

Strategy Synthesis in Markov Decision Processes Under Limited Sampling Access

Citation for published version (APA):

Baier, C., Dubslaff, C., Wienhöft, P., & Kiebel, S. J. (2023). Strategy Synthesis in Markov Decision Processes Under Limited Sampling Access. In K. Y. Rozier, & S. Chaudhuri (Eds.), *NASA Formal Methods: 15th International Symposium, NFM 2023, Houston, TX, USA, May 16–18, 2023, Proceedings* (pp. 86-103). (Lecture Notes in Computer Science (LNCS); Vol. 13903). Springer. https://doi.org/10.1007/978-3-031-33170-1_6

DOI:

[10.1007/978-3-031-33170-1_6](https://doi.org/10.1007/978-3-031-33170-1_6)

Document status and date:

Published: 03/06/2023

Document Version:

Accepted manuscript including changes made at the peer-review stage

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Strategy Synthesis in Markov Decision Processes Under Limited Sampling Access*

Christel Baier¹, Clemens Dubsiaff^{1,2}, Patrick Wienhöft¹, and Stefan J. Kiebel¹

¹ Centre for Tactile Internet with Human-in-the-Loop (CeTI)

Technische Universität Dresden, Dresden, Germany

{christel.baier,patrick.wienhoeft,stefan.kiebel}@tu-dresden.de

² Eindhoven University of Technology, Eindhoven, The Netherlands

c.dubsiaff@tue.nl

Abstract. A central task in control theory, artificial intelligence, and formal methods is to synthesize reward-maximizing strategies for agents that operate in partially unknown environments. In environments modeled by *gray-box* Markov decision processes (MDPs), the impact of the agents' actions are known in terms of successor states but not the stochastics involved. In this paper, we devise a strategy synthesis algorithm for gray-box MDPs via reinforcement learning that utilizes *interval MDPs* as internal model. To compete with limited sampling access in reinforcement learning, we incorporate two novel concepts into our algorithm, focusing on rapid and successful learning rather than on stochastic guarantees and optimality: *lower confidence bound* exploration reinforces variants of already learned practical strategies and *action scoping* reduces the learning action space to promising actions. We illustrate benefits of our algorithms by means of a prototypical implementation applied on examples from the AI and formal methods communities.

1 Introduction

Many machine learning methods take inspiration from the inner-workings of the human brain or human behavior [29]. For instance, learning based on neural networks mimics the human brain at a structural level by explicitly modeling its neurons and their activation. Taking a more high-level view, *reinforcement learning (RL)* [37] formalizes human learning behavior by reinforcing actions that are repeatedly associated with successful task solving [3]. The usual application of RL is to learn reward-optimizing strategies in environments modeled as *Markov decision processes (MDPs)* [31] where the agent has only partial knowledge and learns based on guided exploration through sample runs. Existing RL approaches prioritize stochastic guarantees and convergence to a globally optimal strategy, leading to slow learning performance and infeasibility for small sample sizes [37].

* The authors are supported by the DFG through the Cluster of Excellence EXC 2050/1 (CeTI, project ID 390696704, as part of Germany's Excellence Strategy) and the TRR 248 (see <https://perspicuous-computing.science>, project ID 389792660).

In contrast, human decision making can compete with limited sampling access, not focusing on strict optimality but on efficiency. The more urgent a task and the less time available for its solving, the more humans tend to exploit previously learned strategies – possibly sacrificing optimality but increasing the chance of finishing the task in time [32]. In the extremal case, humans rely on *habits* [42], i.e., sequences of actions that, once triggered, are executed mostly independent from reasoning about the actual task [7]. Habits avoid further costly exploration during learning by restricting the action space.

In this paper, we take inspiration from humans’ ability to reason efficiently with few explorations, shaping novel RL algorithms that rapidly synthesize “good” strategies. Specifically, our learning task amounts to an agent being able to determine a strategy with high expected accumulated reward until reaching a goal, given a limited number of samples. We consider the setting where the environment is modeled as *contracting MDP*, i.e., goal states are almost surely reached under all strategies, on which the agent has a *gray-box* view, i.e., knows the reward structures and the topology but not the exact probabilities [4]. We tackle this task of sample-bounded learning towards nearly-optimal strategies by introducing two new concepts: *lower confidence bound (LCB) sampling* and *action scoping*. Classical reward-based sampling in RL is based on upper confidence bounds (UCB) [2], balancing the *exploration-exploitation dilemma* [37]. In contrast, our LCB sampling method favors situations already shown viable during the learning process. Hence, exploration is limited when there are no good reasons for leaving well-known paths, similar to what humans do with habitual sequences of actions [42]. The second learning component is *action scoping*, restraining exploration actions when shown to be suboptimal in past samples. Scoping is parametrized to tune the degree of exploration and balance between fast strategy synthesis or increasing the chance of learning optimal strategies.

To implement our novel concepts, we provide technical contributions by presenting an RL algorithm on contracting gray-box MDPs with arbitrary rewards and various sampling methods. The learning algorithm is a sample-based approach that generates an *interval MDP (IMDP)* to approximate the environment and whose intervals are iteratively refined. While methods for analyzing IMDPs have already been considered in the literature [17,43], we provide a new connection of their use in RL algorithms. We devise our human-inspired RL algorithms, including LCB and action scoping, by modeling knowledge of the agent as IMDP using concepts from *model-based interval estimation (MBIE)* [35] and *probably almost correct (PAC) statistical model checking (SMC)* [4]. We show that our algorithms on IMDPs are PAC for UCB and LCB sampling, i.e., the probability of a suboptimal strategy can be quantified by an arbitrarily small error tolerance. This, however, cannot be guaranteed in the case of action scoping. Towards an evaluation of LCB and action scoping, we implemented our algorithms in a prototypical tool [1]. By means of several experimental studies from the RL and formal-methods community, e.g., on multi-armed bandits [40] and RACETRACK [10], we show that LCB and action scoping foster fast strategy synthesis, providing better strategies after fewer sample runs than RL-style PAC-SMC methods.

We discuss the impact of scoping parameters and related heuristics, as well as combinations of sampling strategies. In summary, our contributions are:

- A (model-based) RL algorithm for contracting gray-box MDPs with integer rewards relying on IMDP and sampling strategy refinements (see Section 3)
- Instances of this RL algorithm subject to lower and upper confidence bound sampling and tunable action scoping (see Section 4).
- A prototypical implementation of our RL algorithms and an evaluation in examples from both the RL and formal-methods communities.

Related work. SMC [28] for unbounded temporal properties in stochastic systems is most related to our setting, establishing algorithms also in gray-box settings [44,20]. Given a lower bound on transition probabilities, SMC algorithms have been presented for Markov chains [15], MDPs, and even stochastic games [4]. Recent SMC algorithms for MDPs also include learning [13,4] but only for reachability problems. IMDPs have been investigated outside of the RL context in formal verification [33,14] for ω -regular properties, for positive rewards in contracting models [43] by an extension of the well-known value-iteration algorithm [37], and in the performance-evaluation community in the discounted setting [17]. In particular, the RL algorithms we present in this paper use an adaptation of the latter algorithm without discounting as a subroutine to successively tighten bounds on the maximal expected accumulated rewards. More recently, algorithms with convergence guarantees for reachability objectives in (interval) MDPs have been presented [6,19]. *Interval estimation* for RL has been introduced by Kaelbling [24] towards Q-learning [39] and extended to model-based approaches [41] such as MBIE [34] and the UCRL2 algorithm [23] using an error tolerance based on the L_1 -norm opposed to the L_∞ -norm employed in interval MDPs. Besides UCB sampling, the exploration-exploitation dilemma in reward-based learning has also be addressed with exploration bonuses [24,25,36,38], performing well when applied to MBIE [22,35] or in other RL methods such as E^3 [27] and R_{\max} [12].

Supplements. The appendix contains proofs and full experimental evaluations. Our implementation and data sets to reproduce the experiments of this paper is available [1].

2 Preliminaries

A *distribution* over a finite set X is a function $\mu: X \rightarrow [0, 1]$ where $\sum_{x \in X} \mu(x) = 1$. The set of distributions over X is denoted by $Dist(X)$.

Markov decision processes (MDPs). An MDP is a tuple $\mathcal{M} = (S, A, \iota, G, R, T)$ where S , A , and $G \subseteq S$ are finite sets of states, actions, and goal states, respectively, $\iota \in S$ is an initial state, $R: S \rightarrow \mathbb{R}$ is a reward function, and $T: S \times A \rightarrow Dist(S)$ is a partial transition probability function. For state $s \in S$

and action $a \in A$ we say that a is *enabled* in s if $T(s, a)$ is defined. We assume the set $Act(s)$ of all enabled actions to be empty in goal states $s \in G$ and non-empty in all other states. For $(s, a, s') \in S \times A \times S$ we define $T(s, a, s') = T(s, a)(s')$ if $T(s, a)$ is defined and $T(s, a, s') = 0$ otherwise. The *successors* of s via a are denoted by $Post(s, a) = \{s' \mid T(s, a, s') > 0\}$. A *run* of \mathcal{M} is a sequence $\pi = s_0 a_0 s_1 a_1 \dots s_n$ where $s_0 = \iota$, $s_n \in S$, $(s_i, a_i) \in (S \setminus G) \times A$, and $s_{i+1} \in Post(s_i, a_i)$ for $i = 0, \dots, n-1$. The set of all runs in \mathcal{M} is denoted by $Runs(\mathcal{M})$. The *accumulated reward* of π is defined by $R(\pi) = \sum_{i=0}^{n-1} R(s_i)$.

