boldsymbol{\mu}} \rightarrow \boldsymbol{\mu}$. We believe that developing this line of ideas for MCTS would result in a major advance in the quality of tree search algorithms. The main challenge is developing efficient solvers for the general optimization problem (4). For now, even the sparsity pattern revealed by Lemma 8 for depth two does not give rise to efficient solvers. We also do not know how this sparsity pattern evolves for deeper trees, let alone how to compute $\boldsymbol{w}^*(\boldsymbol{\mu})$.

Acknowledgments. Emilie Kaufmann acknowledges the support of the French Agence Nationale de la Recherche (ANR), under grant ANR-16-CE40-0002 (project BADASS). Wouter Koolen acknowledges support from the Netherlands Organization for Scientific Research (NWO) under Veni grant 639.021.439.

References

- [1] J-Y. Audibert, S. Bubeck, and R. Munos. Best Arm Identification in Multi-armed Bandits. In *Proceedings of the 23rd Conference on Learning Theory*, 2010.
- [2] P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2):235–256, 2002.
- [3] L. Borsoniu, R. Munos, and E. Páll. An analysis of optimistic, best-first search for minimax sequential decision making. In *ADPRL14*, 2014.
- [4] C. Browne, E. Powley, D. Whitehouse, S. Lucas, P. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton. A survey of monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in games*, 4(1):1–49, 2012.

- [5] O. Cappé, A. Garivier, O-A. Maillard, R. Munos, and G. Stoltz. Kullback-Leibler upper confidence bounds for optimal sequential allocation. *Annals of Statistics*, 41(3):1516–1541, 2013.
- [6] T. Cazenave. Sequential halving applied to trees. *IEEE Transactions on Computational Intelligence and AI in Games*, 7(1):102–105, 2015.
- [7] E. Even-Dar, S. Mannor, and Y. Mansour. Action Elimination and Stopping Conditions for the Multi-Armed Bandit and Reinforcement Learning Problems. *Journal of Machine Learning Research*, 7:1079–1105, 2006.
- [8] V. Gabillon, M. Ghavamzadeh, and A. Lazaric. Best Arm Identification: A Unified Approach to Fixed Budget and Fixed Confidence. In *Advances in Neural Information Processing Systems*, 2012.
- [9] A. Garivier and E. Kaufmann. Optimal best arm identification with fixed confidence. In *Proceedings of the 29th Conference On Learning Theory (COLT)*, 2016.
- [10] A. Garivier, E. Kaufmann, and W.M. Koolen. Maximin action identification: A new bandit framework for games. In *Proceedings of the 29th Conference On Learning Theory*, 2016.
- [11] Ruitong Huang, Mohammad M. Ajallooeian, Csaba Szepesvári, and Martin Müller. Structured best arm identification with fixed confidence. In *28th International Conference on Algorithmic Learning Theory (ALT)*, 2017.
- [12] K. Jamieson, M. Malloy, R. Nowak, and S. Bubeck. lil'UCB : an Optimal Exploration Algorithm for Multi-Armed Bandits. In *Proceedings of the 27th Conference on Learning Theory*, 2014.
- [13] S. Kalyan Krishnan, A. Tewari, P. Auer, and P. Stone. PAC subset selection in stochastic multi-armed bandits. In *International Conference on Machine Learning (ICML)*, 2012.
- [14] Z. Karnin, T. Koren, and O. Somekh. Almost optimal Exploration in multi-armed bandits. In *International Conference on Machine Learning (ICML)*, 2013.
- [15] E. Kaufmann, O. Cappé, and A. Garivier. On the Complexity of Best Arm Identification in Multi-Armed Bandit Models. *Journal of Machine Learning Research*, 17(1):1–42, 2016.
- [16] E. Kaufmann and S. Kalyan Krishnan. Information complexity in bandit subset selection. In *Proceeding of the 26th Conference On Learning Theory.*, 2013.
- [17] L. Kocsis and C. Szepesvári. Bandit based monte-carlo planning. In *Proceedings of the 17th European Conference on Machine Learning, ECML'06*, pages 282–293, Berlin, Heidelberg, 2006. Springer-Verlag.
- [18] T.L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
- [19] T. Pepels, T. Cazenave, M. Winands, and M. Lanctot. Minimizing simple and cumulative regret in monte-carlo tree search. In *Computer Games Workshop, ECAI*, 2014.
- [20] Aske Plaat, Jonathan Schaeffer, Wim Pijls, and Arie de Bruin. Best-first fixed-depth minimax algorithms. *Artificial Intelligence*, 87(1):255 – 293, 1996.
- [21] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of go with deep neural networks and tree search. *Nature*, 529:484–489, 2016.
- [22] K. Teraoka, K. Hatano, and E. Takimoto. Efficient sampling method for monte carlo tree search problem. *IEICE Transactions on Information and Systems*, pages 392–398, 2014.
- [23] W.R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25:285–294, 1933.

A Numerical Results for a Depth Three Tree

The results for our experiments on a depth-three tree are displayed in Figure 4.

B Confidence Intervals

B.1 Proof of Proposition 1

The proof proceeds by induction. Let the inductive hypothesis be \mathcal{H}_d ="for all the nodes s at (graph) distance d from a leaf, $V_s \in \mathcal{I}_s(t)$ ".

\mathcal{H}_0 clearly holds by definition of \mathcal{E}_t . Now let d such that \mathcal{H}_d holds and let s be at distance $d+1$ of a leaf. Then all $s' \in \mathcal{C}(s)$ are at distance at most d from a leaf and using the inductive hypothesis,

$$\forall c \in \mathcal{C}(s), \quad L_c(t) \leq V_c \leq U_c(t).$$

Assume that s is a MAX node. Using that $U_s(t) = \max_{c \in \mathcal{C}(s)} U_c(t)$, one has $c \in \mathcal{C}(s)$, $V_c \leq U_c(t) \leq U_s(t)$. By definition, $V_s = \max_{c \in \mathcal{C}(s)} V_c$, thus it follows that $V_s \leq U_s(t)$. Still by definition of V_s , it holds that $\forall c \in \mathcal{C}(s)$, $L_c(t) \leq V_c \leq V_s$ and finally, as $L_s(t) = \max_{c \in \mathcal{C}(s)} L_c(t)$, $L_s(t) \leq V_s \leq U_s(t)$. A similar reasoning yields the same conclusion if s is a MIN node, thus \mathcal{H}_{d+1} holds.

As the tree \mathcal{T} is finite, we conclude by induction that $\forall s \in \mathcal{T}$, $V_s \in \mathcal{I}_s(t)$.

