


Monte Carlo Tree Search Guided by Symbolic Advice for MDPs

Damien Busatto-Gaston 

Université Libre de Bruxelles, Brussels, Belgium
damien.busatto-gaston@ulb.ac.be

Debraj Chakraborty 

Université Libre de Bruxelles, Brussels, Belgium
debraj.chakraborty@ulb.ac.be

Jean-Francois Raskin

Université Libre de Bruxelles, Brussels, Belgium
jraskin@ulb.ac.be

Abstract

In this paper, we consider the online computation of a strategy that aims at optimizing the expected average reward in a Markov decision process. The strategy is computed with a receding horizon and using Monte Carlo tree search (MCTS). We augment the MCTS algorithm with the notion of symbolic advice, and show that its classical theoretical guarantees are maintained. Symbolic advice are used to bias the selection and simulation strategies of MCTS. We describe how to use QBF and SAT solvers to implement symbolic advice in an efficient way. We illustrate our new algorithm using the popular game PAC-MAN and show that the performances of our algorithm exceed those of plain MCTS as well as the performances of human players.

2012 ACM Subject Classification Theory of computation → Theory of randomized search heuristics; Theory of computation → Convergence and learning in games

Keywords and phrases Markov decision process, Monte Carlo tree search, symbolic advice, simulation

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2020.40

Supplementary Material <http://di.ulb.ac.be/verif/debraj/pacman/>

Funding This work is partially supported by the ARC project Non-Zero Sum Game Graphs: Applications to Reactive Synthesis and Beyond (Fédération Wallonie-Bruxelles), the EOS project Verifying Learning Artificial Intelligence Systems (F.R.S.-FNRS & FWO), and the COST Action 16228 GAMENET (European Cooperation in Science and Technology).

Acknowledgements Computational resources have been provided by the Consortium des Équipements de Calcul Intensif (CÉCI), funded by the Fonds de la Recherche Scientifique de Belgique (F.R.S.-FNRS) under Grant No. 2.5020.11.

1 Introduction

Markov decision processes (MDP) are an important mathematical formalism for modeling and solving sequential decision problems in stochastic environments [23]. The importance of this model has triggered a large number of works in different research communities within computer science, most notably in formal verification, and in artificial intelligence and machine learning. The works done in these research communities have respective weaknesses and complementary strengths. On the one hand, algorithms developed in formal verification are generally complete and provide strong guarantees on the optimality of computed solutions but they tend to be applicable to models of moderate size only. On the other hand, algorithms developed in artificial intelligence and machine learning usually scale to larger models but only provide weaker guarantees. Instead of opposing the two sets of algorithms, there have



© Damien Busatto-Gaston, Debraj Chakraborty, and Jean-Francois Raskin;
licensed under Creative Commons License CC-BY

31st International Conference on Concurrency Theory (CONCUR 2020).

Editors: Igor Konnov and Laura Kovács; Article No. 40; pp. 40:1–40:24

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

been recent works [2, 14, 6, 12, 11, 19, 1] that try to combine the strengths of the two approaches in order to offer new hybrid algorithms that scale better and provide stronger guarantees. The contributions described in this paper are part of this research agenda: we show how to integrate *symbolic advice* defined by formal specifications into Monte Carlo Tree Search algorithms [7] using techniques such as SAT [21] and QBF [24].


When an MDP is too large to be analyzed *offline* using verification algorithms, receding horizon analysis combined with simulation techniques are used *online* [17]. Receding horizon techniques work as follows. In the current state s of the MDP, for a fixed horizon H , the receding horizon algorithm searches for an action a that is the first action of a plan to act (almost) optimally on the finite horizon H . When such an action is identified, then it is played from s and the state evolves stochastically to a new state s' according to the dynamics specified by the MDP. The same process is repeated from s' . The optimization criterion over the H next step depends on the long run measure that needs to be optimised. The tree unfolding from s that needs to be analyzed is often very large (*e.g.* it may be exponential in H). As a consequence, receding horizon techniques are often coupled with *sampling techniques* that avoid the systematic exploration of the entire tree unfolding at the expense of approximation. The Monte Carlo Tree Search (MCTS) algorithm [7] is an increasingly popular tree search algorithm that implements these ideas. It is one of the core building blocks of the ALPHAGO algorithm [25].

While MCTS techniques may offer reasonable performances out of the shelf, they usually need substantial adjustments that depend on the application to really perform well. One way to adapt MCTS to a particular application is to *bias* the search towards promising subspaces taking into account properties of the application domain [16, 26]. This is usually done by coding directly handcrafted search and sampling strategies. We show in this paper how to use techniques from formal verification to offer a *flexible* and *rigorous* framework to bias the search performed by MCTS using *symbolic advice*. A symbolic advice is a *formal specification*, that can be expressed for example in your favorite linear temporal logic, and which constrain the search and the sampling phases of the MCTS algorithm using QBF and SAT solvers. Our framework offers in principle the ability to easily experiment with precisely formulated bias expressed declaratively using logic.

Contributions. On the theoretical side, we study the impact of using symbolic advice on the guarantees offered by MCTS. We identify sufficient conditions for the symbolic advice to preserve the convergence guarantees of the MCTS algorithm (Theorem 21). Those results are partly based on an analysis of the incidence of sampling on those guarantees (Theorem 11) which can be of independent interest.

On a more practical side, we show how symbolic advice can be implemented using SAT and QBF techniques. More precisely, we use QBF [22] to force that all the prefixes explored by the MCTS algorithm in the partial tree unfolding have the property suggested by the *selection advice* (whenever possible) and we use SAT-based sampling techniques [9] to achieve uniform sampling among paths of the MDP that satisfy the *sampling advice*. The use of this symbolic exploration techniques is important as the underlying state space that we need to analyze is usually huge (*e.g.* exponential in the receding horizon H).

To demonstrate the practical interest of our techniques, we have applied our new *MCTS with symbolic advice algorithm* to play PAC-MAN. Figure 1 shows a grid of the PAC-MAN game. In this version of the classical game, the agent Pac-Man has to eat food pills as fast as possible while avoiding being pinched by ghosts. We have chosen this benchmark to evaluate our algorithm for several reasons. First, the state space of the underlying MDP is



■ **Figure 1** We used two grids of size 9×21 and 27×28 for our experiments. Pac-Man loses if he makes contact with a ghost, and wins if he eats all food pills (in white). The agents can travel in four directions unless they are blocked by the walls in the grid, and ghosts cannot reverse their direction. The score decreases by 1 at each step, and increases by 10 whenever Pac-Man eats a food pill. A win (resp. loss), increases (resp. decreases) the score by 500. The game can be seen as an infinite duration game by saying that whenever Pac-Man wins or loses, the positions of the agents and of the food pills are reset.

way too large for the state-of-the-art implementations of complete algorithms. Indeed, the reachable state space of the small grid shown here has approximately 10^{16} states, while the classical grid has approximately 10^{23} states. Our algorithm can handle both grids. Second, this application not only allows for comparison between performances obtained from several versions of the MCTS algorithm but also with the performances that humans can attain in this game. In the PAC-MAN benchmark, we show that advice that instructs Pac-Man on the one hand to avoid ghosts at all costs during the selection phase of the MCTS algorithm (enforced whenever possible by QBF) and on the other hand to bias the search to paths in which ghosts are avoided (using uniform sampling based on SAT) allow to attain or surpass human level performances while the standard MCTS algorithm performs much worse.

Related works. Our analysis of the convergence of the MCTS algorithm with appropriate symbolic advice is based on extensions of analysis results based on bias defined using UCT (bandit algorithms) [18, 3]. Those results are also related to sampling techniques for finite horizon objectives in MDP [17].

Our concept of selection phase advice is related to the MCTS-minimax hybrid algorithm proposed in [4]. There the selection phase advice is not specified declaratively using logic but encoded directly in the code of the search strategy. No use of QBF nor SAT is advocated there and no use of sampling advice either. In [1], the authors provide a general framework to add safety properties to reinforcement learning algorithms via *shielding*. These techniques analyse statically the full state space of the game in order to compute a set of unsafe actions to avoid. This fits our advice framework, so that such a shield could be used as an online selection advice in order to combine their safety guarantees with our formal results for MCTS. More recently, a variation of shielding called *safe padding* has been studied in [14]. Both works are concerned with reinforcement learning and not with MCTS. Note that in general multiple ghosts may prevent the existence of a strategy to enforce safety, *i.e.* always avoid pincer moves.

Our practical handling of symbolic sampling advice relies on symbolic sampling techniques introduced in [8], while our handling of symbolic selection advice relies on natural encodings via QBF that are similar to those defined in [22].

2 Preliminaries

A *probability distribution* on a finite set S is a function $d : S \rightarrow [0, 1]$ such that $\sum_{s \in S} d(s) = 1$. We denote the set of all probability distributions on set S by $\mathcal{D}(S)$. The support of a distribution $d \in \mathcal{D}(S)$ is $\text{Supp}(d) = \{s \in S \mid d(s) > 0\}$.

2.1 Markov decision process

► **Definition 1** (MDP). A *Markov decision process* is a tuple $M = (S, A, P, R, R_T)$, where S is a finite set of states, A is a finite set of actions, P is a mapping from $S \times A$ to $\mathcal{D}(S)$ such that $P(s, a)(s')$ denotes the probability that action a in state s leads to state s' , $R : S \times A \rightarrow \mathbb{R}$ defines the reward obtained for taking a given action at a given state, and $R_T : S \rightarrow \mathbb{R}$ assigns a terminal reward to each state in S .¹

For a Markov decision process M , a *path* of length $i > 0$ is a sequence of i consecutive states and actions followed by a last state. We say that $p = s_0 a_0 s_1 \dots s_i$ is an i -length path in the MDP M if for all $t \in [0, i - 1]$, $a_t \in A$ and $s_{t+1} \in \text{Supp}(P(s_t, a_t))$, and we denote $\text{last}(p) = s_i$ and $\text{first}(p) = s_0$. We also consider states to be paths of length 0. An infinite path is an infinite sequence $p = s_0 a_0 s_1 \dots$ of states and actions such that for all $t \in \mathbb{N}$, $a_t \in A$ and $s_{t+1} \in \text{Supp}(P(s_t, a_t))$. We denote the finite prefix of length t of a finite or infinite path $p = s_0 a_0 s_1 \dots$ by $p|_t = s_0 a_0 \dots s_t$. Let $p = s_0 a_0 s_1 \dots s_i$ and $p' = s'_0 a'_0 s'_1 \dots s'_j$ be two paths such that $s_i = s'_0$, let a be an action and s be state of M . Then, $p \cdot p'$ denotes $s_0 a_0 s_1 \dots s_i a'_0 s'_1 \dots s'_j$ and $p \cdot a s$ denotes $s_0 a_0 s_1 \dots s_i a s$.

For an MDP M , the set of all finite paths of length i is denoted by Paths_M^i . Let $\text{Paths}_M^i(s)$ denote the set of paths p in Paths_M^i such that $\text{first}(p) = s$. Similarly, if $p \in \text{Paths}_M^i$ and $i \leq j$, then let $\text{Paths}_M^j(p)$ denote the set of paths p' in Paths_M^j such that there exists $p'' \in \text{Paths}_M^{j-i}$ with $p' = p \cdot p''$. We denote the set of all finite paths in M by Paths_M and the set of finite paths of length at most H by $\text{Paths}_M^{\leq H}$.

► **Definition 2.** The *total reward of a finite path* $p = s_0 a_0 \dots s_n$ in M is defined as

$$\text{Reward}_M(p) = \sum_{t=0}^{n-1} R(s_t, a_t) + R_T(s_n).$$

A (probabilistic) *strategy* is a function $\sigma : \text{Paths}_M \rightarrow \mathcal{D}(A)$ that maps a path p to a probability distribution in $\mathcal{D}(A)$. A strategy σ is *deterministic* if the support of the probability distributions $\sigma(p)$ has size 1, it is *memoryless* if $\sigma(p)$ depends only on $\text{last}(p)$, i.e. if σ satisfies that for all $p, p' \in \text{Paths}_M$, $\text{last}(p) = \text{last}(p') \Rightarrow \sigma(p) = \sigma(p')$. For a probabilistic strategy σ and $i \in \mathbb{N}$, let $\text{Paths}_M^i(\sigma)$ denote the paths $p = s_0 a_0 \dots s_i$ in Paths_M^i such that for all $t \in [0, i - 1]$, $a_t \in \text{Supp}(\sigma(p|_t))$. For a finite path p of length $i \in \mathbb{N}$ and some $j \geq i$, let $\text{Paths}_M^j(p, \sigma)$ denote $\text{Paths}_M^j(\sigma) \cap \text{Paths}_M^j(p)$.