An *interval MDP (IMDP)* is a tuple $\mathcal{U} = (S, A, \iota, G, R, \hat{T})$ where S, A, ι, G , and R are as for MDPs, and $\hat{T}: S \times A \rightarrow Intv(S)$ is an *interval transition function*. Here, $Intv(S)$ denotes the set of interval functions $\nu: S \rightarrow \{[a, b] \mid 0 < a \leq b \leq 1\} \cup \{[0, 0]\}$ over S . Note that $Dist(S) \subseteq Intv(S)$, i.e., every distribution over S is also an interval function. A distribution $\mu \in Dist(S)$ is an *instantiation* of $\nu \in Intv(S)$ if $\mu(s) \in \nu(s)$ for all $s \in S$. We again say a is enabled in s if $\hat{T}(s, a)$ is defined and denote the set of enabled actions in s as $Act(s)$, assumed to be non-empty for all $s \in (S \setminus G)$. For each $s \in S$ and $a \in Act(s)$ we denote by T_s^a the set of all instantiations t_s^a of $\hat{T}(s, a)$ and define $Post(s, a) = \{s' \mid \underline{T}(s, a, s') > 0\}$. The MDP \mathcal{M} is an *instantiation* of \mathcal{U} if $T(s, a) \in T_s^a$ for all $s \in S$, $a \in A$. We denote by $[\mathcal{U}]$ the set of all instantiations of \mathcal{U} . Note that as all instantiations of an IMDP \mathcal{U} share the same topology, the set of runs $Runs(\mathcal{M})$ is the same for all instantiations $\mathcal{M} \in [\mathcal{U}]$.

The semantics of the MDP \mathcal{M} is given through *strategies*, i.e., mappings $\sigma: S \rightarrow Dist(A)$ where $\sigma(s)(a) = 0$ for all $a \notin Act(s)$. We call a run $\pi = s_0 a_0 s_1 a_1 \dots s_n$ in \mathcal{M} a σ -run if $\sigma(s_i)(a_i) > 0$ for all $i = 0, \dots, n-1$. The probability of π is defined as $\Pr^\sigma(\pi) = \prod_{i=0}^{n-1} \sigma(s_i)(a_i) \cdot T(s_i, a_i, s_{i+1})$ if π is a σ -run and $\Pr^\sigma(\pi) = 0$ otherwise. The probability of some $B \subseteq Runs(\mathcal{M})$ w.r.t. strategy σ is defined by $\Pr^\sigma(B) = \sum_{\pi \in B} \Pr^\sigma(\pi)$. If $\Pr^\sigma(B) = 1$, then the *expected (accumulated) reward* is defined as $\mathbb{E}^\sigma(B) = \sum_{\pi \in B} \Pr^\sigma(\pi) \cdot R(\pi)$. We call \mathcal{M} *contracting* [26] if $\Pr^\sigma(\diamond G) = 1$ for all strategies σ , i.e., a goal state is almost surely reached for any strategy. The semantics of an IMDP \mathcal{U} is the set of its instantiations $[\mathcal{U}]$. An IMDP \mathcal{U} is *contracting* iff all MDPs in $[\mathcal{U}]$ are contracting. Note that for IMDPs there is also a notion of an operational semantics that lifts strategies to instantiations of transitions [33]. While different in its nature, our contributions of this paper can be easily extended also to the latter semantics.

Value and quality functions. A *value function* $V_{\mathcal{M}}: S \rightarrow \mathbb{R}$ of MDP \mathcal{M} is the solution of the *Bellman equations* [11] given by $V_{\mathcal{M}}(s) = R(s)$ for $s \in G$ and

$$V(s) = R(s) + \max_{a \in Act(s)} \sum_{s' \in S} V_{\mathcal{M}}(s') \cdot T(s, a, s') \quad \text{for } s \notin G.$$

The *quality* $Q_{\mathcal{M}}: S \times A \rightarrow \mathbb{R}$ of \mathcal{M} is defined for all $s \in S$ and $a \in Act(s)$ by

$$Q_{\mathcal{M}}(s, a) = R(s) + \sum_{s' \in Post(s, a)} V_{\mathcal{M}}(s') \cdot T(s, a, s')$$

Intuitively, the quality represents the value of choosing an action a in state s continuing with a reward-maximizing strategy. For an IMDP \mathcal{U} , the value function differs between instantiations, leading to Bellman equations

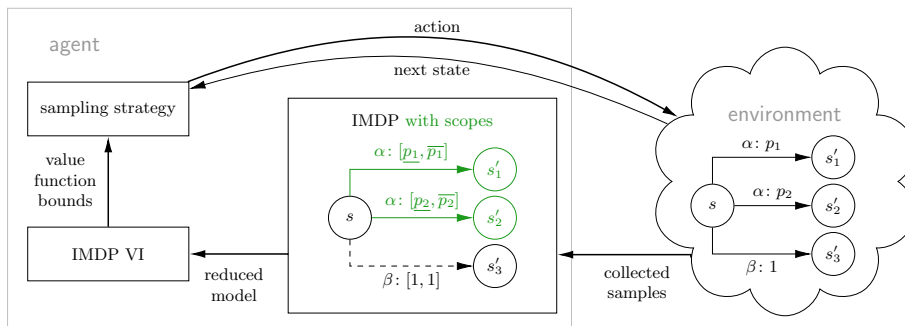


Fig. 1: Schema of reinforcement learning in gray-box MDPs

$$\underline{V}_{\mathcal{U}}(s) = \min_{\mathcal{M} \in [\mathcal{U}]} V_{\mathcal{M}}(s) \quad \bar{V}_{\mathcal{U}}(s) = \max_{\mathcal{M} \in [\mathcal{U}]} V_{\mathcal{M}}(s)$$

for the lower and upper bounds on possible instantiations, respectively. These value functions are naturally lifted to quality functions for IMDPs. We omit subscript \mathcal{M} or \mathcal{U} if clear from the context. Further, we define the *pessimistically optimal strategy* $\underline{\sigma}$ for all $s \in (S \setminus G)$ as $\underline{\sigma}(s) = \arg \max_{a \in \text{Act}(s)} \underline{Q}(s, a)$ and similarly the *optimistically optimal strategy* as $\bar{\sigma}(s) = \arg \max_{a \in \text{Act}(s)} \bar{Q}(s, a)$.

3 Interval MDP reinforcement learning

In this section, we establish an RL algorithm for contracting gray-box MDP that generates a white-box IMDP and successively shrinks the transition probability intervals of the IMDP while updating the sampling strategy. Let $\mathcal{M} = (S, A, \iota, G, R, T)$ be a contracting MDP as above, serving as environmental model. With RL in a *gray-box* setting, the agent's objective is to determine reward-maximizing strategies knowing all components of \mathcal{M} except transition probabilities T . We further make the common assumption [4,6,16] that there is a known constant p_{min} that is a lower bound on the minimal transition probability, i.e., $p_{min} \leq \min\{T(s, a, s') \mid T(s, a, s') > 0\}$.

To learn strategies in \mathcal{M} , samples are generated according to a *sampling strategy* determining the next action an agent performs in each state. Figure 1 shows the overall schema of the algorithm, which runs in *episodes*, i.e., batches of samples. The sampling strategy is updated after each episode by refining an internal IMDP model based on the sample runs and an IMDP value iteration.

3.1 Generating IMDPs from sampled gray-box MDPs

Let $\#(s, a, s')$ denote the number of times the transition (s, a, s') occurred in samples thus far and let $\#(s, a) = \sum_{s' \in \text{Post}(s, a)} \#(s, a, s')$. The goal of each episode is to approximate \mathcal{M} by an IMDP $\mathcal{U} = (S, A, \iota, G, R, \hat{T})$ that is $(1-\delta)$ -correct, i.e., the probability of \mathcal{M} being an instantiation of \mathcal{U} is at least $1-\delta$ for a

Algorithm 1: $\text{IMDP_RL}(\mathcal{M}, \delta, K, N)$

Input : gray-box MDP $\mathcal{M} = (S, A, \iota, G, R, \cdot)$, error tolerance $\delta, K, N \in \mathbb{N}$
Output: pessimistically and optimistically optimal strategies $\underline{\sigma}$ and $\bar{\sigma}$, value function bounds \underline{V} and \bar{V}

```

1 forall  $(s, a) \in S \times A$  do
2    $\sigma(s)(a) := 1/|Act(s)|$  // initialize
3   forall  $s' \in Post(s, a)$  do  $\hat{T}(s, a, s') := [p_{min}, 1]$ 
4  $\mathcal{U} := (S, A, \iota, G, R, \hat{T})$ 
5 forall  $k \in \{1, \dots, K\}$  do
6   forall  $n \in \{1, \dots, N\}$  do SAMPLE( $\mathcal{M}, \sigma$ ) // sample runs
7    $\mathcal{U} := \text{UPDATE\_PROB\_INTERVALS}(\mathcal{U}, \delta)$  // build IMDP model
8    $(\underline{V}, \bar{V}) := \text{COMPUTE\_BOUNDS}(\mathcal{U}, k)$  // IMDP value iteration
9    $\sigma := \text{UPDATE\_STRATEGY}(\mathcal{U}, \underline{V}, \bar{V})$  // compute sampling strategy
10 forall  $s \in (S \setminus G)$  do
11    $(\underline{\sigma}(s), \bar{\sigma}(s)) := (\arg \max_{a \in Act(s)} \underline{Q}(s, a), \arg \max_{a \in Act(s)} \bar{Q}(s, a))$ 
12 return  $\underline{\sigma}, \underline{V}, \bar{\sigma}, \bar{V}$ 