B.2 Proof of Lemma 2

Let

$$\mathcal{E}_t = \bigcap_{\ell \in \mathcal{L}} (\mu_\ell \in \mathcal{I}_\ell(t)) \quad \text{and} \quad \mathcal{E} = \bigcap_{t \in \mathbb{N}} \mathcal{E}_t.$$

Using Proposition 1, on \mathcal{E}_t , for all $s \in \mathcal{T}$, $V_s \in \mathcal{I}_s(t)$. If the algorithm stops at some time t , as $L_{\hat{s}_t}(t) > U_{\underline{c}_t}(t) - \epsilon$, the outputted action, $\hat{s}_\tau = \hat{s}_t$, satisfies $L_{\hat{s}_\tau}(t) > U_{s'}(t) - \epsilon$, for all $s' \neq \hat{s}_\tau$. As \mathcal{E} holds, one obtains

$$V(\hat{s}_\tau) \geq \max_{s' \neq \hat{s}_\tau} V(s') - \epsilon,$$

and \hat{s}_τ is an ϵ -maximin action. Hence, the algorithm is correct on \mathcal{E} . The error probability is thus upper bounded by

$$\begin{aligned} \mathbb{P}(\mathcal{E}^c) &\leq \mathbb{P}\left(\exists \ell \in \mathcal{L}, \exists t \in \mathbb{N} : |\hat{\mu}_\ell(t) - \mu_\ell| > \sqrt{\beta(N_\ell(t), \delta)/(2N_\ell(t))}\right) \\ &\leq \sum_{\ell \in \mathcal{L}} \mathbb{P}\left(\exists s \in \mathbb{N} : 2s(\hat{\mu}_{\ell,s} - \mu_\ell)^2 > \beta(s, \delta)\right) \\ &\leq 2|\mathcal{L}|\mathbb{P}\left(\exists s \in \mathbb{N} : S_s > \sqrt{2\sigma^2 s \beta(s, \delta)}\right), \end{aligned}$$

where $S_s = X_1 + \dots + X_s$ is a martingale with σ^2 -subgaussian increments, with $\sigma^2 = 1/4$. It was shown in [15] that for $\delta \leq 0.1$, if

$$\beta(t, \delta) = \ln\left(\frac{1}{\delta}\right) + 3 \ln \ln\left(\frac{1}{\delta}\right) + (3/2) \ln(\ln(s) + 1),$$

one has $\mathbb{P}\left(\exists s \in \mathbb{N} : S_s > \sqrt{2\sigma^2 s \beta(s, \delta)}\right) \leq \delta$, which concludes the proof.

C Sample complexity analysis of LUCB-MCTS

We provide an analysis of a slight variant of LUCB-MCTS that may stop at even rounds only and for $t \in 2\mathbb{N}$ draws the representative leaf of the two promising depth-one nodes:

$$L_{t+1} = \ell_{b_t}(t) \quad \text{and} \quad L_{t+2} = \ell_{\underline{c}_t}(t). \quad (5)$$

The stopping rule is then $\tau = \inf \{t \in 2\mathbb{N}^* : U_{\underline{c}_t}(t) - L_{b_t}(t) < \epsilon\}$.




Figure 4: Our benchmark 3-way tree of depth 3. Shown below the leaves are the numbers of pulls of 3 algorithms: LUCB-MCTS (0.72% errors, 1551 samples), UGapE-MCTS (0.75%, 1584), and FindTopWinner (0%, 20730). Numbers are averages over 10K repetitions with $\epsilon = 0$ and $\delta = 0.1 \cdot 27$.

For this algorithm, our sample complexity guarantee features a slightly different complexity term. For a leaf $\ell = s_0 s_1 \dots s_D$, we first introduce

$$\tilde{\Delta}_\ell = \max_{s \in \text{Anc}(\ell) \setminus \{s_0, s_1\}} |V(s) - V(\mathbb{P}(s))|,$$

a quantity that differs from Δ_ℓ only by the fact that the maximum doesn't take into account the gap between the root and the depth-one ancestor of ℓ . Then $\tilde{H}_\epsilon^*(\boldsymbol{\mu})$ is defined similarly as $H_\epsilon^*(\boldsymbol{\mu})$ by

$$\tilde{H}_\epsilon^*(\boldsymbol{\mu}) = \sum_{\ell \in \mathcal{L}} \frac{1}{\tilde{\Delta}_\ell^2 \vee \Delta_*^2 \vee \epsilon^2}.$$

Theorem 9. *Let $\delta \leq \min(1, 0.1|\mathcal{L}|)$. LUCB-MCTS using the exploration function (2) and selecting the two promising leaves at each round is such that, with probability larger than $1 - \delta$, $(V(s^*) - V(\hat{s}_\tau) < \epsilon)$ and*

$$\tau \leq 16\tilde{H}_\epsilon^*(\boldsymbol{\mu}) \left[\ln \frac{|\mathcal{L}|}{\delta} + 3 \ln \ln \frac{|\mathcal{L}|}{\delta} + 2 \ln \ln \left(16e\tilde{H}_\epsilon^*(\boldsymbol{\mu}) \left(\ln \frac{|\mathcal{L}|}{\delta} + 3 \ln \ln \frac{|\mathcal{L}|}{\delta} \right) \right) \right].$$

Proof. The analysis follows the same lines as that of UGapE-MCTS, yet it relies on a slightly different key result, proved in the next section. Letting $\mathcal{E}_t = \cap_{\ell \in \mathcal{L}} (\mu_\ell \in \mathcal{I}_\ell(t))$ as in the proof of Theorem 3 and defining $\mathcal{E} = \cap_{t \in 2\mathbb{N}^*} \mathcal{E}_t$, one can state the following.

Lemma 10. *Let $t \in 2\mathbb{N}$.*

$$\mathcal{E}_t \cap (\tau > t) \Rightarrow \exists \ell \in \{L_{t+1}, L_{t+2}\} : N_\ell(t) \leq \frac{8\beta(t, \delta)}{\tilde{\Delta}_\ell^2 \vee \Delta_*^2 \vee \epsilon^2}.$$

Let T be a deterministic time. We upper bound τ assuming the event \mathcal{E} holds. Using Lemma 10 and the fact that for every even t , $(\tau_\delta > t) = (\tau_\delta > t + 1)$ by definition of the algorithm, one has

$$\begin{aligned} \min(\tau, T) &= \sum_{t=0}^T \mathbb{1}_{(\tau > t)} = 2 \sum_{\substack{t \in 2\mathbb{N} \\ t \leq T}} \mathbb{1}_{(\tau_\delta > t)} = 2 \sum_{\substack{t \in 2\mathbb{N} \\ t \leq T}} \mathbb{1}_{(\exists \ell \in \{L_{t+1}, L_{t+2}\} : N_\ell(t) \leq 2\beta(t, \delta) / (\tilde{\Delta}_\ell^2 \vee \Delta_*^2 \vee \epsilon^2))} \\ &\leq 2 \sum_{\substack{t \in 2\mathbb{N} \\ t \leq T}} \sum_{\ell \in \mathcal{L}} \mathbb{1}_{(L_{t+1}=\ell) \cup (L_{t+2}=\ell)} \mathbb{1}_{(N_\ell(t) \leq 8\beta(T, \delta) / (\tilde{\Delta}_\ell^2 \vee \Delta_*^2 \vee \epsilon^2))} \\ &\leq 16 \sum_{\ell \in \mathcal{L}} \frac{1}{\tilde{\Delta}_\ell^2 \vee \Delta_*^2 \vee \epsilon^2} \beta(T, \delta) = 16\tilde{H}_\epsilon^*(\boldsymbol{\mu}) \beta(T, \delta). \end{aligned}$$