For a strategy σ and a path $p \in \text{Paths}_M^i(\sigma)$, let the probability of $p = s_0 a_0 \dots s_i$ in M according to σ be defined as $\mathbb{P}_{M, \sigma}^i(p) = \prod_{t=0}^{i-1} \sigma(p|_t)(a_t) P(s_t, a_t)(s_{t+1})$. The mapping $\mathbb{P}_{M, \sigma}^i$ defines a probability distribution over $\text{Paths}_M^i(\sigma)$.

¹ We assume for convenience that every action in A can be taken from every state. One may need to limit this choice to a subset of *legal* actions that depends on the current state. This concept can be encoded in our formalism by adding a sink state reached with probability 1 when taking an illegal action.

► **Definition 3.** *The expected average reward of a probabilistic strategy σ in an MDP M , starting from state s , is defined as*

$$\text{Val}_M(s, \sigma) = \liminf_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} [\text{Reward}_M(p)],$$

where p is a random variable over $\text{Paths}_M^n(\sigma)$ following the distribution $\mathbb{P}_{M, \sigma}^n$.

► **Definition 4.** *The optimal expected average reward starting from a state s in an MDP M is defined over all strategies σ in M as $\text{Val}_M(s) = \sup_{\sigma} \text{Val}_M(s, \sigma)$.*

One can restrict the supremum to deterministic memoryless strategies [23, Proposition 6.2.1]. A strategy σ is called ϵ -optimal for the expected average reward if $\text{Val}_M(s, \sigma) \geq \text{Val}_M(s) - \epsilon$ for all s .

► **Definition 5.** *The expected total reward of a probabilistic strategy σ in an MDP M , starting from state s and for a finite horizon i , is defined as $\text{Val}_M^i(s, \sigma) = \mathbb{E} [\text{Reward}_M(p)]$, where p is a random variable over $\text{Paths}_M^i(\sigma)$ following the distribution $\mathbb{P}_{M, \sigma}^i$.*

► **Definition 6.** *The optimal expected total reward starting from a state s in an MDP M , with horizon $i \in \mathbb{N}$, is defined over all strategies σ in M as $\text{Val}_M^i(s) = \sup_{\sigma} \text{Val}_M^i(s, \sigma)$.*

One can restrict the supremum to deterministic strategies [23, Theorem 4.4.1.b].

Let $\sigma_{M, s}^{i, *}$ denote a deterministic strategy that maximises $\text{Val}_M^i(s, \sigma)$, and refer to it as an optimal strategy for the expected total reward of horizon i at state s . For $i \in \mathbb{N}$, let σ_M^i refer to a deterministic memoryless strategy that maps every state s in M to the first action of a corresponding optimal strategy for the expected total reward of horizon i , so that $\sigma_M^i(s) = \sigma_{M, s}^{i, *}(s)$. As there may exist several deterministic strategies σ that maximise $\text{Val}_M^i(s, \sigma)$, we denote by $\text{opt}_M^i(s)$ the set of actions a such that there exists an optimal strategy $\sigma_{M, s}^{i, *}$ that selects a from s . A strategy σ_M^i can be obtained by the value iteration algorithm:

► **Proposition 7** (Value iteration [23, Section 5.4]). *For a state s in MDP M , for all $i \in \mathbb{N}$,*

- $\text{Val}_M^{i+1}(s) = \max_{a \in A} [R(s, a) + \sum_{s'} P(s, a)(s') \text{Val}_M^i(s')]$
- $\text{opt}_M^{i+1}(s) = \arg \max_{a \in A} [R(s, a) + \sum_{s'} P(s, a)(s') \text{Val}_M^i(s')]$

Moreover, for a large class of MDPs and a large enough n , the strategy σ_M^n is ϵ -optimal for the expected average reward:

► **Proposition 8** ([23, Theorem 9.4.5]). *For a strongly aperiodic² Markov decision process M , it holds that $\text{Val}_M(s) = \lim_{n \rightarrow \infty} [\text{Val}_M^{n+1}(s) - \text{Val}_M^n(s)]$. Moreover, for any $\epsilon > 0$ there exists $N \in \mathbb{N}$ such that for all $n \geq N$, $\text{Val}_M(s, \sigma_M^n) \geq \text{Val}_M(s) - \epsilon$ for all s .*

A simple transformation can be used to make an MDP strongly aperiodic without changing the optimal expected average reward and the associated optimal strategies. Therefore, one can use an algorithm computing the strategy σ_M^H in order to optimise for the expected average reward, and obtain theoretical guarantees for a horizon H big enough. This is known as the *receding horizon* approach.

Finally, we will use the notation $T(M, s_0, H)$ to refer to an MDP obtained as a tree-shaped *unfolding* of M from state s_0 and for a depth of H . In particular, the states of $T(M, s_0, H)$ correspond to paths in $\text{Paths}_M^{\leq H}(s_0)$. Then, it holds that:

► **Lemma 9.** *$\text{Val}_M^H(s_0)$ is equal to $\text{Val}_{T(M, s_0, H)}^H(s_0)$, and $\text{opt}_M^H(s_0)$ is equal to $\text{opt}_{T(M, s_0, H)}^H(s_0)$.*

The aperiodicity and unfolding transformations are detailed in Appendix A.

² A Markov decision process is strongly aperiodic if $P(s, a)(s) > 0$ for all $s \in S$ and $a \in A$.

2.2 Bandit problems and UCB

In this section, we present bandit problems, whose study forms the basis of a theoretical analysis of Monte Carlo tree search algorithms.

Let A denote a finite set of actions. For each $a \in A$, let $(x_{a,t})_{t \geq 1}$ be a sequence of random payoffs associated to a . They correspond to successive plays of action a , and for every action a and every $t \geq 1$, let $x_{a,t}$ be drawn with respect to a probability distribution $\mathcal{D}_{a,t}$ over $[0, 1]$. We denote by $X_{a,t}$ the random variable associated to this drawing. In a *fixed distributions* setting (the classical bandit problem), every action is associated to a fixed probability distribution \mathcal{D}_a , so that $\mathcal{D}_{a,t} = \mathcal{D}_a$ for all $t \geq 1$.

The *bandit problem* consists of a succession of steps where the player selects an action and observes the associated payoff, while trying to maximise the cumulative gains. For example, selecting action a , then b and then a again would yield the respective payoffs $x_{a,1}$, $x_{b,1}$ and $x_{a,2}$ for the first three steps, drawn from their respective distributions. Let the regret R_n denote the difference, after n steps, between the optimal expected payoff $\max_{a \in A} \mathbb{E}[\sum_{t=1}^n X_{a,t}]$ and the expected payoff associated to our action selection. The goal is to minimise the long-term regret when the number of steps n increases.

The algorithm UCB1 of [3] offers a practical solution to this problem, and offers theoretical guarantees. For an action a and $n \geq 1$, let $\bar{x}_{a,n} = \frac{1}{n} \sum_{t=1}^n x_{a,t}$ denote the average payoff obtained from the first n plays of a . Moreover, for a given step number t let t_a denote how many times action a was selected in the first t steps. The algorithm UCB1 chooses, at step $t+1$, the action a that maximises $\bar{x}_{a,t_a} + c_{t,t_a}$, where c_{t,t_a} is defined as $\sqrt{\frac{2 \ln t}{t_a}}$. This procedure enjoys optimality guarantees detailed in [3], as it keeps the regret R_n below $O(\log n)$.

We will make use of an extension of these results to the general setting of *non-stationary* bandit problems, where the distributions $\mathcal{D}_{a,t}$ are no longer fixed with respect to t . This problem has been studied in [18], and results were obtained for a class of distributions $\mathcal{D}_{a,t}$ that respect assumptions referred to as *drift conditions*.

For a fixed $n \geq 1$, let $\bar{X}_{a,n}$ denote the random variable obtained as the average of the random variables associated with the first n plays of a . Let $\mu_{a,n} = \mathbb{E}[\bar{X}_{a,n}]$. We assume that these expected means eventually converge, and let $\mu_a = \lim_{n \rightarrow \infty} \mu_{a,n}$.

► **Definition 10** (Drift conditions). *For all $a \in A$, the sequence $(\mu_{a,n})_{n \geq 1}$ converges to some value μ_a . Moreover, there exists a constant $C_p > 0$ and an integer N_p such that for $n \geq N_p$ and any $\delta > 0$, if $\Delta_n(\delta) = C_p \sqrt{n \ln(1/\delta)}$ then the tail inequalities $\mathbb{P}[n\bar{X}_{a,n} \geq n\mu_{a,n} + \Delta_n(\delta)] \leq \delta$ and $\mathbb{P}[n\bar{X}_{a,n} \leq n\mu_{a,n} - \Delta_n(\delta)] \leq \delta$ hold.*

We recall in Appendix B the results of [18], and provide an informal description of those results here. Consider using the algorithm UCB1 on a non-stationary bandit problem satisfying the drift conditions, with $c_{t,t_a} = 2C_p \sqrt{\frac{\ln t}{t_a}}$. First, one can bound logarithmically the number of times a suboptimal action is played. This is used to bound the difference between μ_a and $\mathbb{E}[\bar{X}_n]$ by $O(\ln n/n)$, where a is an optimal action and where \bar{X}_n denotes the global average of payoffs received over the first n steps. This is the main theoretical guarantee obtained for the optimality of UCB1. Also for any action a , the authors state a lower bound for the number of times the action is played. The authors also prove a tail inequality similar to the one described in the drift conditions, but on the random variable \bar{X}_n instead of $\bar{X}_{a,n}$. This will be useful for inductive proofs later on, when the usage of UCB1 is nested so that the global sequence \bar{X}_n corresponds to a sequence $\bar{X}_{b,n}$ of an action b played from the next state of the MDP. Finally, it is shown that the probability of making the wrong decision (choosing a suboptimal action) converges to 0 as the number of plays n grows large enough.

3 Monte Carlo tree search with simulation

In a receding horizon approach, the objective is to compute $\text{Val}_M^H(s_0)$ and σ_M^H for some state s_0 and some horizon H . Exact procedures such as the recursive computation of Proposition 7 can not be used on large MDPs, resulting in heuristic approaches. We focus on the Monte Carlo Tree Search (MCTS) algorithm [7], that can be seen as computing approximations of Val_M^H and $\sigma_M^H(s_0)$ on the unfolding $T(M, s_0, H)$. Note that rewards in the MDP M are bounded.³ For the sake of simplicity we assume without loss of generality that for all paths p of length at most H the total reward $\text{Reward}_M(p)$ belongs to $[0, 1]$.

Given an initial state s_0 , MCTS is an iterative process that incrementally constructs a search tree rooted at s_0 describing paths of M and their associated values. This process goes on until a specified budget (of number of iterations or time) is exhausted. An iteration constructs a path in M by following a decision strategy to *select* a sequence of nodes in the search tree. When a node that is not part of the current search tree is reached, the tree is expanded with this new node, whose expected reward is approximated by *simulation*. This value is then used to update the knowledge of all selected nodes in *backpropagation*.

In the search tree, each node represents a path. For a node p and an action $a \in A$, let $\text{children}(p, a)$ be a list of nodes representing paths of the form $p \cdot as'$ where $s' \in \text{Supp}(P(\text{last}(p), a))$. For each node (resp. node-action pair) we store a value $\text{value}(p)$ (resp. $\text{value}(p, a)$) computed for node p (resp. for playing a from node p), meant to approximate $\text{Val}_M^{H-|p|}(\text{last}(p))$ (resp. $R(\text{last}(p), a) + \sum_{s'} P(\text{last}(p), a)(s') \text{Val}_M^{H-|p|-1}(s')$), and a counter $\text{count}(p)$ (resp. $\text{count}(p, a)$), that keeps track of the number of iterations that selected node p (resp. that selected the action a from p). We add subscripts $i \geq 1$ to these notations to denote the number of previous iterations, so that $\text{value}_i(p)$ is the value of p obtained after i iterations of MCTS, among which p was selected $\text{count}_i(p)$ times. We also define $\text{total}_i(p)$ and $\text{total}_i(p, a)$ as shorthand for respectively $\text{value}_i(p) \times \text{count}_i(p)$ and $\text{value}_i(p, a) \times \text{count}_i(p, a)$. Each iteration consists of three phases. Let us describe these phases at iteration number i .