```

given error tolerance $\delta \in \mathbb{R}$. Formally $\prod_{(s,a) \in S \times A} \mathbb{P}(T(s, a) \in \hat{T}(s, a)) \geq 1 - \delta$ [4], where \mathbb{P} refers to the probabilistic behaviour of the algorithm due to sampling the gray-box MDP. The idea towards $(1 - \delta)$ -correct IMDPs is to distribute the error tolerance δ over transitions by defining a *transition error tolerance* $\eta \in \mathbb{R}$. Given a state $s \in S$, an action $s \in Act(s)$ and a successor $s' \in Post(s, a)$, we define the interval transition probability function $\hat{T}_\eta: S \times A \rightarrow Intv(S)$ as

$$\hat{T}_\eta(s, a, s') = \left[\frac{\#(s, a, s')}{\#(s, a)} - c(s, a, \eta), \frac{\#(s, a, s')}{\#(s, a)} + c(s, a, \eta) \right] \cap [p_{min}, 1].$$

where $c(s, a, \eta) = \sqrt{\frac{\log \eta/2}{-2\#(s, a)}}$. Hoeffding's inequality [21] then yields $T(s, a) \in \hat{T}(s, a)$ with probability at least $1 - \eta$. To instantiate an environment approximation, we distribute the error tolerance δ *uniformly*, i.e., to define \hat{T}_η and thus obtain \mathcal{U} we set $\eta = \delta/N_t$ where N_t is the overall number of probabilistic transitions in \mathcal{M} , i.e., $N_t = |\{(s, a, s') \mid s' \in Post(s, a) \text{ and } |Post(s, a)| > 1\}|$. Note that N_t only depends on the topology of the MDP and is thus known in a gray-box setting.

Value iteration on environment approximations. We rely on value iteration for IMDPs [17,43] to solve the interval Bellman equations for all possible instantiations of our environment approximation IMDP \mathcal{U} . Standard value iteration for IMDPs does not exhibit a stopping criterion to guarantee soundness of the results. For soundness, we extend interval value iteration [6,19] with a conservative initialization bound for the value function. For technical details of the value iteration on IMDPs we refer to the appendix.

3.2 Reinforcement learning

Piecing together the parts discussed so far, we obtain an IMDP-based RL algorithm sketched in Algorithm 1 (cf. also Figure 1), comprising K episodes with N sample runs each, updating the model and performing a value iteration (see Line 5 and Line 6, respectively). Both K and N can be seen as parameters limiting the sampling access of the agent. The function SAMPLE in Line 6 interacts with the environment \mathcal{M} and chooses either a yet not sampled action, or samples an action according to σ . A run ends when entering a goal state $s \in G$, or upon reaching a length of $|S|$. The latter is to prevent runs from acquiring a large number of samples by simply staying inside a cycle for as long as possible. In Line 7, the subroutine UPDATE_PROB_INTERVALS incorporates the fresh gathered samples from the environment into the internal IMDP representation as outlined in Section 3.1, updating transition probability intervals. The IMDP value iteration COMPUTE_BOUNDS in Line 8 yields new upper and lower value functions bounds. The number of value iteration steps is $k \cdot |S|$, i.e., increases with each episode to guarantee that the value function is computed with arbitrary precision for a large number of episodes, also known as *bounded value iteration* [4,13]. The computed bounds are then used in UPDATE_STRATEGY in Line 9 to update the sampling strategy for the next episode. The environment approximation \mathcal{U} can be achieved following several strategies according to which samples are generated [2]. A strategy that is widely used in SMC [4] or tabular RL [5,23,34,35] is *upper confidence bound (UCB)* sampling. The UCB strategy samples those actions a in state s that have highest upper bound on the quality $Q(s, a)$, resolving the well-known exploration-exploitation dilemma in RL. This principle is also known as “optimism in the face of uncertainty” (OFU), referring to UCB allocating uncertain probability mass to the best possible outcome [2]. In our framework, standard UCB sampling will serve as the baseline approach. Lastly, we compute and return pessimistic and optimistic strategies along with their value function bounds, before returning them.

Theorem 1. *Let V^* be the solution to the Bellman equations of a given MDP \mathcal{M} . Then for all $\delta \in]0, 1[$ and $K, N \in \mathbb{N}$ the value function bounds \underline{V} and \bar{V} returned by $\text{IMDP_RL}(\mathcal{M}, \delta, K, N)$ as of Algorithm 1 contain V^* with probability at least $1 - \delta$, i.e., $\mathbb{P}(\underline{V}(s) \leq V^*(s) \leq \bar{V}(s)) \geq 1 - \delta$ for all $s \in S$.*

The proof of this theorem is provided in the appendix, essentially showing that (1) $\mathcal{M} \in [\mathcal{U}]$ with probability at least $1 - \delta$ and (2) for each state s the solution to the interval Bellman equation of \mathcal{U} is a subinterval of the computed $[\underline{V}(s), \bar{V}(s)]$.

4 Learning under limited sampling access

Previous work has shown that RL algorithms utilizing the OFU principle converge towards an optimal solution and are also able to perform well in practice [37]. However, they are known to converge rather slowly, requiring lots of sampling data and training time. In this section, we use our IMDP-RL algorithm

presented in Algorithm 1 in a setting where sampling access is limited, i.e., the parameters K and N are fixed. Then, the OFU principle might be not suitable anymore, as the strategy is learnt under an optimistic view for increasing confidence in the actions’ impacts, which requires lots of samples for every action. We propose to focus on finding “good” strategies within the bounded number samples rather than on guaranteed convergence to an optimal strategy. Specifically, we present two complementary methods to reduce the action spaces during sampling: *lower confidence bound sampling* and *action scoping*. Both methods are parametrizable and thus can be adapted to the model size as well as the bound imposed on the number of samples.

4.1 Lower confidence bound sampling

As new sampling strategy incorporated in Line 9 of Algorithm 1, we propose to choose an action a in a state s if it has the highest *lower bound* $\underline{Q}(s, a)$ instead of the highest *upper bound* as within UCB sampling. While then the agent still naturally chooses actions that were already sampled often with high rewards, this avoids further exploring actions with high transition uncertainty. However, such a *lower confidence bound (LCB)* sampling might result in performing exploitations only. Hence, we include an ϵ -greedy strategy [37] into LCB sampling: In each step, with probability $1-\epsilon$ the action with the highest LCB is sampled and with probability ϵ a random action is chosen. In the following, we identify LCB sampling with a degrading ϵ -greedy LCB strategy. Note that also any other exploration strategies, such as sampling with decaying ϵ or *softmax action selection* [37], can easily be integrated into LCB sampling.

While our focus of LCB sampling is on exploiting “good” actions, we can still guarantee convergence towards an optimal strategy in the long run:

Theorem 2. *Algorithm 1 with LCB sampling converges towards an optimal solution, i.e., for $K \rightarrow \infty$ both \underline{V} and \bar{V} converge pointwise towards V^* , and their corresponding strategies $\underline{\sigma}$ and $\bar{\sigma}$ converge towards optimal strategies.*

Similar to how UCB sampling can provide PAC guarantees [4], we can provide PAC guarantees for the value function bounds returned by Algorithm 1 as Theorem 1 guarantees that the solution is in the computed interval with high probability $1 - \delta$ and Theorem 2 guarantees that the interval can become arbitrarily small converging towards the optimal solution from both sides.

4.2 Action scoping

As another approach to compete with resource constraints, we propose to permanently remove unpromising actions from the learned IMDP model, forcing the agent to focus on a subset of enabled actions from the environment MDP. We formalize this idea by setting the *scope* of a state to the set of actions that the agent is allowed to perform in that state.

Scope formation. As depicted in Figure 1, scopes are introduced after each

episode based on the samples of that episode. Initially, all enabled actions are within a scope. Removing an action a from the scope in s is enforced by modifying the interval transition function \hat{T} of \mathcal{U} to the zero interval function at (s, a) , i.e., $\hat{T}(s, a, s') = [0, 0]$ for all $s' \in Post(s, a)$. Scope formation has several notable advantages. First, removing action a from a scope in s reduces the action space $Act(s)$, leading to more sampling data for remaining actions as $\sigma(s)(a) = 0$ for all future episodes. Further, the removal of actions may also reduce the state space in case states are only reachable through specific actions. These positive effects of scoping come at its cost of the algorithm not necessarily converging towards an optimal strategy anymore (cf. Theorem 2). The reason is in possibly removing an optimal action due to unfortunate sampling.

Eager and conservative scopes. We introduce two different scoping schemes: *eager* and *conservative*. Both schemes are tunable by a parameter $h \in]0, 1[$ that specifies the transition error tolerance similar as $\eta = \delta/N_t$ does in our IMDP construction (see Section 3.1). Intuitively, while the formal analysis by means of Line 8 in Algorithm 1 guarantees $1-\delta$ correctness, we allow for different confidence intervals depending on h when forming scopes. Here, greater h corresponds to higher tolerance and hence smaller action scopes.