For any T such that $16\tilde{H}_\epsilon^*(\boldsymbol{\mu}) \beta(T, \delta) < T$, one has $\min(\tau, T) < T$, which implies $\tau < T$. Therefore

$$\tau \leq \sup\{t \in \mathbb{N} : 16\tilde{H}_\epsilon^*(\boldsymbol{\mu}) \beta(t, \delta) \geq t\}.$$

Just like in the analysis of UGapE-MCTS, the conclusion now follows from Lemma 6, applied with $a = 16\tilde{H}_\epsilon^*(\boldsymbol{\mu})$ and $C = \ln(|\mathcal{L}|/\delta) + 3 \ln \ln(|\mathcal{L}|/\delta)$, which yields

$$\tau \leq 16\tilde{H}_\epsilon^*(\boldsymbol{\mu}) \left[\ln \frac{|\mathcal{L}|}{\delta} + 3 \ln \ln \frac{|\mathcal{L}|}{\delta} + 2 \ln \ln \left(16e\tilde{H}_\epsilon^*(\boldsymbol{\mu}) \left(\ln \frac{|\mathcal{L}|}{\delta} + 3 \ln \ln \frac{|\mathcal{L}|}{\delta} \right) \right) \right].$$

Using that $\mathcal{P}(\mathcal{E}) \geq 1 - \delta$ and that the algorithm is correct on \mathcal{E} yields the conclusion.

D Proof of Lemma 5 and Lemma 10

We first state Lemma 12, that holds for both UGapE and LUCB-MCTS and is a consequence of the definition of the exploration procedure. This result builds on the following lemma, that expresses the fact that along a path from the root to a representative leaf, the confidence intervals are nested.

Lemma 11. *Let $t \in \mathbb{N}$ and s_0, s_1, \dots, s_D be a path from the root down to a leaf $\ell = s_D$.*

$$(\ell_{s_1}(t) = s_D) \Rightarrow (\forall k = 2, \dots, D, \mathcal{I}_{s_{k-1}}(t) \subseteq \mathcal{I}_{s_k}(t))$$

Lemma 12. *Let $t \in \mathbb{N}$ and s_0, s_1, \dots, s_D be a path from the root down to a leaf $\ell = s_D$. If \mathcal{E}_t holds and ℓ is selected at round $t + 1$ (UGapE) or if t is even and $\ell \in \{L_{t+1}, L_{t+2}\}$ (LUCB), then*

$$\sqrt{\frac{2\beta(N_\ell(t), \delta)}{N_\ell(t)}} \geq \max_{k=2 \dots D} |V(s_k) - V(s_{k-1})|.$$

UGapE-MCTS: proof of Lemma 5. The following lemma is specific to UGapE-MCTS. We let s_0, s_1, \dots, s_D be a path down to a leaf $\ell = s_D$.

Lemma 13. *Let $t \in \mathbb{N}$. If \mathcal{E}_t holds and UGapE-MCTS has not stopped after t observations, that is $(\tau > t)$,*

$$(L_{t+1} = \ell) \Rightarrow \left(\sqrt{\frac{8\beta(N_\ell(t), \delta)}{N_\ell(t)}} \geq \max(\Delta_*, V(s_0) - V(s_1), \epsilon) \right)$$

Putting together Lemma 12 and Lemma 13 and using that

$$\Delta_\ell = \max\left(V(s_0) - V(s_1), \max_{k=2 \dots D} |V(s_k) - V(s_{k-1})|\right)$$

one obtains

$$\mathcal{E}_t \cap (\tau > t) \cap (L_{t+1} = \ell) \Rightarrow \left(\sqrt{\frac{8\beta(N_\ell(t), \delta)}{N_\ell(t)}} \geq \max(\Delta_\ell, \Delta_*, \epsilon) \right),$$

which yields the proof of Lemma 5 by inverting the bound.

LUCB-MCTS: proof of Lemma 10. The following lemma is specific to the LUCB-MCTS algorithm. It can be viewed as a generalization of Lemma 2 in [13].

Lemma 14. *Let $t \in 2\mathbb{N}$ and let $\gamma \in [V(s_2^*), V(s^*)]$. If \mathcal{E}_t holds and LUCB-MCTS has not stopped after t observations, that is $(\tau > t)$, then*

$$\exists \ell \in \{L_{t+1}, L_{t+2}\} : (\gamma \in \mathcal{I}_\ell(t)) \cap \left(\sqrt{\frac{2\beta(t, \delta)}{N_\ell(t)}} \geq \epsilon \right).$$

Choosing $\gamma = \frac{V(s^*) + V(s_2^*)}{2}$ and letting s_ℓ be the depth-one ancestor of ℓ , on \mathcal{E}_t it holds that $V(s_\ell) \in \mathcal{I}_\ell(t)$ (by Lemma 11) and

$$\begin{aligned} (\gamma \in \mathcal{I}_\ell(t)) &\Rightarrow \left(\sqrt{\frac{2\beta(t, \delta)}{N_\ell(t)}} \geq |V(s_\ell) - \gamma| \right) \\ &\Rightarrow \left(\sqrt{\frac{2\beta(t, \delta)}{N_\ell(t)}} \geq \frac{V(s^*) - V(s_2^*)}{2} \right). \end{aligned}$$

Recall $\Delta_* = V(s^*) - V(s_2^*)$. By Lemma 14, on $\mathcal{E}_t \cap (\tau > t)$, there exists $\ell \in \{L_{t+1}, L_{t+2}\}$ such that

$$\sqrt{\frac{8\beta(t, \delta)}{N_\ell(t)}} \geq \max(\Delta_*, \epsilon). \quad (6)$$

Moreover, noting that for a leaf $\ell = s_0, s_1, \dots, s_D$,

$$\tilde{\Delta}_\ell = \max_{k=2, \dots, D} |V(s_k) - V(s_{k-1})|$$

a consequence of Lemma 12 is that for $\ell \in \{L_{t+1}, L_{t+2}\}$,

$$\sqrt{\frac{2\beta(t, \delta)}{N_\ell(t)}} \geq \tilde{\Delta}_\ell. \quad (7)$$

Combining (6) and (7) yields

$$\mathcal{E}_t \cap (\tau > t) \Rightarrow \exists \ell \in \{L_{t+1}, L_{t+2}\} : \sqrt{\frac{8\beta(t, \delta)}{N_\ell(t)}} \geq \max(\tilde{\Delta}_\ell, \Delta_*, \epsilon),$$

which yields the proof of Lemma 10 by inverting the bound.

D.1 Proof of Lemma 11.