Selection phase. Starting from the root node, MCTS descends through the existing search tree by choosing actions based on the current values and counters and by selecting next states stochastically according to the MDP. This continues until reaching a node q , either outside of the search tree or at depth H . In the former case, the simulation phase is called to obtain a value $\text{value}_i(q)$ that will be backpropagated along the path q . In the latter case, we use the exact value $\text{value}_i(q) = R_T(\text{last}(q))$ instead.

The action selection process needs to balance between the exploration of new paths and the exploitation of known, promising paths. A popular way to balance both is the *upper confidence bound for trees* (UCT) algorithm [18], that interprets the action selection problem of each node of the MCTS tree as a bandit problem, and selects an action a^* in the set $\arg \max_{a \in A} \left[\text{value}_{i-1}(p, a) + C \sqrt{\frac{\ln(\text{count}_{i-1}(p))}{\text{count}_{i-1}(p, a)}} \right]$, for some constant C .

Simulation phase. In the simulation phase, the goal is to get an *initial approximation* for the value of a node p , that will be refined in future iterations of MCTS. Classically, a sampling-based approach can be used, where one computes a fixed number $c \in \mathbb{N}$ of paths $p \cdot p'$ in $\text{Paths}_M^H(p)$. Then, one can compute $\text{value}_i(p) = \frac{1}{c} \sum_{p'} \text{Reward}_M(p')$, and fix $\text{count}_i(p)$ to 1. Usually, the samples are derived by selecting actions uniformly at random in the MDP.

³ There are finitely many paths of length at most H , with rewards in \mathbb{R} .

In our theoretical analysis of MCTS, we take a *more general approach* to the simulation phase, defined by a finite domain $I \subseteq [0, 1]$ and a function $f : \text{Paths}_M^{\leq H} \rightarrow \mathcal{D}(I)$ that maps every path p to a probability distribution on I . In this approach, the simulation phase simply draws a value $\text{value}_i(p)$ at random according to the distribution $f(p)$, and sets $\text{count}_i(p) = 1$.

Backpropagation phase. From the value $\text{value}_i(p)$ obtained at a leaf node $p = s_0 a_0 s_1 \dots s_h$ at depth h in the search tree, let $\text{reward}_i(p_{|k}) = \sum_{l=k}^{h-1} R(s_l, a_l) + \text{value}_i(p)$ denote the reward associated with the path from node $p_{|k}$ to p in the search tree. For k from 0 to $h - 1$ we update the values according to $\text{value}_i(p_{|k}) = \frac{\text{total}_{i-1}(p_{|k}) + \text{reward}_i(p_{|k})}{\text{count}_i(p_{|k})}$. The value $\text{value}_i(p_{|k}, a_k)$ is updated based on $\text{total}_{i-1}(p_{|k}, a_k)$, $\text{reward}_i(p_{|k})$ and $\text{count}_i(p_{|k}, a_k)$ with the same formula.

Theoretical analysis. In the remainder of this section, we prove Theorem 11, that provides theoretical properties of the MCTS algorithm with a general simulation phase (defined by some fixed I and f). This theorem was proven in [18, Theorem 6] for a version of the algorithm that called MCTS recursively until leaves were reached, as opposed to the sampling-based approach that has become standard in practice. Note that sampling-based approaches are captured by our general description of the simulation phase. Indeed, if the number of samples c is set to 1, let I be the set of rewards associated with paths of $\text{Paths}_M^{\leq H}$, and let $f(p)$ be a probability distribution over I , such that for every reward $\text{Reward}_M(p') \in I$, $f(p)(\text{Reward}_M(p'))$ is the probability of path p' being selected with a uniform action selections in $T(M, s_0, H)$, starting from the node p . Then, the value $\text{value}_i(p)$ drawn at random according to the distribution $f(p)$ corresponds to the reward of a random sample $p \cdot p'$ drawn in Paths_M^H . If the number of samples c is greater than 1, one simply needs to extend I to be the set of average rewards over c paths, while $f(p)$ becomes a distribution over average rewards.

► **Theorem 11.** *Consider an MDP M , a horizon H and a state s_0 . Let $V_n(s_0)$ be a random variable that represents the value $\text{value}_n(s_0)$ at the root of the search tree after n iterations of the MCTS algorithm on M . Then, $|\mathbb{E}[V_n(s_0)] - \text{Val}_M^H(s_0)|$ is bounded by $\mathcal{O}(\ln n/n)$. Moreover, the failure probability $\mathbb{P}[\arg \max_a \text{value}_n(s_0, a) \not\subseteq \text{opt}_M^H(s_0)]$ converges to zero.*

Following the proof scheme of [18, Theorem 6], this theorem is obtained from the results mentioned in Section 2.2. To this end, every node p of the search tree is considered to be an instance of a bandit problem with non-stationary distributions. Every time a node is selected, a step is processed in the corresponding bandit problem.

Let $(\mathcal{I}_i(p))_{i \geq 1}$ be a sequence of iteration numbers for the MCTS algorithm that describes when the node p is selected, so that the simulation phase was used on p at iteration number $\mathcal{I}_1(p)$, and so that the i -th selection of node p happened on the iteration number $\mathcal{I}_i(p)$. We define sequences $(\mathcal{I}_i(p, a))_{i \in \mathbb{N}}$ similarly for node-action pairs.

For all paths p and actions a , a payoff sequence $(x_{a,t})_{t \geq 1}$ of associated random variables $(X_{a,t})_{t \geq 1}$ is defined by $x_{a,t} = \text{reward}_{\mathcal{I}_t(p,a)}(p)$. Note that in the selection phase at iteration number $\mathcal{I}_t(p, a)$, p must have been selected and must be a prefix of length k of the leaf node p' reached in this iteration, so that $\text{reward}_{\mathcal{I}_t(p,a)}(p)$ is computed as $\text{reward}_{\mathcal{I}_t(p,a)}(p'_{|k})$ in the backpropagation phase. According to the notations of Section 2.2, for all $t \geq 1$ we have $\text{count}_{\mathcal{I}_t(p)}(p) = t$, $\text{count}_{\mathcal{I}_t(p)}(p, a) = t_a$ and $\text{value}_{\mathcal{I}_t(p)}(p, a) = \bar{x}_{a,t_a}$.

Then, one can obtain Theorem 11 by applying inductively the UCB1 results recalled in Appendix B on the search tree in a bottom-up fashion. Indeed, as the root s_0 is selected at every iteration, $\mathcal{I}_n(s_0) = n$ and $\text{value}_n(s_0) = \bar{x}_n$, while $\text{Val}_M^H(s_0)$ corresponds to recursively selecting optimal actions by Proposition 7.

The main difficulty, and the difference our simulation phase brings compared with the proof of [18, Theorem 6], lies in showing that our payoff sequences $(x_{a,t})_{t \geq 1}$, defined with an initial simulation step, still satisfy the drift conditions of Definition 10. We argue that this is true for all simulation phases defined by any I and f :

► **Lemma 12.** *For any MDP M , horizon H and state s_0 , the sequences $(X_{a,t})_{t \geq 1}$ satisfy the drift conditions.*

Although the long-term guarantees of Theorem 11 hold for any simulation phase independently of the MDP, in practice one would expect better results from a good simulation, that gives a value close to the real value of the current node. Domain-specific knowledge can be used to obtain such simulations, and also to guide the selection phase based on heuristics. Our goal will be to preserve the theoretical guarantees of MCTS in the process.

4 Symbolic advice for MCTS

In this section, we introduce a notion of advice meant to guide the construction of the Monte Carlo search tree. We argue that a symbolic approach is needed in order to handle large MDPs in practice. Let a *symbolic advice* \mathcal{A} be a logical formula over finite paths whose truth value can be tested with an operator \models .

► **Example 13.** A number of standard notions can fit this framework. For example, reachability and safety properties, LTL formulæ over finite traces or regular expressions could be used. We will use a safety property for PAC-MAN as a example (see Figure 1), by assuming that the losing states of the MDP should be avoided. This advice is thus satisfied by every path such that Pac-Man does not make contact with a ghost.

We denote by $\text{Paths}_M^H(\mathcal{A})$ the set of paths $p \in \text{Paths}_M^H$ such that $p \models \mathcal{A}$. For a path $p \in \text{Paths}_M^{\leq H}$, we denote by $\text{Paths}_M^H(p, \mathcal{A})$ the set of paths $p' \in \text{Paths}_M^H(p)$ such that $p' \models \mathcal{A}$.⁴

A *nondeterministic strategy* is a function $\sigma : \text{Paths}_M \rightarrow 2^A$ that maps a finite path p to a subset of A . For a strategy σ' and a nondeterministic strategy σ , $\sigma' \subseteq \sigma$ if for all p , $\text{Supp}(\sigma'(p)) \subseteq \sigma(p)$. Similarly, a nondeterministic strategy for the environment is a function $\tau : \text{Paths}_M \times A \rightarrow 2^S$ that maps a finite path p and an action a to a subset of $\text{Supp}(P(\text{last}(p), a))$. We extend the notations used for probabilistic strategies to nondeterministic strategies in a natural way, so that $\text{Paths}_M^H(\sigma)$ and $\text{Paths}_M^H(\tau)$ denote the paths of length H compatible with the strategy σ or τ , respectively.

For a symbolic advice \mathcal{A} and a horizon H , we define a nondeterministic strategy $\sigma_{\mathcal{A}}^H$ and a nondeterministic strategy $\tau_{\mathcal{A}}^H$ for the environment such that for all paths p with $|p| < H$,

$$\sigma_{\mathcal{A}}^H(p) = \{a \in A \mid \exists s \in S, \exists p' \in \text{Paths}_M^{H-|p|-1}(s), p \cdot as \cdot p' \models \mathcal{A}\},$$

$$\tau_{\mathcal{A}}^H(p, a) = \{s \in S \mid \exists p' \in \text{Paths}_M^{H-|p|-1}(s), p \cdot as \cdot p' \models \mathcal{A}\}.$$

The strategies $\sigma_{\mathcal{A}}^H$ and $\tau_{\mathcal{A}}^H$ can be defined arbitrarily on paths p of length at least H , for example with $\sigma_{\mathcal{A}}^H(p) = A$ and $\tau_{\mathcal{A}}^H(p, a) = \text{Supp}(P(\text{last}(p), a))$ for all actions a . Note that by definition, $\text{Paths}_M^H(s, \mathcal{A}) = \text{Paths}_M^H(s, \sigma_{\mathcal{A}}^H) \cap \text{Paths}_M^H(s, \tau_{\mathcal{A}}^H)$ for all states s .

⁴ In particular, for all $s \in S$, $\text{Paths}_M^H(s, \mathcal{A})$ refers to the paths of length H that start from s and that satisfy \mathcal{A} .

Let \top (resp. \perp) denote the universal advice (resp. the empty advice) satisfied by every finite path (resp. never satisfied), and let σ_\top and τ_\top (resp. σ_\perp and τ_\perp) be the associated nondeterministic strategies. We define a class of advice that can be enforced against an adversarial environment by following a nondeterministic strategy, and that are minimal in the sense that paths that are not compatible with this strategy are not allowed.

► **Definition 14** (Strongly enforceable advice). *A symbolic advice \mathcal{A} is called a strongly enforceable advice from a state s_0 and for a horizon H if there exists a nondeterministic strategy σ such that $\text{Paths}_M^H(s_0, \sigma) = \text{Paths}_M^H(s_0, \mathcal{A})$, and such that $\sigma(p) \neq \emptyset$ for all paths $p \in \text{Paths}_M^{\leq H-1}(s_0, \sigma)$.*

Note that Definition 14 ensures that paths that follow σ can always be extended into longer paths that follow σ . This is a reasonable assumption to make for a nondeterministic strategy meant to enforce a property. In particular, s_0 is a path of length 0 in $\text{Paths}_M^0(s_0, \sigma)$, so that $\sigma(s_0) \neq \emptyset$ and so that by induction $\text{Paths}_M^i(s_0, \sigma) \neq \emptyset$ for all $i \in [0, H]$.