To define scopes, we introduce $\mathcal{U}_h = (S, A, \iota, G, R, \hat{T}_h)$. That is, \mathcal{U}_h is an IMDP with the same topology as the internal model \mathcal{U} , but allows an error tolerance of h in each transition. We denote the corresponding solution to the interval Bellman equations of \mathcal{U}_h by \underline{V}_h and \overline{V}_h , respectively, and the quality functions as \underline{Q}_h and \overline{Q}_h . Additionally, the *mean quality function* \dot{Q} is computed from the solution of the Bellman equations on the maximum likelihood MDP $\mathcal{M} = (S, A, \iota, G, R, \dot{T})$ where $\dot{T}(s, a, s') = \#(s, a, s')/\#(s, a)$ are the maximum likelihood estimates of the transition probabilities. Both functions can be computed, e.g., using standard value iteration approaches.

In state s an action a is *eagerly* removed from its scope if $\dot{Q}(s, a) < \underline{V}_h(s)$, i.e., if the mean quality of a is lower than the lower bound of the (presumably) best action. The idea is that a is most likely not worth exploring if its expected value is lower than what another action provides with high probability. Likewise, an action a is *conservatively* removed from the scope of a state s if $\overline{Q}_h(s, a) < \underline{V}_h(s)$, i.e., the upper bound quality of a is lower than the lower bound of the (presumably) best action. Here the idea is similar as for eager scoping but with a more cautious estimate on the expected value from action a (observe $\overline{Q}_h(s, a) > \dot{Q}(s, a)$). Note that the parameter h is only used as an error tolerance in \mathcal{U}_h in order to reduce the action scopes. The bound \underline{V} and \overline{V} returned in Algorithm 1 still use an error tolerance of δ/N_t per transition.

5 Implementation and evaluation

To investigate properties of the algorithms presented, we developed a prototypical implementation in PYTHON and conducted several experimental studies, driven by the following research questions:

- (RQ1)** How do UCB and LCB influence the quality of synthesized strategies?
(RQ2) Does action scoping contribute to synthesize nearly optimal strategies when limiting the number of samples?

5.1 Experiment setup

We ran our experiments on various community benchmarks from the formal-methods and RL communities. All our experiments were carried out using PYTHON 3.9 on a MacBook Air M1 machine running macOS 11.5.2. For each system variant and scoping parameter, we learn M strategies (i.e., run the algorithm M times) in $K=50$ episodes each with batch size N as the number of state-action pairs that have a probabilistic successor distribution. Plots show results averaged over the M learned strategies. We chose an error tolerance $\delta = 0.1$, a total of $k \cdot |S|$ value iteration steps in the k -th episode, and an exploration of $\epsilon=0.1$.

Models. For an evaluation, we focus on two models: RACETRACK and multi-armed bandits. Results for all other experiments can be found in the appendix.

In RACETRACK [10,30,18,9], an agent controls a vehicle in a two-dimensional grid where the task is to reach a goal position from some start position, not colliding with wall tiles. Figure 2 depicts two example tracks from Barto et al. [10], which we identify as “small track” (left) and “big track” (right). At each step, the movement in the last step is repeated, possibly modified by 1 tile in either direction, leading to 9 possible actions in each state. Environmental noise is modelled by changing the vehicle position by 1 in each direction with small probability. We formulate RACETRACK as RL problem by assigning goal states with one reward and all other states with zero reward. In the case that the vehicle has to cross wall tiles towards the new position, the run ends, not obtaining any reward. In RACETRACK experiments, we learn $M = 10$ strategies constrained by $N = 940$ sample runs.

The second main model is a variant of multi-armed bandits with one initial state having 100 actions, each with a biased coin toss uniformly ranging from 0.25 to 0.75 probability, gaining one reward and returning to the initial state. Here, we learn $M = 100$ strategies constrained by $N = 101$ sample runs.

5.2 Sampling methods (RQ1)

We investigate the differences of UCB and LCB sampling within RACETRACK.

State-space coverage. UCB and LCB sampling differ notably in covering the state space while learning. With UCB sampling, actions with high uncertainty are more likely to be executed, lowering their upper bound and thus increasing the chance of other actions with higher uncertainty in the next sample run. Hence, UCB sampling leads to exploration of many actions and thus to a high coverage of the state space. In contrast, LCB sampling increases confidence in one particular action shown viable in past samples, leading to sample the same action sequences more often. Hence, LCB sampling is likely to cover only those




Fig. 2: RACE TRACK exploration visualization of sampling methods (tiles colored by start (yellow), goal (green), wall (dark gray), and visit frequency (red-white))

states visited by one successful sampling strategy, showing low coverage of the state space. This can be also observed in our experiments. Figure 2 shows the frequency of visiting positions in the small and big example tracks, ranging from high (red) to low (white) frequencies. Both tracks already illustrate that UCB sampling provides higher state-space coverage than LCB sampling. The small track is symmetric and for each strategy striving towards a lower path, there is a corresponding equally performing strategy towards an upper path. UCB sampling treats both directions equally, while the LCB sampling method in essential learns one successful path and increases its confidence, which is further reinforced in the following samples. reached by one of the symmetric strategies.

Robustness. A further difference of the sampling methods is in dealing with less-explored situations, where UCB sampling is likely to explore new situations but LCB sampling prefers actions that increase the likelihood of returning to known states of an already learned viable strategy. This is due to those states having smaller confidence intervals and thus a greater lower bound on the value and quality functions. Figure 2 shows this effect in the frequency plot of the big track: LCB sampling leads to only few isolated positions with high visit frequencies, while UCB shows a trajectory of visited positions.

Guaranteed bounds. The different characteristics of UCB and LCB sampling can also be observed during the learning process in the small track. In Figure 3

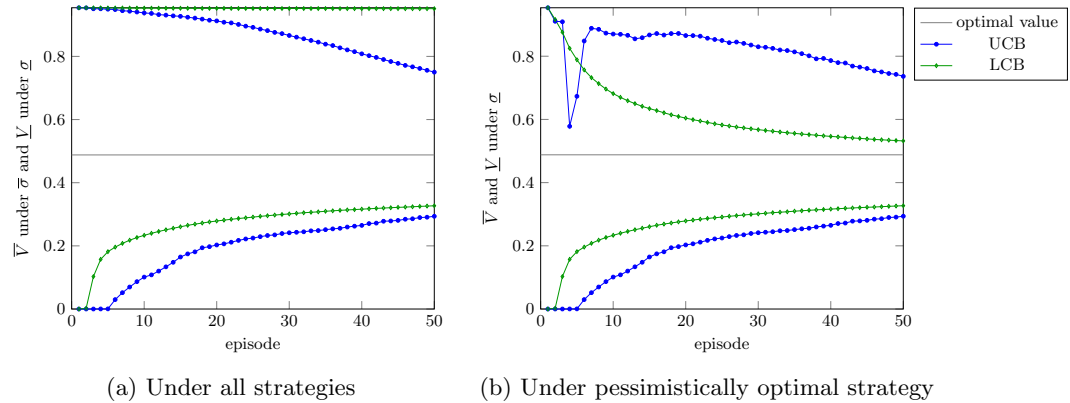


Fig. 3: Comparison of obtained bounds for strategies

on the left we show \underline{V} and \bar{V} after each episode. Note that these bounds apply to different strategies, i.e., the optimistically optimal strategy $\underline{\sigma}$ maximizes \underline{V} , while the pessimistically optimal strategy $\bar{\sigma}$ maximizes \bar{V} . Here, LCB provides values \underline{V} more close to the optimum and, due to its exploitation strategy, gains more confidence in its learned strategy. However, unlike UCB sampling, it cannot improve on \bar{V} significantly, since parts of the environment remain mostly unexplored. We plot bounds under a single fixed strategy $\underline{\sigma}$ on the right. After 50 episodes, UCB then can provide value function bounds $[0.29, 0.75]$, while LCB provides $[0.33, 0.53]$, being more close to the optimal value of 0.49.

LCB is also favorable under limited sampling access, e.g., in (mostly) symmetric environments as the small track: UCB explores the symmetry and requires at least double samples for achieving a similar confidence on the learned strategy.

Concerning **(RQ1)**, we showed that LCB sampling can provide better strategies with high confidence than UCB sampling, while UCB sampling shows better bounds when ranging over all strategies.

5.3 Impact of scoping (RQ2)

We now investigate the impact of action scoping and its parameter h on the multi-armed bandit experiment.

Action-space reduction. Figure 4 shows the number of state-action pairs in the IMDP after each episode w.r.t. UCB and LCB sampling. Here, eager and conservative action scoping is considered with various scoping parameters h . As expected, more actions are removed for greater h . Since $\hat{Q}(s, a) \leq \bar{Q}_h(s, a)$, eager scoping leads to more actions being removed than conservative scoping (cf. eager plots in the lower part of the figures). Observe that the choice of eager or conservative scoping has more impact than the choice of h . In terms of the




Fig. 4: Action-space reduction by action scoping (UCB left, LCB right)




Fig. 5: Bounds of the subsystem obtained by scoping (UCB left, LCB right)

sampling method we observe that for conservative scoping with UCB sampling more actions are removed from scopes than with LCB sampling. A possible explanation is that in LCB sampling, suboptimal actions do not acquire enough samples to sufficiently reduce the upper bound of their expected reward.