The leaf ℓ is the representative of the depth 1 node s_1 , therefore the path s_1, \dots, s_D is such that $c_{s_{k-1}}(t) = s_k$ for all $k = 2, \dots, D$. Using the way the representative are build, we now show that

$$\forall k \in \{2, \dots, D\}, \quad \mathcal{I}_{s_{k-1}}(t) \subseteq \mathcal{I}_{s_k}(t).$$

If s_{k-1} is a MAX node, $U_{s_{k-1}}(t) = U_{s_k}(t)$ by definition and $L_{s_{k-1}}(t) = \max_{s \in \mathcal{C}(s_{k-1})} L_s(t) \geq L_{s_k}(t)$. Similarly, if s_{k-1} is a MIN node, $L_{s_{k-1}}(t) = L_{s_k}(t)$ by definition and $U_{s_{k-1}}(t) = \min_{s \in \mathcal{C}(s_{k-1})} U_s(t) \leq U_{s_k}(t)$, so that in both cases $\overline{\mathcal{I}}_{s_{k-1}}(t) \subseteq \overline{\mathcal{I}}_{s_k}(t)$.

D.2 Proof of Lemma 12.

Let $\ell \in \mathcal{L}$ be a leaf that is sampled based on the information available at round t . In particular, as ℓ is a representative leaf of the depth 1 node s_1 , the path s_1, \dots, s_D is such that $c_{s_{k-1}}(t) = s_k$ for all $k = 2, \dots, D$. Let $k = 2, \dots, D$. If $s_{k-1} \in \{2, \dots, D\}$ is a MAX node, it holds by definition of the representative children that, for all $s' \in \mathcal{C}(s_{k-1})$,

$$U_{s_k}(t) \geq U_{s'}(t).$$

Now, from Lemma 11 one has $U_\ell(t) \geq U_{s_k}(t)$ and from Proposition 1 as \mathcal{E}_t holds, one has

$$\forall s \in \mathcal{T}, \quad V_s \in \mathcal{I}_s(t). \quad (8)$$

Using these two ingredients yields

$$\begin{aligned} U_\ell(t) &\geq V(s') \\ L_\ell(t) + 2\sqrt{\frac{\beta(N_\ell(t), \delta)}{2N_\ell(t)}} &\geq V(s') \\ L_{s_k}(t) + 2\sqrt{\frac{\beta(N_\ell(t), \delta)}{2N_\ell(t)}} &\geq V(s') \\ V(s_k) + 2\sqrt{\frac{\beta(N_\ell(t), \delta)}{2N_\ell(t)}} &\geq V(s'). \end{aligned}$$

Thus

$$\sqrt{\frac{2\beta(N_\ell(t), \delta)}{N_\ell(t)}} \geq \max_{s' \in \mathcal{C}(s_{k-1})} V(s') - V(s_k) = V(s_{k-1}) - V(s_k) \geq 0.$$

If s_{k-1} is a MIN node, a similar reasoning show that

$$\sqrt{\frac{2\beta(N_\ell(t), \delta)}{N_\ell(t)}} \geq V(s_k) - V(s_{k-1}) \geq 0.$$

Putting everything together yields

$$\sqrt{\frac{2\beta(N_\ell(t), \delta)}{N_\ell(t)}} \geq \max_{k=2, \dots, D} |V(s_k) - V(s_{k-1})|.$$

D.3 Proof of Lemma 13.

We first prove the following intermediate result, that generalizes Lemma 4 in [8].

Lemma 15. *For all $t \in \mathbb{N}^*$, the following holds*

$$\begin{aligned} \text{if } R_{t+1} = \underline{b}_t \quad \text{then} \quad U_{\underline{c}_t}(t) &\leq U_{\underline{b}_t}(t) \\ \text{if } R_{t+1} = \underline{c}_t \quad \text{then} \quad L_{\underline{c}_t}(t) &\leq L_{\underline{b}_t}(t) \end{aligned}$$

Proof. Assume \underline{c}_t is selected (i.e. $R_{t+1} = \underline{c}_t$) and $L_{\underline{c}_t}(t) > L_{\underline{b}_t}(t)$. As the confidence interval on $V(\underline{c}_t)$ is larger than the confidence intervals on $V(\underline{b}_t)$ (\underline{c}_t is selected), this also yields $U_{\underline{c}_t}(t) > U_{\underline{b}_t}(t)$. Hence

$$B_{\underline{b}_t}(t) = U_{\underline{c}_t}(t) - L_{\underline{b}_t}(t) > U_{\underline{b}_t}(t) - L_{\underline{c}_t}(t).$$

Also, by definition of \underline{c}_t , $U_{\underline{c}_t}(t) \geq U_{\underline{b}_t}(t)$. Hence

$$B_{\underline{b}_t}(t) > \max_{\underline{b} \neq \underline{c}_t} U_{\underline{b}}(t) - L_{\underline{c}_t}(t) = B_{\underline{c}_t}(t),$$

which contradicts the definition of \underline{b}_t . Thus, we proved by contradiction that $L_{\underline{c}_t}(t) \geq L_{\underline{b}_t}(t)$.

A similar reasoning can be used to prove that $R_{t+1} = \underline{b}_t \Rightarrow U_{\underline{c}_t}(t) \leq U_{\underline{b}_t}(t)$. □

A simple consequence of Lemma 15 is the fact that, on $\mathcal{E}_t \cap (\tau > t)$,

$$(L_{t+1} = \ell) \Rightarrow \left(\sqrt{\frac{2\beta(N_\ell(t), \delta)}{N_\ell(t)}} > \epsilon \right). \quad (9)$$

Indeed, as the algorithm doesn't stop after t rounds, it holds that $U_{\underline{c}_t}(t) - L_{\underline{b}_t}(t) > \epsilon$. If ℓ is the arm selected at round $t+1$, $\ell = \ell_{R_{t+1}}(t)$ and one can prove using Lemma 15 that $U_{R_{t+1}}(t) - L_{R_{t+1}}(t) > \epsilon$ (by distinguishing two cases). Finally, as \mathcal{E}_t holds, by Lemma 11, $\mathcal{I}_{R_{t+1}}(t) \subseteq \mathcal{I}_\ell(t)$. Hence $U_\ell(t) - L_\ell(t) > \epsilon$, and (9) follows using the particular form of the confidence intervals.

To complete the proof, we now show that

$$(L_{t+1} = \ell) \Rightarrow \left(\sqrt{\frac{8\beta(N_\ell(t), \delta)}{N_\ell(t)}} > \max(\Delta_*, V(s_0) - V(s_1)) \right). \quad (10)$$

by distinguishing several cases.

Case 1: $s^* \notin \text{Anc}(\ell)$ and $R_{t+1} = \underline{c}_t$. Using that the algorithm doesn't stop yields

$$L_{\underline{c}_t}(t) - U_{\underline{b}_t}(t) + 2(U_{\underline{c}_t}(t) - L_{\underline{c}_t}(t)) > \epsilon.$$

As \mathcal{E}_t holds, $L_{\underline{c}_t}(t) \leq V(\underline{c}_t) = V(s_1)$ and $U_{\underline{b}_t}(t) \geq V(\underline{b}_t)$.