► **Lemma 15.** *Let \mathcal{A} be a strongly enforceable advice from s_0 with horizon H . It holds that $\text{Paths}_M^H(s_0, \sigma_{\mathcal{A}}^H) = \text{Paths}_M^H(s_0, \mathcal{A})$. Moreover, for all paths $p \in \text{Paths}_M^{\leq H-1}(s_0)$ and all actions a , either $\tau_{\mathcal{A}}^H(p, a) = \tau_\top(p, a)$ or $\tau_{\mathcal{A}}^H(p, a) = \tau_\perp(p, a)$. Finally, for all paths p in $\text{Paths}_M^{\leq H-1}(s_0, \sigma_{\mathcal{A}}^H)$, $\sigma_{\mathcal{A}}^H(p) \neq \emptyset$ and $a \in \sigma_{\mathcal{A}}^H(p)$ if and only if $\tau_{\mathcal{A}}^H(p, a) = \tau_\top(p, a)$.*

Proof. We have $\text{Paths}_M^H(s_0, \mathcal{A}) = \text{Paths}_M^H(s_0, \sigma_{\mathcal{A}}^H) \cap \text{Paths}_M^H(s_0, \tau_{\mathcal{A}}^H)$ for any advice \mathcal{A} . Let us prove that $\text{Paths}_M^H(s_0, \sigma_{\mathcal{A}}^H) \subseteq \text{Paths}_M^H(s_0, \mathcal{A})$ for a strongly enforceable advice \mathcal{A} of associated strategy σ . Let $p = p' \cdot as$ be a path in $\text{Paths}_M^H(s_0, \sigma_{\mathcal{A}}^H)$. By definition of $\sigma_{\mathcal{A}}^H$, there exists $s' \in S$ such that $p' \cdot as' \models \mathcal{A}$, so that $p' \cdot as' \in \text{Paths}_M^H(s_0, \mathcal{A}) = \text{Paths}_M^H(s_0, \sigma)$. Since $s \in \text{Supp}(P(\text{last}(p'), a))$, $p = p' \cdot as$ must also belong to $\text{Paths}_M^H(s_0, \sigma) = \text{Paths}_M^H(s_0, \mathcal{A})$.

Consider a path p and an action a such that $|p| < H$. We want to prove that either all stochastic transitions starting from (p, a) are allowed by \mathcal{A} , or none of them are. By contradiction, let us assume that there exists s_1 and s_2 in $\text{Supp}(P(\text{last}(p), a))$ such that for all $p'_1 \in \text{Paths}_M^{H-|p|-1}(s_1)$, $p \cdot as_1 \cdot p'_1 \not\models \mathcal{A}$, and such that there exists $p'_2 \in \text{Paths}_M^{H-|p|-1}(s_2)$ with $p \cdot as_2 \cdot p'_2 \models \mathcal{A}$. From $p \cdot as_2 \cdot p'_2 \models \mathcal{A}$, we obtain $p \cdot as_2 \cdot p'_2 \in \text{Paths}_M^H(\sigma)$, so that $p \cdot as_2$ is a path that follows σ . Then, $p \cdot as_1$ is a path that follows σ as well. It follows that $\sigma(p \cdot as_1) \neq \emptyset$, and $p \cdot as_1$ can be extended in to a path $p \cdot as_1 p'_3 \in \text{Paths}_M^H(\sigma)$. This implies the contradiction $p \cdot as_1 p'_3 \models \mathcal{A}$.

Finally, consider a path p in $\text{Paths}_M^{\leq H-1}(s_0, \sigma_{\mathcal{A}}^H)$. By the definitions of $\sigma_{\mathcal{A}}^H$ and $\tau_{\mathcal{A}}^H$, $a \in \sigma_{\mathcal{A}}^H(p)$ if and only if $\tau_{\mathcal{A}}^H(p, a) \neq \emptyset$, so that $\tau_{\mathcal{A}}^H(p, a) = \tau_\top(p, a)$. Then, let us write $p = p' \cdot as$. From $p \in \text{Paths}_M^{\leq H-1}(s_0, \sigma_{\mathcal{A}}^H)$ we get $a \in \sigma_{\mathcal{A}}^H(p')$, so that $s \in \tau_{\mathcal{A}}^H(p', a)$, and therefore $\sigma_{\mathcal{A}}^H(p) \neq \emptyset$. ◀

A strongly enforceable advice is encoding a notion of guarantee, as $\sigma_{\mathcal{A}}^H$ is a winning strategy for the reachability objective on $T(M, s_0, H)$ defined by the set $\text{Paths}_M^H(\mathcal{A})$.

We say that the strongly enforceable advice \mathcal{A}' is *extracted* from a symbolic advice \mathcal{A} for a horizon H and a state s_0 if \mathcal{A}' is the greatest part of \mathcal{A} that can be guaranteed for the horizon H starting from s_0 , *i.e.* if $\text{Paths}_M^H(s_0, \mathcal{A}')$ is the greatest subset of $\text{Paths}_M^H(s_0, \mathcal{A})$ such that $\sigma_{\mathcal{A}'}^H$ is a winning strategy for the reachability objective $\text{Paths}_M^H(s_0, \mathcal{A})$ on $T(M, s_0, H)$. This greatest subset always exists because if \mathcal{A}'_1 and \mathcal{A}'_2 are strongly enforceable advice in \mathcal{A} , then $\mathcal{A}'_1 \cup \mathcal{A}'_2$ is strongly enforceable by union of the nondeterministic strategies associated with \mathcal{A}'_1 and \mathcal{A}'_2 . However, this greatest subset may be empty, and as \perp is not a strongly enforceable advice we say that in this case \mathcal{A} cannot be enforced from s_0 with horizon H .

► **Example 16.** Consider a symbolic advice \mathcal{A} described by the safety property for PAC-MAN of Example 13. For a fixed horizon H , the associated nondeterministic strategies $\sigma_{\mathcal{A}}^H$ and $\tau_{\mathcal{A}}^H$ describe action choices and stochastic transitions compatible with this property. Notably, \mathcal{A} may not be a strongly enforceable advice, as there may be situations (p, a) where some stochastic transitions lead to bad states and some do not. In the small grid of Figure 1, the path of length 1 that corresponds to Pac-Man going left and the red ghost going up is allowed by the advice \mathcal{A} , but not by any safe strategy for Pac-Man as there is a possibility of losing by playing left. If a strongly enforceable advice \mathcal{A}' can be extracted from \mathcal{A} , it is a more restrictive safety property, where the set of bad states is obtained as the attractor [20, Section 2.3] for the environment towards the bad states defined in \mathcal{A} . In this setting, \mathcal{A}' corresponds to playing according to a strategy for Pac-Man that ensures not being eaten by adversarial ghosts for the next H steps.

► **Definition 17 (Pruned MDP).** For an MDP $M = (S, A, P, R, R_T)$ a horizon $H \in \mathbb{N}$, a state s_0 and an advice \mathcal{A} , let the pruned unfolding $T(M, s_0, H, \mathcal{A})$ be defined as a sub-MDP of $T(M, s_0, H)$ that contains exactly all paths in $\text{Paths}_M^H(s_0)$ satisfying \mathcal{A} . It can be obtained by removing all action transitions that are not compatible with $\sigma_{\mathcal{A}}^H$, and all stochastic transitions that are not compatible with $\tau_{\mathcal{A}}^H$. The distributions $P(p, a)$ are then normalised over the stochastic transitions that are left.

Note that by Lemma 15, if \mathcal{A} is a strongly enforceable advice then $\tau_{\mathcal{A}}^H(p, a) = \tau_{\top}(p, a)$ for all paths p in $\text{Paths}_M^{\leq H-1}(s_0, \sigma_{\mathcal{A}}^H)$, so that the normalisation step for the distributions $P(p, a)$ is not needed. It follows that for all nodes p in $T(M, s_0, H, \mathcal{A})$ and all actions a , the distributions $P(p, a)$ in $T(M, s_0, H, \mathcal{A})$ are the same as in $T(M, s_0, H)$. Thus, for all strategies σ in $T(M, s_0, H, \mathcal{A})$, $\text{Val}_{T(M, s_0, H, \mathcal{A})}^H(s_0, \sigma) = \text{Val}_{T(M, s_0, H)}^H(s_0, \sigma)$, so that $\text{Val}_{T(M, s_0, H, \mathcal{A})}^H(s_0) \leq \text{Val}_{T(M, s_0, H)}^H(s_0) = \text{Val}_M^H(s_0)$ by Lemma 9.

► **Definition 18 (Optimality assumption).** An advice \mathcal{A} satisfies the optimality assumption for horizon H if $\sigma_{M, s}^{H, *} \subseteq \sigma_{\mathcal{A}}^H$ for all $s \in S$, where $\sigma_{M, s}^{H, *}$ is an optimal strategy for the expected total reward of horizon H at state s .

► **Lemma 19.** Let \mathcal{A} be a strongly enforceable advice that satisfies the optimality assumption. Then, $\text{Val}_M^H(s_0)$ equals $\text{Val}_{T(M, s_0, H, \mathcal{A})}^H(s_0)$. Moreover, $\text{opt}_{T(M, s_0, H, \mathcal{A})}^H(s_0) \subseteq \text{opt}_M^H(s_0)$.

Proof. By the optimality assumption $\sigma_{M, s}^{H, *}$ is a strategy that can be followed in $T(M, s_0, H, \mathcal{A})$. Indeed, from a path p in $\text{Paths}_M^{\leq H-1}(s_0, \sigma_{\mathcal{A}}^H)$ any action a in the support of $\sigma_{M, s}^{H, *}(\text{last}(p))$ satisfies $a \in \sigma_{\mathcal{A}}^H(p)$. Thus, by Lemma 9 $\text{Val}_M^H(s_0) = \text{Val}_{T(M, s_0, H)}^H(s_0, \sigma_{M, s}^{H, *}) = \text{Val}_{T(M, s_0, H, \mathcal{A})}^H(s_0, \sigma_{M, s}^{H, *})$. By definition of the optimal expected total reward, $\text{Val}_{T(M, s_0, H, \mathcal{A})}^H(s_0, \sigma_{M, s}^{H, *}) \leq \text{Val}_{T(M, s_0, H, \mathcal{A})}^H(s_0)$, so that $\text{Val}_M^H(s_0) = \text{Val}_{T(M, s_0, H, \mathcal{A})}^H(s_0)$. Let a be an action in $\text{opt}_{T(M, s_0, H, \mathcal{A})}^H(s_0)$. There exists an optimal strategy σ that maximises $\text{Val}_{T(M, s_0, H, \mathcal{A})}^H(s_0, \sigma)$ so that $\sigma(s_0) = a$. It follows from $\text{Val}_{T(M, s_0, H)}^H(s_0) = \text{Val}_{T(M, s_0, H, \mathcal{A})}^H(s_0)$ that σ is also an optimal strategy in $T(M, s_0, H)$, so that $a \in \text{opt}_{T(M, s_0, H)}^H(s_0) = \text{opt}_M^H(s_0)$ by Lemma 9. ◀

► **Example 20.** Let \mathcal{A}' be a strongly enforceable safety advice for PAC-MAN as described in Example 16. Assume that visiting a bad state leads to an irrecoverably bad reward, so that taking an unsafe action (*i.e.* an action such that there is a non-zero probability of losing associated with all Pac-Man strategies) is always worse (on expectation) than taking a safe action. Then, the optimality assumption holds for the advice \mathcal{A}' . This can be achieved by giving a penalty score for losing that is low enough.

4.1 MCTS under symbolic advice

We will augment the MCTS algorithm using two advice: a selection advice φ to guide the MCTS tree construction, and a simulation advice ψ to prune the sampling domain. We assume that the selection advice is a strongly enforceable advice that satisfies the optimality assumption. Notably, we make no such assumption for the simulation advice, so that any symbolic advice can be used.

Selection phase under advice. We use the advice φ to prune the tree according to σ_φ^H . Therefore, from any node p our version of UCT selects an action a^* in the set

$$\arg \max_{a \in \sigma_\varphi^H(p)} \left[\text{value}(p, a) + C \sqrt{\frac{\ln(\text{count}(p))}{\text{count}(p, a)}} \right].$$

Simulation phase under advice. For the simulation phase, we use a sampling-based approach biased by the simulation advice: paths are sampled by picking actions uniformly at random in the pruned MDP $T(M, s_0, H, \psi)$, with a fixed prefix p defined by the current node in the search tree. This can be interpreted as a probability distribution over $\text{Paths}_M^H(p, \psi)$. If $p \notin T(M, s_0, H, \psi)$, the simulation phase outputs a value of 0 as it is not possible to satisfy ψ from p . Another approach that does not require computing the pruned MDP repeats the following steps for a bounded number of time before returning 0 if no valid sample is found:

1. Pick a path $p \cdot p' \in \text{Paths}_M^H(p)$ using a uniform sampling method;
2. If $p \cdot p' \not\models \psi$, reject and try again, otherwise output p' as a sample.