Strategy bounds. Next, we investigate the bounds obtained by the strategies returned by Algorithm 1. For brevity, we focus here on the cases $h=\delta/N_t$ and $h=0.05$. Our results are plotted in Figure 11 and Figure 12 for \underline{V} and \bar{V} on the subsystem obtained by applying action scopes with both $\underline{\sigma}$ and $\bar{\sigma}$ and solely $\underline{\sigma}$, respectively. For UCB sampling, we observe that bounds tighten faster the more actions are removed from scopes and reduce the system size, i.e., particularly for eager scoping and for $h=0.05$. For LCB, scopes do not have such a drastic influence, since actions are only leaving the scope if there is an alternative action with high \underline{V} , in which case the latter action is sampled mostly anyway.





Fig. 6: Bounds for pessimistically optimal strategy (UCB left, LCB right)

Fig. 7: Expected total reward w.r.t. sampling strategy σ (UCB left, LCB right)

Sampling strategy quality. In Figure 13 we plot the expected total reward of the employed sampling strategy σ in Algorithm 1 after each episode. Eager scoping tremendously improves the quality of the sampling strategy for both UCB and LCB sampling. For the UCB strategy we observe an initial monotonic increase of the online performance that eventually drops off. This is because a lot of actions cannot improve on the trivial upper bound of 1 until a lot of samples are acquired. In the first roughly 20 episodes the number of actions with trivial upper bound decreases as more samples are collected. Additionally, the actions with upper bound less than 1 are those which were less successful in the past, i.e., likely suboptimal actions, explaining the monotonic increase. Afterwards, the fluctuations within UCB sampling stem from the fact that after reducing all upper bounds below 1, UCB sampling may result in a fully deterministic sampling strategy that only updates after each episode. Hence, exploring a particular action takes the full N runs of an episode, even if the action is sub-

Table 1: Values of the optimal strategy in the multi-armed bandit model

	no	UCB-cons.		UCB-eager		LCB-cons.		LCB-eager	
	scoping	$h=\delta/N_t$	$h=0.05$	$h=\delta/N_t$	$h=0.05$	$h=\delta/N_t$	$h=0.05$	$h=\delta/N_t$	$h=0.05$
value	0.75	0.75	0.75	0.75	0.742	0.75	0.75	0.746	0.739

optimal and only has a high upper bound due to a lack of samples. Especially here, scoping helps to eliminate such actions and avoids sampling them for a full episode just to confirm the action was indeed suboptimal. For LCB sampling the better performance with scoping is due to suboptimal actions being removed and thus not eligible in the exploration step with probability ϵ .

Subsystem bounds. With the introduction of scopes, our RL algorithm is not guaranteed to converge to optimal values. To determine how far the learned strategies depart from the optimum, we compute the optimal strategy within the subsystem generated by only considering actions within the computed scope. Note that for the transition function we are using the exact probabilities as in the environment MDP. The results are given in Table 1. Without scope reduction the subsystem is just the entire environment. When introducing scopes, we did not remove the optimal action via conservative scoping a single time with either sampling method, even for $h = 0.05$. With eager scoping, the optimal action may not be in the scope anymore, but the optimal strategy in the subsystem still performs reasonably well compared to the overall optimal strategy. The fact we observe this only with eager scoping is not surprising, as removing more actions from scopes (recall Figure 4) of course increases the chance of removing the optimal action in a state.

For **(RQ2)**, we conclude that for both UCB and LCB sampling, scoping and especially eager scoping significantly improves the quality of learned strategies after few samples, while only slightly deviating from the optimal strategy. In the UCB setting, scoping leads to further exploitation and thus better bounds.

5.4 Further examples

We ran our algorithms on several other environment MDPs from the RL and formal-methods communities. The bounds of the expected reward obtained for the reduced subsystem, as well as for the strategy maximizing the lower bound, are given in the appendix (cf. Table 3). In general, the strategy learned from LCB yields equal or higher lower bounds and tighter bounds for single strategies, while UCB sampling gives tighter bounds for the entire system. Employing action scoping generally tightens the bounds further with the eager scoping emphasizing this effect. The margins of the differences vary between the examples. In general, both LCB sampling and scoping have biggest impact on large action spaces and on models with high probability deviations such as with small error probabilities.

On the flipside, we observed that LCB performs poorly when mostly or fully deterministic actions or even runs are available, as those incur little uncertainty and thus tend to have relatively large lower bounds even with few samples.

6 Concluding remarks

We devised novel model-based RL algorithms that are inspired by efficient human reasoning under time constraints. Similar to humans tending to stick and return to known situations during strategy learning, LCB exploration favors to return to states with high confidence and proceeding with viable learned strategies. On the action level, scoping implements a reduction of the exploration space as humans do when favoring known actions without further exploration. As for humans acting under resource constraints, both ingredients have been shown to yield better strategies after few sample runs than classical RL methods, especially when choosing high scoping parameters that foster action scoping. While our methods synthesize good strategies faster, an optimal strategy might be not achievable in the limit. We would like to emphasize that also using IMDPs as internal model for RL is a result by itself, incorporating the knowledge about action scopes and confidences in the learned probabilities of the environmental MDP model.

We mainly discussed applications of our techniques in the setting of reinforcement learning. Nevertheless, they can well be utilized also in the formal methods domain, providing an existential variant for statistical model checking of MDPs, asking for the existence of a strategy to reach a goal with accumulating a certain reward.

In future work, we will further investigate adaptations of our algorithms to reflect main properties of human reasoning, e.g., by establishing real-world neuroscientific experiments that involve humans. Our vision is to use formal methods to support explanations of human behaviors, to which we see the present paper providing a first step in this direction [8]. Future adaptations of our proposed algorithm may focus on more sophisticated exploration schemes than dithering, such as softmax selection or exploration bonuses, that allow for deeper exploration of the environment [37]. While not the main focus of this paper, it is well possible to extend our approach also to a black-box setting, i.e., without knowledge about the topology of the MDP, using similar techniques as in [4]. One advantage of using the gray-box setting is in also ensuring applicability to the instance of infinite state MDPs with finitely many actions if this MDP can be effectively explored. For this, it suffices to consider only a finite fragment of the MDP given that our algorithms restrict the sample lengths to a fixed bound.

References

1. https://osf.io/r24mu/?view_only=b44cec578cce44e5920f150940f68230
2. Amin, S., Gomrokchi, M., Satija, H., van Hoof, H., Precup, D.: A survey of exploration methods in reinforcement learning (2021)
3. Anderson, J.R.: Learning and Memory: An Integrated Approach, 2nd Ed. Learning and Memory: An Integrated Approach, 2nd Ed, John Wiley & Sons Inc, Hoboken, NJ, US (2000)
4. Ashok, P., Křetínský, J., Weininger, M.: Pac statistical model checking for markov decision processes and stochastic games. In: Dillig, I., Tasiran, S. (eds.) Computer Aided Verification. pp. 497–519. Springer International Publishing, Cham (2019)
5. Auer, P., Cesa-Bianchi, N., Fischer, P.: Finite-time analysis of the multiarmed bandit problem. *Machine Learning* **47**, 235–256 (2004)
6. Baier, C., Klein, J., Leuschner, L., Parker, D., Wunderlich, S.: Ensuring the reliability of your model checker: Interval iteration for markov decision processes. In: Proc. 28th International Conference on Computer Aided Verification (CAV'17). LNCS, vol. 10426, pp. 160–180. Springer (2017)
7. Baier, C., Cuevas Rivera, D., Dubsloff, C., Kiebel, S.J.: Human-inspired models for tactile computing, chap. 8, pp. 173–200. Academic Press (2021)
8. Baier, C., Dubsloff, C., Funke, F., Jantsch, S., Majumdar, R., Piribauer, J., Ziemek, R.: From Verification to Causality-Based Explications. In: 48th International Colloquium on Automata, Languages, and Programming (ICALP 2021). Leibniz International Proceedings in Informatics (LIPIcs), vol. 198, pp. 1:1–1:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2021)
9. Baier, C., Dubsloff, C., Hermanns, H., Klauck, M., Klüppelholz, S., Köhl, M.A.: Components in probabilistic systems: Suitable by construction. In: Leveraging Applications of Formal Methods, Verification and Validation: Verification Principles. pp. 240–261. Lecture Notes in Computer Science, Springer International Publishing, Cham (2020)
10. Barto, A.G., Bradtke, S.J., Singh, S.P.: Learning to act using real-time dynamic programming. *Artif. Intell.* **72**(1-2), 81–138 (1995)
11. Bertsekas, D.P., Tsitsiklis, J.N.: An analysis of stochastic shortest path problems. *Mathematics of Operations Research* **16**(3), 580–595 (1991). <https://doi.org/10.1287/moor.16.3.580>
12. Brafman, R.I., Tennenholtz, M.: R-max - a general polynomial time algorithm for near-optimal reinforcement learning. *J. Mach. Learn. Res.* **3**(null), 213–231 (Mar 2003). <https://doi.org/10.1162/153244303765208377>
13. Brázdil, T., Chatterjee, K., Chmelník, M., Forejt, V., Křetínský, J., Kwiatkowska, M., Parker, D., Ujma, M.: Verification of markov decision processes using learning algorithms. In: Cassez, F., Raskin, J.F. (eds.) Automated Technology for Verification and Analysis. pp. 98–114. Springer International Publishing, Cham (2014)
14. Chatterjee, K., Sen, K., Henzinger, T.A.: Model-checking omega-regular properties of interval markov chains. pp. 302–317. FOSSACS'08/ETAPS'08, Springer-Verlag, Berlin, Heidelberg (2008)
15. Daca, P., Henzinger, T.A., Křetínský, J., Petrov, T.: Faster statistical model checking for unbounded temporal properties (2016)
16. Daca, P., Henzinger, T.A., Křetínský, J., Petrov, T.: Faster statistical model checking for unbounded temporal properties. *ACM Trans. Comput. Logic* **18**(2) (may 2017). <https://doi.org/10.1145/3060139>