Therefore, if $\underline{b}_t = s^*$ it holds that

$$2(U_{\underline{c}_t}(t) - L_{\underline{c}_t}(t)) > V(s^*) - V(s_1) + \epsilon.$$

If $\underline{b}_t \neq s^*$, by definition of \underline{c}_t one has

$$U_{\underline{c}_t}(t) \geq U_{s^*}(t) \geq V(s^*),$$

hence

$$\begin{aligned} L_{\underline{c}_t}(t) + (U_{\underline{c}_t}(t) - L_{\underline{c}_t}(t)) &\geq V(s^*) \\ V(s_1) + (U_{\underline{c}_t}(t) - L_{\underline{c}_t}(t)) &\geq V(s^*). \end{aligned}$$

Thus, recalling that $V(s_0) = V(s^*)$, whatever the value of \underline{b}_t , one obtains

$$2(U_{\underline{c}_t}(t) - L_{\underline{c}_t}(t)) \geq V(s_0) - V(s_1).$$

From Lemma 11 the width of $\mathcal{I}_{\underline{c}_t}(t)$ is upper bounded by the width of $\mathcal{I}_\ell(t)$, hence

$$2(U_\ell(t) - L_\ell(t)) \geq V(s_0) - V(s_1). \quad (11)$$

Case 2: $s^* \notin \text{Anc}(\ell)$ and $R_{t+1} = \underline{b}_t$. As $s^* \neq \underline{b}_t$, by definition of \underline{c}_t one has

$$U_{\underline{c}_t}(t) \geq U_{s^*}(t) \geq V(s^*).$$

Hence, using Lemma 15,

$$U_{\underline{b}_t}(t) - L_{\underline{b}_t}(t) \geq U_{\underline{c}_t}(t) - L_{\underline{b}_t}(t) \geq V(s^*) - V(s_1),$$

as \mathcal{E}_t holds. Finally, by Lemma 11,

$$U_\ell(t) - L_\ell(t) \geq V(s_0) - V(s_1). \quad (12)$$

Case 3: $s^* \in \text{Anc}(\ell)$ and $R_{t+1} = \underline{b}_t$. One has $\underline{b}_t = s^*$. Using that the algorithm doesn't stop yields

$$L_{\underline{c}_t}(t) - U_{\underline{b}_t}(t) + 2(U_{\underline{b}_t}(t) - L_{\underline{b}_t}(t)) > \epsilon.$$

As \mathcal{E}_t holds, $U_{\underline{b}_t}(t) \leq V(s^*)$ and $L_{\underline{c}_t}(t) \leq V(\underline{c}_t) \leq V(s_2^*)$. Therefore, if $\underline{b}_t = s^*$ it holds that

$$2(U_{\underline{b}_t}(t) - L_{\underline{b}_t}(t)) > V(s^*) - V(s_2^*) + \epsilon.$$

and by Lemma 11

$$2(U_\ell(t) - L_\ell(t)) \geq V(s^*) - V(s_2^*). \quad (13)$$

Case 4: $s^* \in \text{Anc}(\ell)$ and $R_{t+1} = \underline{c}_t$. One has $\underline{c}_t = s^*$. Using Lemma 15 yields

$$U_{\underline{c}_t}(t) - L_{\underline{c}_t}(t) \geq U_{\underline{c}_t}(t) - L_{\underline{b}_t}(t) \geq V(s^*) - V(s_2^*),$$

as \mathcal{E}_t holds and $V(\underline{b}_t) \leq V(s_2^*)$. Finally, by Lemma 11,

$$U_\ell(t) - L_\ell(t) \geq V(s^*) - V(s_2^*). \quad (14)$$

Combining (11)-(14), we see that in all four cases

$$2(U_\ell(t) - L_\ell(t)) \geq \max(V(s^*) - V(s_2^*), V(s_0) - V(s_1)),$$

as for $s^* \notin \text{Anc}(\ell)$, $V(s_0) - V(s_1) = V(s^*) - V(s_1) \geq V(s^*) - V(s_2^*)$, and for $s^* \in \text{Anc}(\ell)$, $V(s_0) - V(s_1) = 0$. Using the expression of the confidence intervals and recalling that $\Delta_* = V(s^*) - V(s_2^*)$, one obtains

$$4\sqrt{\frac{\beta(N_\ell(t), \delta)}{2N_\ell(t)}} \geq \max(\Delta_*, V(s_0) - V(s_1))$$

which proves (10).

D.4 Proof of Lemma 14.

Fix $\gamma \in [V(s_2^*), V(s^*)]$ and assume $\mathcal{E}_t \cap (\tau > t)$ holds. We assume (by contradiction) that γ doesn't belong to $\mathcal{I}_{L_{t+1}}(t)$ nor to $\mathcal{I}_{L_{t+2}}(t)$. There are four possibilities:

- $L_{L_{t+1}}(t) > \gamma$ and $L_{L_{t+2}}(t) > \gamma$. As \mathcal{E}_t holds and L_{t+1} and L_{t+2} are representative, it yields that there exists two nodes $s \in \mathcal{C}(s_0)$ such that $V_s > \gamma$, which contradicts the definition of γ .
- $U_{L_{t+1}}(t) < \gamma$ and $U_{L_{t+2}}(t) < \gamma$. From the definition of \underline{c}_{t+1} , it yields that for all $s \in \mathcal{C}(s_0)$, $U_s(t) < \gamma$ and as \mathcal{E}_t holds one obtains $V_s < \gamma$ for all $s \in \mathcal{C}(s_0)$, which contradicts the definition of γ .
- $L_{L_{t+1}}(t) > \gamma$ and $\gamma > U_{L_{t+2}}(t)$. This implies that $L_{L_{t+1}}(t) > U_{L_{t+2}}(t)$ and that $L_{\underline{b}_t}(t) > U_{\underline{c}_t}(t)$ (by Lemma 11 and the fact that L_{t+1} and L_{t+2} are representative leaves). This yields $(\tau \leq t)$ and a contradiction.
- $U_{L_{t+1}}(t) < \gamma$ and $\gamma < L_{L_{t+2}}(t)$. This implies in particular that $\hat{\mu}_{L_{t+1}}(t) < \hat{\mu}_{L_{t+2}}(t)$. Thus $\hat{V}(\underline{b}_t, t) < \hat{V}(\underline{c}_t, t)$, which contradicts the definition of \underline{b}_t .

Hence, we just proved by contradiction that there exists $\ell \in \{L_{t+1}, L_{t+2}\}$ such that $\gamma \in \mathcal{I}_\ell(t)$. To prove Lemma 14, it remains to establish the following three statements.