We compute $\text{value}_i(p)$ by averaging the rewards of these samples.

Theoretical analysis. We show that the theoretical guarantees of the MCTS algorithm developed in Section 3 are maintained by the MCTS algorithm under symbolic advice.

► **Theorem 21.** *Consider an MDP M , a horizon H and a state s_0 . Let $V_n(s_0)$ be a random variable that represents the value $\text{value}_n(s_0)$ at the root of the search tree after n iterations of the MCTS algorithm under a strongly enforceable advice φ satisfying the optimality assumption and a simulation advice ψ . Then, $|\mathbb{E}[V_n(s_0)] - \text{Val}_M^H(s_0)| = \mathcal{O}((\ln n)/n)$. Moreover, the failure probability $\mathbb{P}[\arg \max_a \text{value}_n(s_0, a) \not\subseteq \text{opt}_M^H(s_0)]$ converges to zero.*

In order to prove Theorem 21, we argue that running MCTS under a selection advice φ and a simulation advice ψ is equivalent to running the MCTS algorithm of Section 3 on the pruned MDP $T(M, s_0, H, \varphi)$, with a simulation phase defined using the advice ψ .

The simulation phase biased by ψ can be described in the formalism of Section 3, with a domain $I = \{\frac{1}{c} \sum_{i=1}^c \text{Reward}_M(p_i) \mid p_1, \dots, p_c \in \text{Paths}_{T(M, s_0, H, \varphi)}^{\leq H}\}$, and a mapping f_ψ from paths p in $\text{Paths}_{T(M, s_0, H, \varphi)}^{\leq H}$ to a probability distribution on I describing the outcome of a sampling phase launched from the node p . Formally, the weight of $\frac{1}{c} \sum_{i=1}^c \text{Reward}_M(p_i) \in I$ in $f(p)$ is the probability of sampling the sequence of paths p_1, \dots, p_c in the simulation phase under advice launched from p . Then, from Theorem 11 we obtain convergence properties of MCTS under symbolic advice towards the value and optimal strategy in the pruned MDP, and Lemma 19 lets us conclude the proof of Theorem 21 as those values and strategies are maintained in M by the optimality assumption. In particular, the failure probability $\mathbb{P}[\arg \max_a \text{value}_n(s_0, a) \not\subseteq \text{opt}_M^H(s_0)]$ is upper bounded by $\mathbb{P}[\arg \max_a \text{value}_n(s_0, a) \not\subseteq \text{opt}_{T(M, s_0, H, \varphi)}^H(s_0)]$ since $\text{opt}_{T(M, s_0, H, \varphi)}^H \subseteq \text{opt}_M^H$.

4.2 Using satisfiability solvers

We will now discuss the use of general-purpose solvers to implement symbolic advice according to the needs of MCTS.

A symbolic advice \mathcal{A} describes a finite set of paths in Paths_M^H , and as such can be encoded as a Boolean formula over a set of variables V , such that satisfying assignments $v : V \rightarrow \{\text{true}, \text{false}\}$ are in bijection with paths in $\text{Paths}_M^H(\mathcal{A})$.

If a symbolic advice is described in Linear Temporal Logic, and a symbolic model of the MDP M is available, one can encode \mathcal{A} as a Boolean formula of size linear in the size of the LTL formula and H [5].

► **Example 22.** In practice, one can use Boolean variables to encode the positions of Pac-Man and ghosts in the next H steps of the game, then construct a CNF formula with clauses that encode the game rules and clauses that enforce the advice. The former clauses are implications such as “if Pac-Man is in position (x, y) and plays the action a , then it must be in position (x', y') at the next game step”, while the latter clauses state that the position of Pac-Man should never be equal to the position of one of the Ghosts.

On-the-fly computation of a strongly enforceable advice. A direct encoding of a strongly enforceable advice may prove impractically large. We argue for an on-the-fly computation of $\sigma_{\mathcal{A}}^H$ instead, in the particular case where the strongly enforceable advice is extracted from a symbolic advice \mathcal{A} with respect to the initial state s_0 and with horizon H .

► **Lemma 23.** *Let \mathcal{A}' be a strongly enforceable advice extracted from \mathcal{A} for horizon H . Consider a node p at depth i in $T(M, s_0, H, \mathcal{A}')$, for all $a_0 \in A$, $a_0 \in \sigma_{\mathcal{A}'}^H(p)$ if and only if*

$$\forall s_1 \exists a_1 \forall s_2 \dots \forall s_{H-i+1}, p \cdot a_0 s_1 a_1 s_2 \dots s_{H-i+1} \models \mathcal{A},$$

where actions are quantified over A and every s_k is quantified over $\text{Supp}(P(s_{k-1}, a_{k-1}))$.

Proof. The proof is a reverse induction on the depth i of p . For the initialisation step with $i = H$, let us prove that $\forall s_1, p \cdot a_0 s_1 \models \mathcal{A}$ if and only if $a_0 \in \sigma_{\mathcal{A}'}^H(p)$. On the one hand, if \mathcal{A} is guaranteed by playing a_0 from p , then a_0 must be allowed by the greatest strongly enforceable subset of \mathcal{A} . On the other hand, $a_0 \in \sigma_{\mathcal{A}'}^H(p)$ implies $\forall s_1, p \cdot a_0 s_1 \models \mathcal{A}'$ as \mathcal{A}' is strongly enforceable, and finally $\mathcal{A}' \Rightarrow \mathcal{A}$. We now assume the property holds for $1 \leq i \leq H$, and prove it for $i - 1$. If $a_0 \in \sigma_{\mathcal{A}'}^H(p)$, then for all s_1 we have $s_1 \in \tau_{\mathcal{A}'}^H(p, a_0)$, so that there exists a_1 with $a_1 \in \sigma_{\mathcal{A}'}^H(p \cdot a_0 s_1)$. As $p \cdot a_0 s_1$ is at depth i we can conclude that $\forall s_1 \exists a_1 \forall s_2 \dots \forall s_{H-i+1}, p \cdot a_0 s_1 a_1 s_2 \dots s_{H-i+1} \models \mathcal{A}$ by assumption. For the converse direction, the alternation of quantifiers states that \mathcal{A} can be guaranteed from p by some deterministic strategy that starts by playing a_0 , and therefore a_0 must be allowed by the strongly enforceable advice extracted from \mathcal{A} . ◀

Therefore, given a Boolean formula encoding \mathcal{A} , one can use a Quantified Boolean Formula (QBF) solver to compute $\sigma_{\mathcal{A}'}^H$, the strongly enforceable advice extracted from \mathcal{A} : this computation can be used whenever MCTS performs an action selection step under the advice \mathcal{A}' , as described in Section 4.1.

The performance of this approach will crucially depend on the number of alternating quantifiers, and in practice one may limit themselves to a smaller depth $h < H - i$ in this step, so that safety is only guaranteed for the next h steps.

Some properties can be inductively guaranteed, so that satisfying the QBF formula of Lemma 23 with a depth $H - i = 1$ is enough to guarantee the property globally. For example, if there always exists an action leading to states that are not bad, it is enough to check for safety locally with a depth of 1. This is the case in PAC-MAN for a deadlock-free layout when there is only one ghost.

Weighted sampling under a symbolic advice. Given a symbolic advice \mathcal{A} as a Boolean formula, and a probability distribution $w \in \mathcal{D}(\text{Paths}_M^H)$, our goal is to sample paths of M that satisfy \mathcal{A} with respect to w .⁵ Let ω denote a weight function over Boolean assignments that matches w . This reduces our problem to the weighted sampling of satisfying assignments in a Boolean formula. An exact solver for this problem may not be efficient, but one can use the techniques of [8] to perform approximate sampling in polynomial time:

► **Proposition 24** ([8]). *Given a CNF formula \mathcal{A} , a tolerance $\epsilon > 0$ and a weight function ω , we can construct a probabilistic algorithm which outputs a satisfying assignment z such that for all y that satisfies \mathcal{A} :*

$$\frac{\omega(y)}{(1 + \epsilon) \sum_{x \models \psi} \omega(x)} \leq \Pr[z = y] \leq \frac{(1 + \epsilon)\omega(y)}{\sum_{x \models \psi} \omega(x)}.$$

The above algorithm occasionally “fails” (outputs no assignment even though there are satisfying assignments) but its failure probability can be bounded by any given δ . Given an oracle for SAT, the above algorithm runs in time polynomial in $\ln(\frac{1}{\delta})$, $|\psi|$, $\frac{1}{\epsilon}$ and r where r is the ratio between highest and lowest weight according to ω .

In particular, this algorithm uses ω as a black-box, and thus does not require precomputing the probabilities of all paths satisfying \mathcal{A} . In our particular application of Proposition 24, the value r can be bounded by $\left(\frac{p_{\max}|A|}{p_{\min}}\right)^H$ where p_{\min} and p_{\max} are the smallest and greatest probabilities for stochastic transitions in M .

Note that if we wish to sample from a given node p of the search tree, we can force p as a mandatory prefix of satisfying assignments by fixing the truth value of relevant variables in the Boolean formula.

5 A Pac-Man case study

We performed our experiments on the multi-agent game PAC-MAN, using the code of [13]. The ghosts can have different strategies where they take actions based on their own position as well as position of Pac-Man. In our experiments, we used two different types of ghosts, the *random ghosts* (in green) always choose an action uniformly at random from the legal actions available, while the *directional ghosts* (in red) take the legal action that minimises the Manhattan distance to Pac-Man with probability 0.9, and move randomly otherwise.

The game can be seen as a Markov decision process, where states encode a position for each agent⁶ and for the food pills in the grid, where actions encode individual Pac-Man moves, and where stochastic transitions encode the moves of ghosts according to their probabilistic models. For each state and action pair, we define a reward based on the score gained or lost by this move, as explained in the caption of Figure 1. We also assign a terminal reward to each state, so as to allow MCTS to compare paths of length H which would otherwise obtain

⁵ The probability of a path p being sampled should be equal to $w(p) / \sum_{p' \models \mathcal{A}} w(p')$.

⁶ The last action played by ghosts should be stored as well, as they are not able to reverse their direction.

the same score. Intuitively, better terminal rewards are given to states where Pac-Man is closer to the food pills and further away from the ghosts, so that terminal rewards play the role of a static evaluation of positions.

Experiments. We used a receding horizon $H = 10$. The baseline is given by a standard implementation of the algorithm described in Section 3. A search tree is constructed with a maximum depth H , for 100 iterations, so that the search tree constructed by the MCTS algorithm contains up to 100 nodes. At the first selection of every node, 100 samples are obtained by using a uniform policy. Overall, this represents a tiny portion of the tree unfolding of depth 10, which underlines the importance of properly guiding the search to the most interesting neighborhoods. As a point of reference, we also had human players take control of Pac-Man, and computed the same statistics. The players had the ability to slow down the game as they saw fit, as we aimed for a comparison between the quality of the strategic decisions made by these approaches, and not of their reaction speeds.

We compare these baselines with the algorithm of Section 4.1, using the following advice. The *simulation advice* ψ that we consider is defined as a safety property satisfied by every path such that Pac-Man does not make contact with a ghost, as in Example 13. We provide a Boolean formula encoding ψ , so that one can use a SAT solver to obtain samples, or sampling tools as described in Proposition 24, such as WEIGHTGEN [8]. We use UNIGEN [9] to sample almost uniformly over the satisfying assignments of ψ .⁷

From this simulation advice, we extract whenever possible a strongly enforceable *selection advice* φ that guarantees that Pac-Man will not make contact with a ghost, as described in Example 16. If safety cannot be enforced, \top is used as a selection advice, so that no pruning is performed. This is implemented by using the Boolean formula ψ in a QBF solver according to Lemma 23. For performance reasons, we guarantee safety for a smaller horizon $h < 10$, that we fixed at 3 in our experiments.

Several techniques were used to reduce the state-space of the MDP in order to obtain smaller formulæ. For example, a ghost that is too far away with respect to H or h can be safely ignored, and the current positions of the food pills is not relevant for safety.