17. Givan, R., Leach, S., Dean, T.: Bounded-parameter markov decision processes. *Artificial Intelligence* **122**(1), 71–109 (2000). [https://doi.org/https://doi.org/10.1016/S0004-3702\(00\)00047-3](https://doi.org/https://doi.org/10.1016/S0004-3702(00)00047-3)
18. Gros, T.P., Hermanns, H., Hoffmann, J., Klauck, M., Steinmetz, M.: Deep statistical model checking. In: Gotsman, A., Sokolova, A. (eds.) *Formal Techniques for Distributed Objects, Components, and Systems*. pp. 96–114. Springer International Publishing, Cham (2020)
19. Haddad, S., Monmege, B.: Interval Iteration Algorithm for MDPs and IMDPs. *Theoretical Computer Science* **735**, 111 – 131 (Jul 2018). <https://doi.org/10.1016/j.tcs.2016.12.003>
20. He, R., Jennings, P., Basu, S., Ghosh, A., Wu, H.: A bounded statistical approach for model checking of unbounded until properties. pp. 225–234 (01 2010)
21. Hoeffding, W.: Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association* **58**(301), 13–30 (1963)
22. Ishii, S., Yoshida, W., Yoshimoto, J.: Control of exploitation–exploration meta-parameter in reinforcement learning. *Neural Networks* **15**(4), 665–687 (2002). [https://doi.org/https://doi.org/10.1016/S0893-6080\(02\)00056-4](https://doi.org/https://doi.org/10.1016/S0893-6080(02)00056-4)
23. Jaksch, T., Ortner, R., Auer, P.: Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research* **11**(51), 1563–1600 (2010)
24. Kaelbling, L.P.: *Learning in Embedded Systems*. The MIT Press (05 1993). <https://doi.org/10.7551/mitpress/4168.001.0001>
25. Kaelbling, L.P., Littman, M.L., Moore, A.W.: Reinforcement learning: A survey. *J. Artif. Int. Res.* **4**(1), 237–285 (may 1996)
26. Kallenberg, L.: *Lecture Notes Markov Decision Problems - version 2020* (02 2020)
27. Kearns, M., Singh, S.: Near-optimal reinforcement learning in polynomial time. *Machine Learning* **49** (01 2002). <https://doi.org/10.1023/A:1017984413808>
28. Legay, A., Lukina, A., Traonouez, L.M., Yang, J., Smolka, S.A., Grosu, R.: *Statistical Model Checking*, pp. 478–504. Springer International Publishing, Cham (2019)
29. Mitchell, T.: *Machine Learning*. McGraw Hill (1997)
30. Pineda, L.E., Zilberstein, S.: Planning under uncertainty using reduced models: Revisiting determinization. In: *ICAPS* (2014)
31. Puterman, M.: *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc. (1994)
32. Schwoebel, S., Markovic, D., Smolka, M.N., Kiebel, S.J.: Balancing control: A Bayesian interpretation of habitual and goal-directed behavior. *Journal of Mathematical Psychology* **100** (FEB 2021). <https://doi.org/10.1016/j.jmp.2020.102472>
33. Sen, K., Viswanathan, M., Agha, G.: Model-checking markov chains in the presence of uncertainties. In: Hermanns, H., Palsberg, J. (eds.) *Tools and Algorithms for the Construction and Analysis of Systems*. pp. 394–410. Springer Berlin Heidelberg, Berlin, Heidelberg (2006)
34. Strehl, A., Littman, M.: An empirical evaluation of interval estimation for markov decision processes. pp. 128– 135 (12 2004). <https://doi.org/10.1109/ICTAI.2004.28>
35. Strehl, A., Littman, M.: An analysis of model-based interval estimation for markov decision processes. *Journal of Computer and System Sciences* **74**, 1309–1331 (12 2008). <https://doi.org/10.1016/j.jcss.2007.08.009>
36. Sutton, R.S.: Dyna, an integrated architecture for learning, planning, and reacting. *SIGART Bull.* **2**(4), 160–163 (Jul 1991). <https://doi.org/10.1145/122344.122377>
37. Sutton, R.S., Barto, A.G.: *Reinforcement Learning: An Introduction*. The MIT Press, second edn. (2018)

38. Thrun, S.B., Möller, K.: Active exploration in dynamic environments. In: Moody, J., Hanson, S., Lippmann, R.P. (eds.) *Advances in Neural Information Processing Systems*. vol. 4. Morgan-Kaufmann (1992)
39. Watkins, C.J.C.H., Dayan, P.: Q-learning. In: *Machine Learning*. pp. 279–292 (1992)
40. Weber, R.: On the Gittins Index for Multiarmed Bandits. *The Annals of Applied Probability* **2**(4), 1024 – 1033 (1992). <https://doi.org/10.1214/aoap/1177005588>
41. Wiering, M., Schmidhuber, J.: Efficient model-based exploration. In: *Proceedings of the sixth intercontinental conference on Simulation of Adaptive Behaviour: From Animals to Animats 6*. pp. 223–228. MIT Press/Bradford Books (1998)
42. Wood, W., Rüdiger, D.: Psychology of Habit. *Annual Review of Psychology* **67**(1), 289–314 (2016). <https://doi.org/10.1146/annurev-psych-122414-033417>
43. Wu, D., Koutsoukos, X.: Reachability analysis of uncertain systems using bounded-parameter markov decision processes. *Artificial Intelligence* **172**(8), 945–954 (2008). <https://doi.org/https://doi.org/10.1016/j.artint.2007.12.002>
44. Younes, H.L.S., Clarke, E.M., Zuliani, P.: Statistical verification of probabilistic properties with unbounded until. In: *SBMF* (2010)

A Appendix

In this appendix, we provide content that had to be omitted from the main paper due to the lack of space.

A.1 Environment approximation

While the approximation approach by Ashok et al. [4] does not explicitly operate on IMDPs, it can be expressed in our setting by

$$\hat{T}(s, a, s') = \left[\max(p_{min}, \frac{\#(s, a, s')}{\#(s, a)} - c'), 1 \right]$$

where $c' = \sqrt{\frac{\log \delta_T}{-2\#(s, a)}}$. The latter is the confidence interval obtained by applying the Hoeffding inequality [21] that guarantees $(1-\delta_T)$ -correctness for each transition triple $(s, a, s') \in S \times A \times S$ when viewing one sample of action a in state s as a Bernoulli trial. Their method takes only the lower bound of the transition probability into account, which we extend to also consider the upper bound.

We now prove the correctness of our definition of

Lemma 1. *Let $\mathcal{M} = (S, A, \iota, G, R, T)$ be an MDP, and $\eta \in]0, 1[$. Assume that for a fixed state $s \in S$ an action $a \in Act(s)$ has been sampled $\#(s, a) \in \mathbb{N}_{>0}$ times and successor s' was observed $\#(s, a, s')$ times. Then, with probability $1 - \eta$ it holds that*

$$T(s, a, s') \in \left[\frac{\#(s, a, s')}{\#(s, a)} - c, \frac{\#(s, a, s')}{\#(s, a)} + c \right]$$

where $c = \sqrt{\frac{\log \eta/2}{-2\#(s, a)}}$

Proof. We can directly bound the probability of $T(s, a, s')$ being outside the interval for a given s, a, s' by applying the two-sided Hoeffding bound, proving the Lemma.

$$\begin{aligned} P \left(\left| T(s, a, s') - \frac{\#(s, a, s')}{\#(s, a)} \right| \geq c \right) &\leq 2 \exp(-2\#(s, a)c^2) \\ &= 2 \exp \left(-2\#(s, a) \sqrt{\frac{\log \eta/2}{-2\#(s, a)}}^2 \right) \\ &= \eta \end{aligned}$$

□

Using this result we can also easily show that $\mathcal{M} \in \mathcal{U} = (S, A, \iota, G, R, T_\eta)$ where $\eta = \delta/N_t$. As \mathcal{M} and \mathcal{U} by definition have the same graph structure and rewards, we only need to check that $T(s, a, s') \in T_\eta(s, a, s')$ for all $s \in S, a \in Act(s)$ and $s' \in Post(s, a)$.