1. $(\gamma \in \mathcal{I}_{L_{t+1}}(t)) \cap (\gamma \in \mathcal{I}_{L_{t+2}}(t)) \Rightarrow \left(\exists \ell \in \{L_{t+1}, L_{t+2}\} : \sqrt{\frac{2\beta(t, \delta)}{N_\ell(t)}} > \epsilon \right)$
2. $(\gamma \in \mathcal{I}_{L_{t+1}}(t)) \cap (\gamma \notin \mathcal{I}_{L_{t+2}}(t)) \Rightarrow \left(\sqrt{\frac{2\beta(t, \delta)}{N_{L_{t+1}}(t)}} > \epsilon \right)$
3. $(\gamma \notin \mathcal{I}_{L_{t+1}}(t)) \cap (\gamma \in \mathcal{I}_{L_{t+2}}(t)) \Rightarrow \left(\sqrt{\frac{2\beta(t, \delta)}{N_{L_{t+2}}(t)}} > \epsilon \right)$

Statement 1. As the algorithm doesn't stop, $U_{\underline{c}_t}(t) - L_{\underline{b}_t}(t) > \epsilon$. Hence

$$\begin{aligned} U_{L_{t+2}}(t) - L_{L_{t+1}}(t) &> \epsilon \\ \hat{\mu}_{L_{t+2}}(t) + \sqrt{\frac{\beta(N_{L_{t+2}}(t), \delta)}{2N_{L_{t+2}}(t)}} - \hat{\mu}_{L_{t+1}}(t) + \sqrt{\frac{\beta(N_{L_{t+1}}(t), \delta)}{2N_{L_{t+1}}(t)}} &> \epsilon \\ \sqrt{\frac{\beta(N_{L_{t+2}}(t), \delta)}{2N_{L_{t+2}}(t)}} + \sqrt{\frac{\beta(N_{L_{t+1}}(t), \delta)}{2N_{L_{t+1}}(t)}} &> \epsilon \end{aligned}$$

using that by definition of \underline{b}_t , $\hat{\mu}_{L_{t+2}}(t) < \hat{\mu}_{L_{t+1}}(t)$. Hence, there exists $\ell \in \{L_{t+1}, L_{t+2}\}$ such that

$$\sqrt{\frac{\beta(N_\ell(t), \delta)}{2N_\ell(t)}} > \frac{\epsilon}{2} \Rightarrow \sqrt{\frac{2\beta(t, \delta)}{N_\ell(t)}} > \epsilon.$$

Statement 2. We consider two cases and first assume that $\gamma \in \mathcal{I}_{L_{t+1}}(t)$ and $\gamma \geq U_{L_{t+2}}(t)$. Using the fact that the algorithm doesn't stop at round t , the following events hold

$$\begin{aligned} &(U_{L_{t+2}}(t) - L_{L_{t+1}}(t) > \epsilon) \cap (U_{L_{t+1}}(t) > \gamma) \cap (U_{L_{t+2}}(t) \leq \gamma) \\ \Rightarrow &(U_{L_{t+2}}(t) - L_{L_{t+1}}(t) > \epsilon) \cap (U_{L_{t+1}}(t) > \gamma) \cap (U_{L_{t+2}}(t) - L_{L_{t+1}}(t) + L_{L_{t+1}}(t) \leq \gamma) \\ \Rightarrow &(U_{L_{t+1}}(t) > \gamma) \cap (L_{L_{t+1}}(t) \leq \gamma - \epsilon) \\ \Rightarrow &(U_{L_{t+1}}(t) - L_{L_{t+1}}(t) > \epsilon) \\ \Rightarrow &\left(\sqrt{\frac{2\beta(t, \delta)}{N_{L_{t+1}}(t)}} > \epsilon \right). \end{aligned}$$

The second case is $\gamma \in \mathcal{I}_{L_{t+1}}(t)$ and $\gamma \leq L_{L_{t+2}}(t)$. Then the following holds

$$\begin{aligned} &(U_{L_{t+2}}(t) - L_{L_{t+1}}(t) > \epsilon) \cap (L_{L_{t+1}}(t) \leq \gamma) \cap (L_{L_{t+2}}(t) \geq \gamma) \\ \Rightarrow &(L_{L_{t+1}}(t) \leq \gamma) \cap (U_{L_{t+2}}(t) + L_{L_{t+2}}(t) - L_{L_{t+1}}(t) > \gamma + \epsilon) \\ \Rightarrow &(L_{L_{t+1}}(t) \leq \gamma) \cap \left(2\hat{\mu}_{L_{t+2}}(t) - \hat{\mu}_{L_{t+1}}(t) + \sqrt{\frac{2\beta(N_{L_{t+1}}(t), \delta)}{2N_{L_{t+1}}(t)}} > \gamma + \epsilon \right) \\ \Rightarrow &(L_{L_{t+1}}(t) \leq \gamma) \cap \left(\hat{\mu}_{L_{t+1}}(t) + \sqrt{\frac{2\beta(N_{L_{t+1}}(t), \delta)}{2N_{L_{t+1}}(t)}} > \gamma + \epsilon \right) \\ \Rightarrow &(U_{L_{t+1}}(t) - L_{L_{t+1}}(t) > \epsilon), \end{aligned}$$

where the third implication uses the fact that $\hat{\mu}_{L_{t+2}}(t) \leq \hat{\mu}_{L_{t+1}}(t)$.

Statement 3. We consider two cases and first assume that $\gamma \in \mathcal{I}_{L_{t+2}}(t)$ and $\gamma \leq L_{L_{t+1}}(t)$. Using the fact that the algorithm doesn't stop at round t , the following events hold

$$\begin{aligned} &(U_{L_{t+2}}(t) - L_{L_{t+1}}(t) > \epsilon) \cap (L_{L_{t+1}}(t) \geq \gamma) \cap (L_{L_{t+2}}(t) \leq \gamma) \\ \Rightarrow &(U_{L_{t+2}}(t) > \gamma + \epsilon) \cap (L_{L_{t+2}}(t) \leq \gamma) \\ \Rightarrow &(U_{L_{t+2}}(t) - L_{L_{t+2}}(t) > \epsilon) \\ \Rightarrow &\left(\sqrt{\frac{2\beta(t, \delta)}{N_{L_{t+2}}(t)}} > \epsilon \right). \end{aligned}$$

The second case is $\gamma \in \mathcal{I}_{L_{t+2}}(t)$ and $\gamma \geq U_{L_{t+1}}(t)$. Then the following holds

$$\begin{aligned}
& (U_{L_{t+2}}(t) - L_{L_{t+1}}(t) > \epsilon) \cap (U_{L_{t+2}}(t) \geq \gamma) \cap (\gamma \geq U_{L_{t+1}}(t)) \\
\Rightarrow & (U_{L_{t+2}}(t) - (L_{L_{t+1}}(t) + U_{L_{t+1}}(t)) > \epsilon - \gamma) \cap (U_{L_{t+2}}(t) \geq \gamma) \\
\Rightarrow & \left(\hat{\mu}_{L_{t+2}}(t) + \sqrt{\frac{2\beta(N_{L_{t+2}}(t), \delta)}{2N_{L_{t+2}}(t)}} - 2\hat{\mu}_{L_{t+1}}(t) > \epsilon - \gamma \right) \cap (U_{L_{t+2}}(t) \geq \gamma) \\
\Rightarrow & \left(-\hat{\mu}_{L_{t+2}}(t) + \sqrt{\frac{2\beta(N_{L_{t+2}}(t), \delta)}{2N_{L_{t+2}}(t)}} > \epsilon - \gamma \right) \cap (U_{L_{t+2}}(t) \geq \gamma) \\
\Rightarrow & (U_{L_{t+2}}(t) - L_{L_{t+2}}(t) > \epsilon),
\end{aligned}$$

where the third implication uses the fact that $\hat{\mu}_{L_{t+2}}(t) \leq \hat{\mu}_{L_{t+1}}(t)$.