Results. For each experiment, we ran 100 games in a high-end cluster using AMD Opteron Processors 6272 at 2.1 GHz. A summary of our results is displayed in Table 1. We mainly use the number of games won out of 100 to evaluate the performance of our algorithms.⁸ In the small grid with four random ghosts, the baseline MCTS algorithm wins 17% of games. Adding the selection advice results in a slight increase of the win rate to 25%. The average score is improved as expected, but even if one ignores the ± 500 score associated with a win or a loss, we observe that more food pills were eaten on average as well. The simulation advice provides in turn a sizeable increase in both win rate (achieving 71%) and average score. Using both advice at the same time gave the best results overall, with a win rate of 85%. The same observations can be made in other settings as well, either with a directional ghost model or on a large grid. Moreover, the simulation advice significantly reduces the number of game turns Pac-Man needs to win, resulting in fewer game draws, most notably on the large grid.

⁷ The distribution over path is slightly different than when sampling uniformly over actions in the pruned MDP $T(M, s_0, H, \psi)$, but UNIGEN enjoys better performances than WEIGHTGEN.

⁸ We do not evaluate the accuracy in terms of making optimal choices because those cannot be computed due to the size of the MDPs (about 10^{16} states).

■ **Table 1** Summary of experiments with different ghost models, algorithms and grid size. The win, loss and draw columns denote win/loss/draw rates in percents (the game ends in a draw after 300 game steps). The food eaten column refers to the number of food pills eaten on average, out of 25 food pills in total. Score refers to the average score obtained over all runs.

Grid	Ghosts	Algorithm	win	loss	draw	food	score
9 x 21	4 x Random	MCTS	17	59	24	16.65	-215.32
		MCTS+Selection advice	25	54	21	17.84	-146.44
		MCTS+Simulation advice	71	29	0	22.11	291.80
		MCTS+both advice	85	15	0	23.42	468.74
		Human	44	56	0	18.87	57.76
	1 x Directional + 3 x Random	MCTS	11	85	4	14.86	-339.99
		MCTS+Selection advice	16	82	2	15.25	-290.6
		MCTS+Simulation advice	27	70	3	17.14	-146.79
		MCTS+both advice	33	66	1	17.84	-92.47
		Human	24	76	0	15.10	-166.28
27 x 28	4 x Random	MCTS	1	10	89	14.85	-182.77
		MCTS+both advice	95	5	0	24.10	517.04

In the baseline experiments without any advice, computing the next action played by Pac-Man from a given state takes 200 seconds of runtime on average.⁹ Comparatively, the algorithm with both advice is about three times slower than the baseline. If we make use of the same time budget in the standard MCTS algorithm (roughly increasing the number of nodes in the MCTS tree threefold), the win rate climbs to 26%, which is still significantly below the 85% win rate achieved with advice. Moreover, while this experiment was not designed to optimise the performance of these approaches in terms of computing time, we were able to achieve a performance of 5s per move on a standard laptop by reducing the number of iterations and samples in MCTS. This came at the cost of a decreased win-rate of 76% with both advice. Further code improvements *e.g.* using parallelism as in [10] could reasonably lead to real-time performances.

Supplementary material is available at <http://di.ulb.ac.be/verif/debraj/pacman/>.

6 Conclusion and future works

In this paper, we have introduced the notion of symbolic advice to guide the selection and the simulation phases of the MCTS algorithm. We have identified sufficient conditions to preserve the convergence guarantees offered by the MCTS algorithm while using symbolic advice. We have also explained how to implement them using SAT and QBF solvers in order to apply symbolic advice to large MDP defined symbolically rather than explicitly. We believe that the generality, flexibility and precision offered by logical formalism to express symbolic advice in MCTS can be used as the basis of a methodology to systematically inject domain knowledge into MCTS. We have shown that domain knowledge expressed as simple symbolic advice (safety properties) improves greatly the efficiency of the MCTS algorithm in the PAC-MAN application. This application is challenging as the underlying MDPs have huge state spaces, *i.e.* up to 10^{23} states. In this application, symbolic advice allow the MCTS algorithm to reach or even surpass human level in playing.

⁹ This holds for both the small and large grids, as in both cases we consider the next 10 game steps only, resulting in MCTS trees of similar size.

As further work, we plan to offer a compiler from LTL to symbolic advice, in order to automate their integration in the MCTS algorithm for diverse application domains. We also plan to work on the efficiency of the implementation. So far, we have developed a prototype implementation written in Python (an interpreted language). This implementation cannot be used to evaluate performances in absolute terms but it was useful to show that if the same amount of resources is allocated to the two algorithms the one with advice performs much better. We believe that by using a well-optimised code base and by exploiting parallelism, we should be able to apply our algorithm in real-time and preserve the level of quality reported in the experimental section. Finally, we plan to study how learning can be incorporated in our framework. One natural option is to replace the static reward function used after H steps by a function learned from previous runs of the algorithm and implemented using a neural network (as it is done in AlphaGo [25] for example).

We thank Gilles Geeraerts for fruitful discussions in the early phases of this work.

References

- 1 Mohammed Alshiekh, Roderick Bloem, Rüdiger Ehlers, Bettina Könighofer, Scott Niekum, and Ufuk Topcu. Safe reinforcement learning via shielding. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence, (AAAI 2018)*, pages 2669–2678. AAAI Press, 2018.
- 2 Pranav Ashok, Tomáš Brázdil, Jan Kretínský, and Ondrej Slámecka. Monte Carlo tree search for verifying reachability in Markov decision processes. In *Proceedings of the 8th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA 2018)*, volume 11245 of *Lecture Notes in Computer Science*, pages 322–335. Springer, 2018. doi:10.1007/978-3-030-03421-4_21.
- 3 Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2-3):235–256, 2002. doi:10.1023/A:1013689704352.
- 4 Hendrik Baier and Mark H. M. Winands. MCTS-minimax hybrids. *IEEE Transactions on Computational Intelligence and AI in Games*, 7(2):167–179, 2015. doi:10.1109/TCIAIG.2014.2366555.
- 5 Armin Biere, Keijo Heljanko, Tommi Junttila, Timo Latvala, and Viktor Schuppan. Linear encodings of bounded LTL model checking. *Logical Methods in Computer Science*, Volume 2, Issue 5, November 2006. doi:10.2168/LMCS-2(5:5)2006.
- 6 Tomáš Brázdil, Krishnendu Chatterjee, Martin Chmelik, Vojtech Forejt, Jan Kretínský, Marta Z. Kwiatkowska, David Parker, and Mateusz Ujma. Verification of Markov decision processes using learning algorithms. In *Proceedings of the 12th International Symposium on Automated Technology for Verification and Analysis (ATVA 2014)*, volume 8837 of *Lecture Notes in Computer Science*, pages 98–114. Springer, 2014. doi:10.1007/978-3-319-11936-6_8.
- 7 Cameron Browne, Edward Jack Powley, Daniel Whitehouse, Simon M. Lucas, Peter I. Cowling, Philipp Rohlfshagen, Stephen Tavener, Diego Perez Liebana, Spyridon Samothrakis, and Simon Colton. A survey of Monte Carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in Games*, 4(1):1–43, 2012. doi:10.1109/TCIAIG.2012.2186810.
- 8 Supratik Chakraborty, Daniel J. Fremont, Kuldeep S. Meel, Sanjit A. Seshia, and Moshe Y. Vardi. Distribution-aware sampling and weighted model counting for SAT. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence, 2014 (AAAI 2014)*, pages 1722–1730. AAAI Press, 2014. URL: <http://www.aaai.org/ocs/index.php/AAAI/AAAI14/paper/view/8364>.
- 9 Supratik Chakraborty, Daniel J. Fremont, Kuldeep S. Meel, Sanjit A. Seshia, and Moshe Y. Vardi. On parallel scalable uniform SAT witness generation. In *Proceedings of the 21st International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2015), Held as Part of the European Joint Conferences on Theory and Practice of Software (ETAPS 2015)*, volume 9035 of *Lecture Notes in Computer Science*, pages 304–319. Springer, 2015. doi:10.1007/978-3-662-46681-0_25.

- 10 Guillaume M. J. B. Chaslot, Mark H. M. Winands, and H. Jaap van den Herik. Parallel Monte-Carlo tree search. In H. Jaap van den Herik, Xinhe Xu, Zongmin Ma, and Mark H. M. Winands, editors, *Proceedings of the 6th International Conference on Computers and Games (CG 2008)*, volume 5131 of *Lecture Notes in Computer Science*, pages 60–71. Springer, 2008. doi:10.1007/978-3-540-87608-3_6.
- 11 Krishnendu Chatterjee, Petr Novotný, Guillermo A. Pérez, Jean-François Raskin, and Dorde Zikelic. Optimizing expectation with guarantees in POMDPs. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI 2017)*, pages 3725–3732. AAAI Press, 2017.
- 12 Przemyslaw Daca, Thomas A. Henzinger, Jan Kretínský, and Tatjana Petrov. Faster statistical model checking for unbounded temporal properties. In *Proceedings of the 22nd International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2016), Held as Part of the European Joint Conferences on Theory and Practice of Software (ETAPS 2016)*, volume 9636 of *Lecture Notes in Computer Science*, pages 112–129. Springer, 2016. doi:10.1007/978-3-662-49674-9_7.
- 13 John DeNero and Dan Klein. CS 188 : Introduction to artificial intelligence. URL: <https://inst.eecs.berkeley.edu/~cs188>.
- 14 Mohammadhossein Hasanbeig, Alessandro Abate, and Daniel Kroening. Cautious reinforcement learning with logical constraints. In Amal El Fallah Seghrouchni, Gita Sukthankar, Bo An, and Neil Yorke-Smith, editors, *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems, (AAMAS 2020)*, pages 483–491. International Foundation for Autonomous Agents and Multiagent Systems, 2020. URL: <https://dl.acm.org/doi/abs/10.5555/3398761.3398821>.
- 15 Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. doi:10.1080/01621459.1963.10500830.
- 16 Jing Huang, Zhiqing Liu, Benjie Lu, and Feng Xiao. Pruning in UCT algorithm. In *Proceedings of the International Conference on Technologies and Applications of Artificial Intelligence (TAAI 2010)*, pages 177–181, 2010. doi:10.1109/TAAI.2010.38.
- 17 Michael J. Kearns, Yishay Mansour, and Andrew Y. Ng. A sparse sampling algorithm for near-optimal planning in large Markov decision processes. *Machine Learning*, 49(2-3):193–208, 2002. doi:10.1023/A:1017932429737.
- 18 Levente Kocsis and Csaba Szepesvári. Bandit based Monte-Carlo planning. In *Proceedings of the 17th European Conference on Machine Learning (ECML 2006)*, volume 4212 of *Lecture Notes in Computer Science*, pages 282–293. Springer, 2006. doi:10.1007/11871842_29.
- 19 Jan Kretínský, Guillermo A. Pérez, and Jean-François Raskin. Learning-based mean-payoff optimization in an unknown MDP under omega-regular constraints. In *Proceedings of the 29th International Conference on Concurrency Theory (CONCUR 2018)*, volume 118 of *Leibniz International Proceedings in Informatics*, pages 8:1–8:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.
- 20 Christof Löding. Infinite games and automata theory. In Krzysztof R. Apt and Erich Grädel, editors, *Lectures in Game Theory for Computer Scientists*, pages 38–73. Cambridge University Press, 2011.
- 21 João Marques-Silva and Sharad Malik. Propositional SAT solving. In Edmund M. Clarke, Thomas A. Henzinger, Helmut Veith, and Roderick Bloem, editors, *Handbook of Model Checking*, pages 247–275. Springer, 2018.
- 22 Nina Narodytska, Alexander Legg, Fahiem Bacchus, Leonid Ryzhyk, and Adam Walker. Solving games without controllable predecessor. In *Proceedings of the 26th International Conference on Computer Aided Verification (CAV 2014)*, volume 8559 of *Lecture Notes in Computer Science*, pages 533–540. Springer, 2014. doi:10.1007/978-3-319-08867-9_35.
- 23 Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley Series in Probability and Statistics. Wiley, 1994. doi:10.1002/9780470316887.

- 24 Ankit Shukla, Armin Biere, Luca Pulina, and Martina Seidl. A survey on applications of quantified Boolean formulas. In *Proceedings of the 31st IEEE International Conference on Tools with Artificial Intelligence (ICTAI 2019)*, pages 78–84. IEEE, 2019. doi:10.1109/ICTAI.2019.00020.
- 25 David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Vedavyas Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy P. Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016. doi:10.1038/nature16961.
- 26 David Silver and Gerald Tesauro. Monte-Carlo simulation balancing. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 945–952, 2009.