Algorithm 2: MINIMIZING_TRANSITIONS(\mathcal{U}, \hat{V})

Input : IMDP $\mathcal{U} = (S, A, \iota, G, R, \hat{T})$, value function estimate \hat{V}
Output: value minimizing transition function T^-

```

1 forall  $s \in S$  do
2   forall  $a \in Act(s)$  do
3     forall  $s' \in Post(s, a)$  do  $T^-(s, a, s') := \underline{T}(s, a, s')$ 
4     while  $\sum_{s' \in Post(s, a)} T^-(s, a, s') < 1$  do
5       if  $1 - \sum_{u \in Post(s, a)} T^-(s, a, u) > \bar{T}(s, a, s') - \underline{T}(s, a, s')$  then
6         |  $T^-(s, a, s') := \underline{T}(s, a, s')$ 
7       else
8         |  $T^-(s, a, s') := 1 - \sum_{u \in Post(s, a), u \neq s'} T^-(s, a, u)$ 
9 return  $T^-$ 

```

First, note that for a non-probabilistic transition (i.e., s, a with $Post(s, a) = \{s'\}$) the statement also holds trivially as $T(s, a, s') = \frac{\#(s, a, s')}{\#(s, a)} \in T_\eta(s, a, s')$.

Using Lemma 1 the probability that $T(s, a, s')$ is outside the interval for any of the N_t probabilistic transitions can thus be bounded by $\eta N_t = \delta$.

A.2 Building minimizing and maximizing instantiations

A core idea in the adapted value iteration algorithm for IMDPs is to find an instantiation that minimizes or maximizes the value in each state with respect to a previously computed value function [17,43].

Formally, we can state the problem as follows: Given an IMDP $\mathcal{U} = (S, A, \iota, G, R, \hat{T})$ with value function estimate $\hat{V} : S \rightarrow Intv(S)$, choose for each state s and action $a \in Act(s)$ a transition instantiation $t_s^a \in T_s^a$ such that $\hat{Q}(s, a) = \sum_{s' \in Post(s, a)} \hat{V}(s') t_s^a$ is minimized (resp. maximized).

As the instantiation of t_s^a only impacts any $\hat{Q}(s, a)$, and \hat{Q} for no other parameters, this can be optimized locally in a straight forward way. We exemplify this for the minimization case in Algorithm 2. The maximization case is analogous. First, in an effort to ensure that $t_s^a \in T_s^a$, we shift the minimum amount of probability mass into each successor state, i.e., we initialize the value of $t_s^a(s')$ to $\underline{T}(s, a, s')$ for each $s' \in Post(s, a)$. We then distribute the remaining probability mass into the state $s^* \in Post(s, a)$ for which $\hat{V}(s^*)$ is minimal (resp. maximal) until either the total probability mass of t_s^a reaches 1, or until $t_s^a(s^*)$ reaches $\bar{T}(s, a, s')$. If the former is the case, we are done. Otherwise, we repeat the procedure with the state out of the other states for which \hat{V} is minimal (resp. maximal).

As we initialized all transition probabilities with $\underline{T}(s, a, s')$ and only increase them, but to no more than $\bar{T}(s, a, s')$, and as we stop when t_s^a reaches a probability mass of 1 we guarantee that $t_s^a \in T_s^a$. The result also clearly minimizes $\hat{Q}(s, a)$.

Algorithm 3: COMPUTE_BOUNDS(\mathcal{U}, k)

Input : IMDP $\mathcal{U} = (S, A, \iota, G, R, \hat{T})$, parameter k
Output: Lower and upper bound value functions \underline{V} and \overline{V}

- 1 $\underline{V}, \overline{V} = \text{INITIALIZE_BOUNDS}(\mathcal{U})$
- 2 **repeat** $k \cdot |S|$ **times**
- 3 $T^- := \text{MINIMIZING_TRANSITIONS}(\mathcal{U}, \underline{V})$
- 4 $T^+ := \text{MAXIMIZING_TRANSITIONS}(\mathcal{U}, \overline{V})$
- 5 **forall** $s \in S$ **do**
- 6 $\underline{V}'(s) := R(s) + \max_{a \in \text{Act}(s)} \sum_{s' \in S} \underline{V}(s') \cdot T^-(s, a, s')$
- 7 $\overline{V}'(s) := R(s) + \max_{a \in \text{Act}(s)} \sum_{s' \in S} \overline{V}(s') \cdot T^+(s, a, s')$
- 8 $\underline{V} := \underline{V}'$
- 9 $\overline{V} := \overline{V}'$
- 10 **return** $\underline{V}', \overline{V}'$

A.3 Value iteration for IMDPs

We briefly recall the IMDP value-iteration algorithm outlined in Algorithm 3 first developed by Givan et al. [17] for discounted rewards, and later applied by Wu et al. [43] for positive rewards.

An important observation for that is that for any IMDP \mathcal{U} there is an instantiation that is minimizing (or maximizing) the value function for all states simultaneously. The main idea behind the algorithm is then to guess a total order \leq_V over all states $s \in S$ and assume that $s \leq_V s' \iff V(s) \leq V(s')$. Given this assumption one can easily perform a value iteration step by first computing the minimizing transition instantiation for a state s and an action a : Transitions are initially instantiated with the lower bound of $\hat{T}(s, a, s')$ for each successor state $s' \in \text{Post}(s, a)$ and then probability mass is added to the transitions to the successors which are minimal w.r.t. \leq_V until the sum of the instantiated probabilities reaches 1. The computation for the maximizing MDP is analogous. Following that, the order \leq_V is redefined according the value function computed in the VI step.

In total, we perform $k \cdot |S|$ such value iteration steps where k gives the number of episodes that have passed.³

Recall that there is no stopping criterion with formal guarantees. Towards proving that the algorithm can be seen as an any-time algorithm, we first show the following Lemma:

Lemma 2. *Let \underline{V}^* and \overline{V}^* be the solution of the interval Bellman equations of an IMDP \mathcal{U} , respectively, and $([\underline{V}_i(s), \overline{V}_i(s)])_{t \in \mathbb{N}}$ be a sequence of interval value functions s.t. $[\underline{V}_0(s), \overline{V}_0(s)] \supseteq [\underline{V}^*(s), \overline{V}^*(s)]$ for all $s \in S$ and $\dot{V}_{i+1}(s)$ where*

³ Technically any function $f: \mathbb{N} \rightarrow \mathbb{N}$ with the restriction that $\lim_{k \rightarrow \infty} f(k) = \infty$ could be used here, but we choose $k \cdot |S|$ as linear growth in k avoids unnecessarily many iterations and scaling by $|S|$ adapts the algorithm to the environment size

$\dot{V} \in \{\underline{V}, \overline{V}\}$ is computed according to Algorithm 3, i.e.,

$$\begin{aligned}\underline{V}_{i+1}(s) &= R(s) + \max_{a \in \text{Act}(s)} \sum_{s' \in S} \underline{V}_i(s') \cdot T^-(s, a, s') \\ \overline{V}_{i+1}(s) &= R(s) + \max_{a \in \text{Act}(s)} \sum_{s' \in S} \overline{V}_i(s') \cdot T^+(s, a, s')\end{aligned}$$

Then $[\underline{V}_i(s), \overline{V}_i(s)] \supseteq [\underline{V}^*(s), \overline{V}^*(s)]$ for all $i \in \mathbb{N}, s \in S$.

Proof. We first show the statement for the lower bound, i.e., that $\underline{V}_i(s) \leq \underline{V}^*(s)$.

By induction. The base case \underline{V}_0 holds by assumption.

For the induction step, we assume \underline{V}_i satisfies the condition, i.e., $\underline{V}_i(s) \leq \underline{V}^*(s)$ for all $s \in S$.

First, notice that by definition of MINIMIZING_TRANSITIONS, we have $T^-(s, a) = \arg \min_{t_s^a \in T_s^a} \sum_{s' \in S} \underline{V}(s') \cdot t_s^a(s')$. This is the case since the function shifts as much probability mass as possible into states with a low value function in the previous step.

Then, by induction hypothesis and because $t_s^a(s') \geq 0$ the following chain holds:

$$\begin{aligned}\underline{V}^*(s) &= \max_{a \in \text{Act}(s)} \min_{t_s^a \in T_s^a} \sum_{s' \in S} \underline{V}^*(s') \cdot t_s^a(s') \\ &\geq \max_{a \in \text{Act}(s)} \min_{t_s^a \in T_s^a} \sum_{s' \in S} \underline{V}_i(s') \cdot t_s^a(s') \\ &= \underline{V}_{i+1}(s)\end{aligned}$$

The analogous case for upper bounds can be proven in the same way, i.e., for any $i \in \mathbb{N}$ and $s \in S$ we can guarantee that $\overline{V}_i(s) \geq \overline{V}^*(s)$.

From the inequalities for both bounds we can deduce that indeed $[\underline{V}_i(s), \overline{V}_i(s)] \supseteq [\underline{V}^*(s), \overline{V}^*(s)]$ for all $t \in \mathbb{N}$. \square

Note that this requires a safe upper and lower bound, $\underline{V}_0(s)$ and $\overline{V}_0(s)$, of the value function in each state $s \in S$. For reachability tasks the trivial bounds 0 and 1 suffices while for general reward structures the topology of the environment along with p_{min} can be used to compute such bounds [6].