E Proof of the lower bounds

E.1 Proof of Theorem 7

Theorem 7 follows from considering the best possible change of distribution $\lambda \in \text{Alt}(\mu)$. The expected log-likelihood ratio of the observations until τ under a model parameterized by μ and a model parameterized by λ is

$$\mathbb{E}_{\mu}[L_{\tau}(\mu, \lambda)] = \sum_{\ell \in \mathcal{L}} \mathbb{E}_{\mu}[N_{\ell}(\tau)] d(\mu_{\ell}, \lambda_{\ell}),$$

where $N_{\ell}(t)$ is the number of draws of the leaf ℓ until round t . Using Lemma 1 of [15], for any event \mathcal{E} in the filtration generated by τ ,

$$\mathbb{E}_{\mu}[L_{\tau}(\mu, \lambda)] \geq d(\mathbb{P}_{\mu}(\mathcal{E}), \mathbb{P}_{\lambda}(\mathcal{E})).$$

As the strategy is δ -correct, letting $\mathcal{E} = (\hat{s}_{\tau} = s^*(\mu))$ one has $\mathbb{P}_{\mu}(\mathcal{E}) \geq 1 - \delta$ and $\mathbb{P}_{\lambda}(\mathcal{E}) \leq \delta$ (under this model, $s^*(\mu)$ is not the best action at the root under the model parameterized by λ). Using monotonicity properties of the Bernoulli KL-divergence, one obtains, for any $\lambda \in \text{Alt}(\mu)$,

$$\sum_{\ell \in \mathcal{L}} \mathbb{E}_{\mu}[N_{\ell}(\tau)] d(\mu_{\ell}, \lambda_{\ell}) \geq d(1 - \delta, \delta).$$

Then, one can write

$$\begin{aligned}
& \inf_{\lambda \in \text{Alt}(\mu)} \sum_{\ell \in \mathcal{L}} \mathbb{E}_{\mu}[N_{\ell}(\tau)] d(\mu_{\ell}, \lambda_{\ell}) \geq d(1 - \delta, \delta) \\
& \mathbb{E}_{\mu}[\tau] \inf_{\lambda \in \text{Alt}(\mu)} \sum_{\ell \in \mathcal{L}} \frac{\mathbb{E}_{\mu}[N_{\ell}(\tau)]}{\mathbb{E}_{\mu}[\tau]} d(\mu_{\ell}, \lambda_{\ell}) \geq d(1 - \delta, \delta) \\
& \mathbb{E}_{\mu}[\tau] \left(\sup_{w \in \Sigma_{|\mathcal{L}|}} \inf_{\lambda \in \text{Alt}(\mu)} \sum_{\ell \in \mathcal{L}} \frac{\mathbb{E}_{\mu}[N_{\ell}(\tau)]}{\mathbb{E}_{\mu}[\tau]} d(\mu_{\ell}, \lambda_{\ell}) \right) \geq d(1 - \delta, \delta),
\end{aligned}$$

using that $\sum_{\ell \in \mathcal{L}} \frac{\mathbb{E}_{\mu}[N_{\ell}(\tau)]}{\mathbb{E}_{\mu}[\tau]} = 1$. This concludes the proof.

One can also note that for an algorithm to match the lower bound, all the inequalities above should be equalities. In particular one would need $w_{\ell}^*(\mu) \simeq \frac{\mathbb{E}_{\mu}[N_{\ell}(\tau)]}{\mathbb{E}_{\mu}[\tau]}$, where $w_{\ell}^*(\mu)$ is a maximizer in the definition of $T^*(\mu)^{-1}$ in (4).

E.2 Proof of Lemma 8

In the particular case of a depth-two tree with K actions for player A and M actions for player B,

$$T^*(\mu)^{-1} = \sup_{w \in \Sigma_{K \times M}} \inf_{\lambda \in \text{Alt}(\mu)} \sum_{k=1}^K \sum_{m=1}^M w_{k,m} d(\mu_{k,m}, \lambda_{k,m}).$$

From the particular structure of μ , the best action at the root is action $i = 1$. Hence

$$\text{Alt}(\mu) = \{\lambda : \exists a \in \{1, \dots, M\}, \exists i \in \{2, \dots, K\} : \forall j \in \{1, \dots, M\}, \lambda_{1,a} < \lambda_{i,j}\}.$$

It follows that

$$\begin{aligned}
T^*(\boldsymbol{\mu})^{-1} &= \sup_{\boldsymbol{w} \in \Sigma_{K \times M}} \min_{\substack{a \in \{1, \dots, M\} \\ i \in \{2, \dots, K\}}} \inf_{\boldsymbol{\lambda}: \forall j, \lambda_{1,a} < \lambda_{i,j}} \sum_{k=1}^K \sum_{m=1}^M w_{k,m} d(\mu_{k,m}, \lambda_{k,m}) \\
&= \sup_{\boldsymbol{w} \in \Sigma_{K \times M}} \min_{\substack{a \in \{1, \dots, M\} \\ i \in \{2, \dots, K\}}} \underbrace{\inf_{\boldsymbol{\lambda}: \forall j, \lambda_{1,a} < \lambda_{i,j}} \left[w_{1,a} d(\mu_{1,a}, \lambda_{1,a}) + \sum_{j=1}^M w_{i,j} d(\mu_{i,j}, \lambda_{i,j}) \right]}_{:= F_{a,i}(\boldsymbol{\mu}, \boldsymbol{w})}. \quad (15)
\end{aligned}$$

Indeed, in the rightmost constrained minimization problem, all the $\lambda_{k,m}$ on which no constraint lie can be set to $\mu_{k,m}$ to minimize the corresponding term in the sum. Using tools from constrained optimization, one can prove that

$$F_{a,i}(\boldsymbol{\mu}, \boldsymbol{w}) = \inf_{(\lambda_{1,a}, (\lambda_{i,j})_j) \in \mathcal{C}} \left[w_{1,a} d(\mu_{1,a}, \lambda_{1,a}) + \sum_{j=1}^M w_{i,j} d(\mu_{i,j}, \lambda_{i,j}) \right],$$

where \mathcal{C} is the set

$$\mathcal{C} = \{(\boldsymbol{\mu}', \boldsymbol{\mu}') \in [0, 1]^{M+1} : \exists j_0 \in \{1, \dots, K\}, \exists c \in [\mu_{i,j_0}, \mu_{1,a}] : \boldsymbol{\mu}' = \boldsymbol{\mu}'_1 = \dots = \boldsymbol{\mu}'_{j_0} = c \text{ and } \boldsymbol{\mu}'_j = \mu_{i,j} \text{ for } j > j_0\}.$$

Letting $H(\boldsymbol{\mu}', \boldsymbol{\mu}', \boldsymbol{w}, \boldsymbol{w}) = w_{1,a} d(\mu_{1,a}, \boldsymbol{\mu}') + \sum_{j=1}^M w_{i,j} d(\mu_{i,j}, \boldsymbol{\mu}'_j)$ one can easily show that for all $(\boldsymbol{\mu}', \boldsymbol{\mu}') \in \mathcal{C}$,

$$H(\boldsymbol{\mu}', \boldsymbol{\mu}', \boldsymbol{w}, \boldsymbol{w}) \leq H(\boldsymbol{\mu}', \boldsymbol{\mu}', \boldsymbol{w}, \tilde{\boldsymbol{w}}),$$

where $\tilde{\boldsymbol{w}}$ is constructed from \boldsymbol{w} by putting all the weight on the smallest arm:

$$\tilde{w}_{i,1} = \sum_{j \geq 1} w_{1,j} \text{ and } \tilde{w}_{i,j} = 0 \text{ for } j \geq 2.$$