A Markov decision processes

One can make an MDP strongly aperiodic without changing the optimal expected average reward and its optimal strategies with the following transition:

► **Definition 25** (Aperiodic transformation [23, Section 8.5.4]). *For an MDP $M = (S, A, P, R)$, we define a new MDP $M_\alpha = (S, A, P_\alpha, R_\alpha)$ for $0 < \alpha < 1$, with $R_\alpha(s, a) = R(s, a)$, $P_\alpha(s, a)(s) = \alpha + (1 - \alpha)P(s, a)(s)$ and $P_\alpha(s, a)(s') = (1 - \alpha)P(s, a)(s')$. Notice that M_α is strongly aperiodic.*

Every finite path in M is also in M_α . Thus for a strategy $\hat{\sigma}$ in M_α , there is a σ in M whose domain is restricted to the paths in M .

► **Proposition 26** ([23, Section 8.5.4]). *Let M be an MDP. M_α is a new MDP generated by applying the aperiodic transformation mentioned above. Then the set of memoryless strategies that optimises the expected average reward in M_α is the same as the set of memoryless strategies that optimises the expected average reward in M . Also from any s , $\text{Val}_M(s) = \text{Val}_{M_\alpha}(s)$.*

► **Definition 27** (Finite horizon unfolding of an MDP). *For an MDP $M = (S, A, P, R, R_T)$, a horizon depth $H \in \mathbb{N}$ and a state s_0 , the unfolding of M from s_0 and with horizon H is a tree-shaped MDP defined as $T(M, s_0, H) = (S' = S_0 \cup \dots \cup S_H, A, P', R', R'_T)$, where for all $i \in [0, H]$, $S_i = \text{Paths}^i(s_0)$. The mappings P' , R' and R'_T are inherited from P , R and R_T in a natural way with additional self-loops at the leaves of the unfolding, so that for all $i \in [0, H]$, $p \in S_i$, $a \in A$ and $p' \in S'$,*

$$P'(p, a)(p') = \begin{cases} P(\text{last}(p), a)(\text{last}(p')) & \text{if } i < H \text{ and } \exists s' \in S, p' = p \cdot as' \\ 1 & \text{if } i = H \text{ and } p' = p \\ 0 & \text{otherwise,} \end{cases}$$

$$R'(p, a) = \begin{cases} R(\text{last}(p), a) & \text{if } i < H \\ 0 & \text{otherwise.} \end{cases}$$

$$R'_T(p) = R_T(\text{last}(p))$$

Proof of Lemma 9. Let us prove that for all $i \in [0, H]$ and all $p \in S_i$,

- $\text{Val}_M^{H-i}(\text{last}(p)) = \text{Val}_{T(M, s_0, H)}^{H-i}(p)$, and
- $\text{opt}_M^{H-i}(\text{last}(p)) = \text{opt}_{T(M, s_0, H)}^{H-i}(p)$.

40:20 Monte Carlo Tree Search Guided by Symbolic Advice for MDPs

We prove the first statement by induction on $H - i$. For $H - i = 0$, for all $p \in S_i$, $\text{Val}_M^{H-i}(\text{last}(p)) = \text{Val}_{T(M,s_0,H)}^{H-i}(p) = R_T(\text{last}(p))$. Assume the statement is true for $H - i = k$, so that for all $p \in S_{H-k}$, $\text{Val}_M^k(\text{last}(p)) = \text{Val}_{T(M,s_0,H)}^k(p)$. Then for all $p \in S_{H-k-1}$, we have for all $a \in A$ and $s \in \text{Supp}(P(\text{last}(p), a))$, $\text{Val}_M^k(s) = \text{Val}_{T(M,s_0,H)}^k(p \cdot as)$. It follows that

$$\begin{aligned} \text{Val}_M^{k+1}(\text{last}(p)) &= \max_{a \in A} (R(\text{last}(p), a) + \sum_s P(\text{last}(p), a) \text{Val}_M^k(s)) \\ &= \max_{a \in A} (R(\text{last}(p), a) + \sum_s P(\text{last}(p), a) \text{Val}_{T(M,s_0,H)}^k(p \cdot as)) \\ &= \text{Val}_{T(M,s_0,H)}^{k+1}(p). \end{aligned}$$

From $\text{Val}_M^{H-i}(\text{last}(p)) = \text{Val}_{T(M,s_0,H)}^{H-i}(p)$ and $\text{opt}_M^H(\text{last}(p)) = \arg \max_{a \in A} (R(\text{last}(p), a) + \sum_s P(\text{last}(p), a) \text{Val}_M^{H-1}(s))$ we derive $\text{opt}_M^{H-i}(\text{last}(p)) = \text{opt}_{T(M,s_0,H)}^{H-i}(p)$. ◀

B UCB

Let $\bar{X}_{a,n} = \frac{1}{n} \sum_{t=1}^n X_{a,t}$ denote the average of the first n plays of action a . Let $\mu_{a,n} = \mathbb{E}[\bar{X}_{a,n}]$. We assume that these expected means eventually converge, and let $\mu_a = \lim_{n \rightarrow \infty} \mu_{a,n}$.

► **Definition 28** (Drift conditions).

- For all $a \in A$, the sequence $(\mu_{a,n})_{n \geq 1}$ converges to some value μ_a .
- There exists a constant $C_p > 0$ and an integer N_p such that for $n \geq N_p$ and any $\delta > 0$, $\Delta_n(\delta) = C_p \sqrt{n \ln(1/\delta)}$, the following bounds hold:

$$\begin{aligned} \mathbb{P}[n\bar{X}_{a,n} \geq n\mu_{a,n} + \Delta_n(\delta)] &\leq \delta, \\ \mathbb{P}[n\bar{X}_{a,n} \leq n\mu_{a,n} - \Delta_n(\delta)] &\leq \delta. \end{aligned}$$

We define $\delta_{a,n} = \mu_{a,n} - \mu_a$. Then, μ^* , μ_n^* , δ_n^* are defined as μ_j , $\mu_{j,n}$, $\delta_{j,n}$ where j is the optimal action.¹⁰ Moreover, let $\Delta_a = \mu^* - \mu_a$.

As $\delta_{a,n}$ converges to 0 by assumption, for all $\epsilon > 0$ there exists $N_0(\epsilon) \in \mathbb{N}$, such that for $t > N_0(\epsilon)$, then $2|\delta_{a,t}| \leq \epsilon \Delta_a$ and $2|\delta_t^*| \leq \epsilon \Delta_a$ for all all suboptimal actions $a \in A$.

The authors start by bounding the number of time a suboptimal action is played:

► **Theorem 29** ([18, Theorem 1]). Consider UCB1 applied to a non-stationary bandit problem with $c_{t,s} = 2C_p \sqrt{\frac{\ln t}{s}}$. Fix $\epsilon > 0$. Let $T_a(n)$ denote number of times action a has been played at time n . Then under the drift conditions, there exists N_p such that for all suboptimal actions $a \in A$,

$$\mathbb{E}[T_a(n)] \leq \frac{16C_p^2 \ln n}{(1-\epsilon)^2 \Delta_a^2} + N_0(\epsilon) + N_p + 1 + \frac{\pi^2}{3}.$$

Let $\bar{X}_n = \sum_{a \in A} \frac{T_a(n)}{n} \bar{X}_{a,T_a(n)}$ denote the global average of payoffs received up to time n . Then, one can bound the difference between μ^* and \bar{X}_n :

► **Theorem 30** ([18, Theorem 2]). Under the drift conditions of Definition 28, it holds that

$$|\mathbb{E}[\bar{X}_n] - \mu^*| \leq |\delta_n^*| + O\left(\frac{|A|(C_p^2 \ln n + N_0(1/2))}{n}\right).$$

¹⁰ It is assumed for simplicity that a single action a is optimal, i.e. maximises $\mathbb{E}[X_{a,n}]$ for n large enough.

The following theorem shows that the number of times an action is played can be lower bounded:

► **Theorem 31** ([18, Theorem 3]). *Under the drift conditions of Definition 28, there exists some positive constant ρ such that after n iterations for all action a , $T_a(n) \geq \lceil \rho \ln(n) \rceil$.*

Then, the authors also prove a tail inequality similar to the one described in the drift conditions, but on the random variable \bar{X}_n instead of $\bar{X}_{a,n}$.

► **Theorem 32** ([18, Theorem 4]). *Fix an arbitrary $\delta > 0$ and let $\Delta_n = 9\sqrt{2n \ln(2/\delta)}$. Let n_0 be such that $\sqrt{n_0} \geq O(|A|(C_p^2 \ln n_0 + N_0(1/2)))$. Then under the drift conditions, for any $n \geq n_0$, the following holds true:*

$$\mathbb{P}[n\bar{X}_n \geq n\mathbb{E}[\bar{X}_n] + \Delta_n(\delta)] \leq \delta$$

$$\mathbb{P}[n\bar{X}_n \leq n\mathbb{E}[\bar{X}_n] - \Delta_n(\delta)] \leq \delta$$

Finally, the authors argue that the probability of making the wrong decision (choosing a suboptimal action) converges to 0 as the number of plays grows:

► **Theorem 33** ([18, Theorem 5]). *Let I_t be the action chosen at time t , and let a^* be the optimal action. Then $\lim_{t \rightarrow \infty} \Pr(I_t \neq a^*) = 0$.*

C MCTS with Simulation

After n iterations of MCTS, we have $\text{total}_n(p) = \sum_{i|I_i(p) \leq n} \text{reward}_{I_i(p)}(p)$ and $\text{total}(p, a) = \sum_{i|I_i(p,a) \leq n} \text{reward}_{I_i(p,a)}(p, a)$.

We use the following observations, derived from the structure of the MCTS algorithm. For all nodes p in the search tree, after n iterations, we have:

$$\begin{aligned} \text{total}_n(p) &= \text{reward}_{I_1(p)}(p) + \sum_{a \in A} \text{total}_n(p, a) \\ \text{total}_n(p, a) &= \sum_{s \in \text{Supp}(P(\text{last}(p), a))} \text{total}_n(p \cdot as) + R(\text{last}(p), a) \cdot \text{count}_n(p, a) \\ \text{value}_n(p) &= \frac{\text{total}_n(p)}{\text{count}_n(p)} \\ \text{count}_n(p) &= 1 + \sum_a \text{count}_n(p, a) \\ \text{count}_n(p, a) &= \sum_s \text{count}_n(p \cdot as) \end{aligned}$$

In the following proof we will abuse notations slightly and conflate the variables and counters used in MCTS with their associated random variables, *e.g.* we write $\mathbb{E}[\text{value}_n(s_0)]$ instead of $\mathbb{E}[V_n(s_0)]$ with $V_n(s_0)$ a random variable that represents the value $\text{value}_n(s_0)$.

Proof of Lemma 12. We use the following inequality (Chernoff-Hoeffding inequality)[15, Theorem 2] throughout the proof:

Let X_1, X_2, \dots, X_n be independent random variables in $[0, 1]$. Let $S_n = \sum_n X_i$. Then for all $a > 0$, $\mathbb{P}\left[S_n \geq \mathbb{E}[S_n] + t\right] \leq \exp\left(-\frac{2t^2}{n}\right)$ and $\mathbb{P}\left[S_n \leq \mathbb{E}[S_n] - t\right] \leq \exp\left(-\frac{2t^2}{n}\right)$.

40:22 Monte Carlo Tree Search Guided by Symbolic Advice for MDPs

We need to show that the following conditions hold:

1. $\lim_{\text{count}_n(p) \rightarrow \infty} \mathbb{E}[\text{value}_n(p, a)]$ exists for all a .
2. There exists a constant $C_p > 0$ such that for $\text{count}_n(p, a)$ big enough and any $\delta > 0$, $\Delta_{\text{count}_n(p, a)}(\delta) = C_p \sqrt{\text{count}_n(p, a) \ln(1/\delta)}$, the following bounds hold:

$$\mathbb{P} \left[\text{total}_n(p, a) \geq \mathbb{E}[\text{total}_n(p, a)] + \Delta_{\text{count}_n(p, a)}(\delta) \right] \leq \delta$$

$$\mathbb{P} \left[\text{total}_n(p, a) \leq \mathbb{E}[\text{total}_n(p, a)] - \Delta_{\text{count}_n(p, a)}(\delta) \right] \leq \delta$$

We show it by induction on $H - |p|$.

For $|P| = H - 1$: $\text{reward}_i(p, a)$ follows a stationary distribution: $\text{reward}_i(p, a) = R(\text{last}(p), a) + R_T(s)$ with probability $P(\text{last}(p), a)(s)$. Thus

$$\begin{aligned} \mathbb{E}[\text{total}_n(p, a)] &= \mathbb{E} \left[\sum_{i|I_i(p, a) \leq n} \text{reward}_{I_i(p, a)}(p, a) \right] \\ &= \text{count}_n(p, a) \left(\sum_s R_T(s) P(\text{last}(p), a)(s) + R(\text{last}(p), a) \right). \end{aligned}$$

Thus $\mathbb{E}[\text{value}_n(p, a)] = \sum_s R_T(s) P(\text{last}(p), a)(s) + R(\text{last}(p), a)$.

From the Chernoff-Hoeffding inequality,

$$\begin{aligned} \mathbb{P} \left[\sum_{i|I_i(p, a) \leq n} \text{reward}_{I_i(p, a)}(p, a) \geq \mathbb{E} \left[\sum_{i|I_i(p, a) \leq n} \text{reward}_{I_i(p, a)}(p, a) \right] + \sqrt{\frac{\text{count}_n(p, a)}{2} \ln \frac{1}{\delta}} \right] &\leq \delta, \\ \mathbb{P} \left[\sum_{i|I_i(p, a) \leq n} \text{reward}_{I_i(p, a)}(p, a) \leq \mathbb{E} \left[\sum_{i|I_i(p, a) \leq n} \text{reward}_{I_i(p, a)}(p, a) \right] - \sqrt{\frac{\text{count}_n(p, a)}{2} \ln \frac{1}{\delta}} \right] &\leq \delta. \end{aligned}$$

Therefore, condition 2 also holds with $C_p = \frac{1}{\sqrt{2}}$.

Assume that the conditions are true for all $p \cdot as$. Then, from Theorem 30 we get:

$$\begin{aligned} &\left| \mathbb{E} \left[\frac{\sum_{a'} \text{total}_n(pas, a')}{\sum_{a'} \text{count}_n(pas, a')} \right] - \lim_{\text{count}_n(p \cdot as) \rightarrow \infty} \mathbb{E}[\text{value}_n(p \cdot as, a^*)] \right| \\ &\leq \left| \mathbb{E}[\text{value}_n(p \cdot as, a^*)] - \lim_{\text{count}_n(p \cdot as) \rightarrow \infty} \mathbb{E}[\text{value}_n(p \cdot as, a^*)] \right| + \mathcal{O} \left(\frac{\ln(\text{count}_n(p \cdot as) - 1)}{\text{count}_n(p \cdot as) - 1} \right), \end{aligned}$$

where a^* is the optimal action from $p \cdot as$. Now,

$$\begin{aligned} \lim_{\text{count}_n(p) \rightarrow \infty} \mathbb{E}[\text{value}_n(p \cdot as)] &= \lim_{\text{count}_n(p) \rightarrow \infty} \mathbb{E} \left[\frac{\text{total}_n(p \cdot as)}{\text{count}_n(p \cdot as)} \right] \\ &= \lim_{\text{count}_n(p) \rightarrow \infty} \mathbb{E} \left[\frac{\text{total}_n(p \cdot as) - \text{reward}_{\mathcal{I}_1(p)}(p \cdot as)}{\text{count}_n(p \cdot as) - 1} \right] \\ &= \lim_{\text{count}_n(p) \rightarrow \infty} \mathbb{E} \left[\frac{\sum_{a'} \text{total}_n(p \cdot as, a')}{\sum_{a'} \text{count}_n(p \cdot as, a')} \right]. \end{aligned}$$

Let $\lim_{\text{count}_n(p \cdot as) \rightarrow \infty} \mathbb{E}[\text{value}_n(p \cdot as, a^*)]$ be denoted by $\mu_{p \cdot as}$ (we know that this limit exists by the induction hypothesis). From Theorem 31, we know that $\text{count}_n(p, a) \rightarrow \infty$ for all a when $\text{count}_n(p) \rightarrow \infty$. And as for all states s , state s is chosen according to distribution $P(p, a)(s)$, $\text{count}_n(p \cdot as) \rightarrow \infty$ with probability 1. Then,

$$\begin{aligned} \lim_{\text{count}_n(p) \rightarrow \infty} \mathbb{E}[\text{value}_n(p \cdot a)] &= \lim_{\text{count}_n(p) \rightarrow \infty} \mathbb{E} \left[\sum_s \text{value}_n(p \cdot as) \frac{\text{count}_n(p \cdot as)}{\text{count}_n(p, a)} + R(\text{last}(p), a) \right] \\ &= R(\text{last}(p), a) + \sum_s \mu_{p \cdot as} \cdot P(\text{last}(p), a)(s). \end{aligned}$$

So $\lim_{\text{count}_n(p) \rightarrow \infty} \mathbb{E}[\text{value}_n(p \cdot a)]$ exists.

From Theorem 32, when $\text{count}_n(p \cdot as)$ is big enough, for all $\delta > 0$, $\mathbb{P} \left[\sum_{a'} \text{total}_n(pas, a') \geq \mathbb{E}[\sum_{a'} \text{total}_n(pas, a')] + \Delta_1^s(\delta) \right] \leq \frac{\delta}{2|S|}$ where $\Delta_1^s(\delta) = 9 \left(\sqrt{\text{count}_n(p \cdot as) \ln \left(\frac{4|S|}{\delta} \right)} \right)$.

Therefore $\mathbb{P} \left[\text{total}_n(p \cdot as) - \text{reward}_{\mathcal{I}_1(p \cdot as)}(p \cdot as) \geq \mathbb{E}[\text{total}_n(p \cdot as) - \text{reward}_{\mathcal{I}_1(p \cdot as)}(p \cdot as)] + \Delta_1^s(\delta) \right] \leq \frac{\delta}{2|S|}$. Also the random variable associated to $\text{reward}_{\mathcal{I}_1(p \cdot as)}(p \cdot as)$ following a fixed stationary distribution $f(p)$ in $[0, 1]$. So from the Chernoff-Hoeffding inequality, $\mathbb{P} \left[\text{reward}_{\mathcal{I}_1(p \cdot as)}(p \cdot as) \geq \mathbb{E}[\text{reward}_{\mathcal{I}_1(p \cdot as)}(p \cdot as)] + \Delta_2(\delta) \right] \leq \frac{\delta}{2|S|}$ where $\Delta_2(\delta) = \frac{1}{\sqrt{2}} \left(\sqrt{\ln \left(\frac{2|S|}{\delta} \right)} \right)$.

Now using the fact that for n random variables $\{A_i\}_{i \leq n}$ and n random variables $\{B_i\}_{i \leq n}$, $\mathbb{P}[\sum_i A_i \geq \sum_i B_i] \leq \sum_i \mathbb{P}[A_i \geq B_i]$, we get:

$$\begin{aligned} &\mathbb{P} \left[\text{total}_n(p \cdot as) \geq \mathbb{E}[\text{total}_n(p \cdot as)] + \Delta_1^s(\delta) + \Delta_2(\delta) \right] \\ &\leq \mathbb{P} \left[\text{total}_n(p \cdot as) - \text{reward}_{\mathcal{I}_1(p \cdot as)}(p \cdot as) \geq \mathbb{E}[\text{total}_n(p \cdot as) - \text{reward}_{\mathcal{I}_1(p \cdot as)}(p \cdot as)] + \Delta_1^s(\delta) \right] \\ &\quad + \mathbb{P} \left[\text{reward}_{\mathcal{I}_1(p \cdot as)}(p \cdot as) \geq \mathbb{E}[\text{reward}_{\mathcal{I}_1(p \cdot as)}(p \cdot as)] + \Delta_2(\delta) \right] \leq \frac{\delta}{|S|}. \end{aligned}$$

As $\text{count}_n(p, a) \mathbb{E}[R(\text{last}(p), a)] = \mathbb{E}[\text{count}_n(p, a) \cdot R(\text{last}(p), a)]$,

$$\begin{aligned} &\mathbb{P} \left[\text{total}_n(p, a) \geq \mathbb{E}[\text{total}_n(p, a)] + \sum_s (\Delta_1^s(\delta) + \Delta_2(\delta)) \right] \\ &\leq \sum_s \mathbb{P} \left[\text{total}_n(p \cdot as) \geq \mathbb{E}[\text{total}_n(p \cdot as)] + (\Delta_1^s(\delta) + \Delta_2(\delta)) \right] \leq \delta. \end{aligned}$$

Similarly, when $\text{count}_n(p \cdot as)$ is big enough, for all $\delta > 0$ it holds that $\mathbb{P} \left[\text{total}_n(p \cdot as) - \text{reward}_{\mathcal{I}_1(p \cdot as)}(p \cdot as) \leq \mathbb{E}[\text{total}_n(p \cdot as) - \text{reward}_{\mathcal{I}_1(p \cdot as)}(p \cdot as)] - \Delta_1^s(\delta) \right] \leq \frac{\delta}{2|S|}$ and $\mathbb{P} \left[\text{reward}_{\mathcal{I}_1(p \cdot as)}(p \cdot as) \leq \mathbb{E}[\text{reward}_{\mathcal{I}_1(p \cdot as)}(p \cdot as)] - \Delta_2(\delta) \right] \leq \frac{\delta}{2|S|}$.

Thus $\mathbb{P} \left[\text{total}_n(p, a) \leq \mathbb{E}[\text{total}_n(p, a)] - \sum_s (\Delta_1^s(\delta) + \Delta_2(\delta)) \right] \leq \delta$.

As $\text{count}_n(p \cdot as) \leq \text{count}_n(p, a)$, there exists $C \in \mathbb{N}$ such that for $\text{count}_n(p \cdot as)$ big enough and for all $\delta > 0$:

$$\begin{aligned} \sum_s (\Delta_1^s(\delta) + \Delta_2(\delta)) &\leq C \sum_s \sqrt{\text{count}_n(p \cdot as) \ln \left(\frac{1}{\delta} \right)} \\ &\leq C \sum_s \sqrt{\text{count}_n(p, a) \ln \left(\frac{1}{\delta} \right)} \\ &\leq C|S| \sqrt{\text{count}_n(p, a) \ln \left(\frac{1}{\delta} \right)}. \end{aligned}$$

40:24 Monte Carlo Tree Search Guided by Symbolic Advice for MDPs

So, there is a constant C_p such that for $\text{count}_n(p, a)$ big enough and any $\delta > 0$, it holds that $\Delta_{\text{count}_n(p, a)}(\delta) = C_p \sqrt{\text{count}_n(p, a) \ln(1/\delta)} \geq \sum_s (\Delta_1^s(\delta) + \Delta_2(\delta))$. Therefore, the following bound hold: $\mathbb{P}[\text{total}_n(p, a) \geq \mathbb{E}[\text{total}_n(p, a)] + \Delta_{\text{count}_n(p, a)}(\delta)]$ is upper bounded by $\mathbb{P}[\text{total}_n(p, a) \geq \mathbb{E}[\text{total}_n(p, a)] + \sum_s (\Delta_1^s(\delta) + \Delta_2(\delta))]$. It follows that $\mathbb{P}[\text{total}_n(p, a) \geq \mathbb{E}[\text{total}_n(p, a)] + \Delta_{\text{count}_n(p, a)}(\delta)] \leq \delta$. Similarly, $\mathbb{P}[\text{total}_n(p, a) \leq \mathbb{E}[\text{total}_n(p, a)] - \Delta_{\text{count}_n(p, a)}(\delta)]$ is upper bounded by $\mathbb{P}[\text{total}_n(p, a) \leq \mathbb{E}[\text{total}_n(p, a)] - \sum_s (\Delta_1^s(\delta) + \Delta_2(\delta))]$. It follows that $\mathbb{P}[\text{total}_n(p, a) \leq \mathbb{E}[\text{total}_n(p, a)] - \Delta_{\text{count}_n(p, a)}(\delta)] \leq \delta$.

This proves that for any p , the sequences $(x_{a,t})_{t \geq 1}$ associated with $\text{reward}_{\mathcal{X}_t(p, a)}(p)$ satisfy the drift conditions. \blacktriangleleft