This is because the largest $d(\mu_{i,j}, c)$ is $d(\mu_{i,1}, c)$ as $\mu_{i,1} \leq \mu_{i,j} \leq c$ for $j \leq j_0$. Hence taking the infimum, one obtains

$$F_{a,i}(\boldsymbol{\mu}, \boldsymbol{w}) \leq F_{a,i}(\boldsymbol{\mu}, \tilde{\boldsymbol{w}}).$$

Repeating this argument for all i , one can construct $\tilde{\boldsymbol{w}}$ such that

$$\forall i \geq 2, \tilde{w}_{i,1} = \sum_{j \geq 1} w_{i,j} \text{ and } \tilde{w}_{i,j} = 0 \text{ for } j \geq 2$$

and $F_{a,i}(\boldsymbol{\mu}, \boldsymbol{w}) \leq F_{a,i}(\boldsymbol{\mu}, \tilde{\boldsymbol{w}})$ for all a, i . Thus, the supremum in 15 is necessarily attained for \boldsymbol{w} in the set $\tilde{\Sigma}_{K \times M} = \{\boldsymbol{w} \in \Sigma_{K, M} : w_{i,j} = 0 \text{ for } i \geq 2, j \geq 2\}$. It follows that

$$\begin{aligned}
T^*(\boldsymbol{\mu})^{-1} &= \sup_{\boldsymbol{w} \in \tilde{\Sigma}_{M+K-1}} \min_{\substack{a \in \{1, \dots, M\} \\ i \in \{2, \dots, K\}}} \inf_{\boldsymbol{\lambda}: \forall j, \lambda_{1,a} < \lambda_{i,1}} [w_{1,a} d(\mu_{1,a}, \lambda_{1,a}) + w_{i,1} d(\mu_{i,1}, \lambda_{i,1})] \\
&= \sup_{\boldsymbol{w} \in \tilde{\Sigma}_{K+M-1}} \min_{\substack{a=1, \dots, M \\ i=2, \dots, K}} \left[w_{1,a} d\left(\mu_{1,a}, \frac{w_{1,a} \mu_{1,a} + w_{i,1} \mu_{i,1}}{w_{1,a} + w_{i,1}}\right) + w_{i,1} d\left(\mu_{i,1}, \frac{w_{1,a} \mu_{1,a} + w_{i,1} \mu_{i,1}}{w_{1,a} + w_{i,1}}\right) \right],
\end{aligned}$$

which concludes the proof.

F Inverting Bounds

Consider the exploration rate

$$\beta(s) = C + \frac{3}{2} \ln(1 + \ln s) \quad \text{where} \quad C = \ln \frac{|\mathcal{L}|}{\delta} + 3 \ln \ln \frac{|\mathcal{L}|}{\delta}$$

where we assume $C \geq 1$, so that $\beta(1) = C \geq 1$. Now fix some $a \geq 1$, and let us define

$$S = \sup \{s \geq 1 | a\beta(s) \geq s\}.$$

The goal is to get a tight upper bound on S . We claim that

Theorem 16.

$$aC + \frac{3}{2}a \ln(1 + \ln(aC)) \leq S \leq aC + \frac{3}{2}a \ln(1 + \ln(aC)) \underbrace{\frac{C(1 + \ln(aC))}{C(1 + \ln(aC)) - \frac{3}{2}}}_{\rightarrow 1},$$

where the upper bound is only non-trivial if $C(1 + \ln(aC)) > \frac{3}{2}$, which, for example, is implied by $C > 1.23696$.

Proof. The requirement $a\beta(s) \geq s$ is monotone in s , in that it holds for small s (including $s = 1$) and fails for large s . So to show the theorem it suffices to show that $a\beta(s) \geq s$ holds at s given by the left-hand-side, while it fails for s equal to the right-hand side.

First, we need to establish that $a\beta(s) \geq s$ for s equal to the left-hand side expression of the theorem. That is, we need to show

$$a \left(C + \frac{3}{2} \ln \left(1 + \ln \left(a \left(C + \frac{3}{2} \ln(1 + \ln(aC)) \right) \right) \right) \right) \geq a \left(C + \frac{3}{2} \ln(1 + \ln(aC)) \right)$$

which is equivalent (the simplification is entirely mechanical) to

$$aC \geq 1$$

which holds by assumption. Then we need to plug in the right hand side. Here we need to show that

$$\begin{aligned} a \left(C + \frac{3}{2} \ln \left(1 + \ln \left(aC + \frac{3}{2} a \ln(1 + \ln(aC)) \frac{C(1 + \ln(aC))}{C(1 + \ln(aC)) - \frac{3}{2}} \right) \right) \right) \\ \leq aC + \frac{3}{2} a \ln(1 + \ln(aC)) \frac{C(1 + \ln(aC))}{C(1 + \ln(aC)) - \frac{3}{2}}, \end{aligned}$$

that is

$$\ln \left(1 + \ln \left(aC + \frac{3}{2} a \ln(1 + \ln(aC)) \frac{C(1 + \ln(aC))}{C(1 + \ln(aC)) - \frac{3}{2}} \right) \right) \leq \ln(1 + \ln(aC)) \frac{C(1 + \ln(aC))}{C(1 + \ln(aC)) - \frac{3}{2}}.$$

To show this, we use $\ln(1 + x) \leq x$ twice to show

$$\begin{aligned} & \ln \left(1 + \ln \left(aC + \frac{3}{2} a \ln(1 + \ln(aC)) \frac{C(1 + \ln(aC))}{C(1 + \ln(aC)) - \frac{3}{2}} \right) \right) \\ &= \ln \left(1 + \ln(aC) + \ln \left(1 + \frac{3}{2} \ln(1 + \ln(aC)) \frac{(1 + \ln(aC))}{C(1 + \ln(aC)) - \frac{3}{2}} \right) \right) \\ &\leq \ln \left(1 + \ln(aC) + \frac{3}{2} \ln(1 + \ln(aC)) \frac{(1 + \ln(aC))}{C(1 + \ln(aC)) - \frac{3}{2}} \right) \\ &= \ln(1 + \ln(aC)) + \ln \left(1 + \frac{\frac{3}{2} \ln(1 + \ln(aC))}{C(1 + \ln(aC)) - \frac{3}{2}} \right) \\ &\leq \ln(1 + \ln(aC)) + \frac{\frac{3}{2} \ln(1 + \ln(aC))}{C(1 + \ln(aC)) - \frac{3}{2}} \\ &= \frac{C(1 + \ln(aC))}{C(1 + \ln(aC)) - \frac{3}{2}} \ln(1 + \ln(aC)) \end{aligned}$$

as desired. □