The sequences $(\underline{V}_i)_{i \in \mathbb{N}}$ and $(\overline{V}_i)_{i \in \mathbb{N}}$ are not necessarily monotonic. However, one could adapt the definition to

$$\underline{V}_{i+1}(s) = \max\{\underline{V}_i(s), \underline{W}_i(s)\}$$

with

$$\underline{W}_i(s) = R(s) + \max_{a \in \text{Act}(s)} \sum_{s' \in S} \underline{V}_i(s') \cdot T^-(s, a, s')$$

and analogously (using the minimum between \overline{W} and \overline{V}) for the upper bound, to enforce monotonicity. Under this definition Lemma 2 still holds, as the proof already handles the case $\underline{W}_i(s) \geq \underline{V}_i(s)$, and the case $\underline{W}_i(s) < \underline{V}_i(s)$ can be proven directly from the induction hypothesis as it implies $\underline{V}_{i+1}(s) = \underline{V}_i(s)$

Algorithm 4: UPDATE_LCB_STRATEGY($\mathcal{U}, \underline{V}, \epsilon$)

Input : An IMDP $\mathcal{U} = (S, A, \iota, G, R, \hat{T})$ with lower bounds on its value function \underline{V} , exploration parameter ϵ

Output: a strategy σ

- 1 $\leq_V := \{(s, s') \mid V(s) \leq V(s')\}$
- 2 **forall** $s \in S$ **do**
- 3 **forall** $a \in Act(s)$ **do**
- 4 $T^- := \text{MINIMIZING_TRANSITIONS}(\mathcal{U}, \leq_V)$
- 5 $Q(s, a) := R(s) + \sum_{s' \in S} \underline{V}(s') \cdot T^-(s, a, s')$
- 6 $\sigma(s, a) := \epsilon / |Act(s)|$
- 7 $a^* := \arg \max_{a \in Act(s)} Q(s, a)$
- 8 $\sigma(s, a^*) := \sigma(s, a^*) + (1 - \epsilon)$
- 9 **return** σ

Theorem 1. *Let V^* be the solution to the Bellman equations of a given MDP \mathcal{M} . Then for all $\delta \in]0, 1[$ and $K, N \in \mathbb{N}$ the value function bounds \underline{V} and \bar{V} returned by $\text{IMDP_RL}(\mathcal{M}, \delta, K, N)$ as of Algorithm 1 contain V^* with probability at least $1 - \delta$, i.e., $\mathbb{P}(\underline{V}(s) \leq V^*(s) \leq \bar{V}(s)) \geq 1 - \delta$ for all $s \in S$.*

Proof. By Lemma 1 we have $\mathcal{M} \in [\mathcal{U}]$ with probability at least $1 - \delta$. Assume $\mathcal{M} \in [\mathcal{U}]$ and let \underline{V}^* and \bar{V}^* be the solutions to the interval Bellman equations. Then, clearly $\underline{V}^*(s) \leq V^*(s) \leq \bar{V}^*(s)$ for all $s \in S$ and further by Lemma 2 $\underline{V}_i(s) \leq V^*(s) \leq \bar{V}_i(s)$ for all $t \in \mathbb{N}$ where $\underline{V}_i(s)$ and $\bar{V}_i(s)$ denote the computed lower and upper bounds of value function in the t -th iteration step. Note that Lemma 2 is applicable since we pick safe initial bounds in Line 1 in Algorithm 3.

Hence, $\mathbb{P}(\underline{V}(s) \leq V^*(s) \leq \bar{V}(s)) \geq 1 - \delta$ for all $s \in S$. \square

A.4 LCB sampling

After computing the bounds of the value function for a given dataset over the unknown MDP (cf. Algorithm 3), we can utilize the obtained bounds to define the strategy to be used to sample the MDP in further episodes. We outline this procedure in Algorithm 4. Similar to the idea behind value iteration on IMDPs, we first define a state ordering according to the given lower bounds of the value function. We then also compute the instantiation that minimizes the value function in all states (cf. algorithm 2). This yields the lower bound of the expected reward for each state-action pair. While each action $a \in Act(s)$ is allocated a probability of $\epsilon / |Act(s)|$, the action a^* with the highest lower bound is allocated the the remaining $(1 - \epsilon)$ probability mass of the strategy is state s .

As described in the main part of the paper, the ϵ -sampling serves as an exploration mechanism while at the same time guaranteeing the algorithm converges in the long run.

Theorem 2. *Algorithm 1 with LCB sampling converges towards an optimal solution, i.e., for $K \rightarrow \infty$ both \underline{V} and \bar{V} converge pointwise towards V^* , and their corresponding strategies $\underline{\sigma}$ and $\bar{\sigma}$ converge towards optimal strategies.*

Proof. Let \mathcal{M} be an MDP and \mathcal{U} the IMDP computed according to Line 7, calling Algorithm 3. We first show the claim that as $K \rightarrow \infty$, the computed bound of the value function \underline{V} and \bar{V} converge towards the correct solution to the interval Bellman equations, i.e., for all $s \in S$:

$$\begin{aligned}\underline{V}(s) &= \min_{\mathcal{M} \in [\mathcal{U}]} V_{\mathcal{M}}(s) \\ \bar{V}(s) &= \max_{\mathcal{M} \in [\mathcal{U}]} V_{\mathcal{M}}(s)\end{aligned}$$

As that the main loop, i.e., the value iteration step, is executed $k \cdot |S|$ times in the k -th episode, and the value iteration step is a point-wise contraction mapping (cf. [43, Lemma 5]), we can apply Banach fixed-point theorem to show that the value iteration converges as $k \rightarrow \infty$.

Further, it is known that there are globally reward-minimizing (resp. reward maximizing) MDP \mathcal{M}^- and \mathcal{M}^+ , i.e.,

$$\begin{aligned}\underline{V}(s) &= V_{\mathcal{M}^-}(s) \\ \bar{V}(s) &= V_{\mathcal{M}^+}(s)\end{aligned}$$

for all $s \in S$ [17,43]. As both \mathcal{M}^- and \mathcal{M}^+ have the same graph structure as \mathcal{M} , they are contracting as well and thus there is a unique fixed point of the Bellman equations [11]. Hence, the value iteration converges the the (unique) solution to the Bellman equations.

Next, we show that \underline{V} and \bar{V} converge towards V^* under LCB sampling.

It is well known that, as long as each run has the chance to visit any state, an ϵ -greedy sampling strategy is guaranteed to visit each state infinitely often. More precisely, the probability of reaching any state s in $L \geq |S|$ steps can be bounded from below by a constant $(p_{\min} \frac{\epsilon}{|A|})^{|S|} > 0$. Thus, the expected number of times s is visited as well as the number of times any action $a \in Act(s)$ is sampled in S go towards infinity. Hence, the confidence interval size approaches 0, i.e., $\lim_{t \rightarrow \infty} c(s, a) = 0$. Further, by the law of large numbers, $\lim_{t \rightarrow \infty} \frac{\#(s, a, s')}{\#(s, a)} = T(s, a, s')$ and thus $\lim_{t \rightarrow \infty} \hat{T}(s, a, s') = [T(s, a, s'), T(s, a, s')]$.

In particular, this implies that

$$\lim_{t \rightarrow \infty} \sum_{s \in S} \underline{T}(s, a, s') = \lim_{t \rightarrow \infty} \sum_{s \in S} \bar{T}(s, a, s') = 1$$

and thus the distributions T^- and T^+ in COMPUTE_BOUNDS(\mathcal{U}), converge pointwise towards the true distribution T which means the solution the interval Bellman equations the algorithm is solving coincide with the true Bellman

equations of the environment, and thus \underline{V} and \overline{V} converge pointwise towards the solution of the Bellman equations V^* of the MDP \mathcal{M} .

Lastly, as in the long run $\underline{V} = \overline{V} = V^*$, and $\underline{\sigma}$ and $\overline{\sigma}$ maximize \underline{V} and \overline{V} , respectively, clearly both also maximize V^* and thus converge towards optimal strategies. \square

Such random exploration schemes are also called *dithering*. One well known disadvantage of such approaches is that this unguided exploration does not differentiate between suboptimal actions [37].

Another viable sampling strategy would be to use the softmax action selection where, for each state-action pair (s, a) we define a *temperature* parameter $\tau_{s,a}$ and compute the strategy as

$$\sigma(s)(a) = \frac{e^{\underline{Q}(s,a)/\tau_{s,a}}}{\sum_{a' \in Act(s)} e^{\underline{Q}(s,a')/\tau_{s,a}}}$$

Specifically, this allows of to perform exploration probabilistically, but unlike ϵ -learning introduces a weight to distinguish between non-optimal actions. One problem that arises however is that finding good parameters $\tau_{s,a}$ requires some degree of knowledge about possible action quality values. In a resource constrained setting there is most likely not enough time to find suitable parameters, thus we focus on ϵ -learning in the main part of the paper. If however good estimates of action quality are present a priori, one might want to consider other exploration schemes, such as the softmax sampling strategy.

A.5 Runtime

Table 2: Model sizes and runtimes

model	optimal solution	model size			runtime [s]	
		states	state-action pairs	transitions	no scoping	scoping
Bandit	0.99	3	99	198	0.17	0.25
Bandit25-75	0.75	3	101	202	0.17	0.27
BanditGauss	0.7	3	100	200	0.15	0.24
BettingFav	17.5	235	771	1542	2.2	4.2
BettingUnfav	10	235	771	1542	2.2	3.2
consensus	0.11	272	400	492	23	34
csma	0.50	1038	1054	1282	27	30
Gridworld	-5.5	20	48	132	3.1	5.2
pacman	0.45	6854	8484	8889	7.6	10
RacetrackSmall	0.49	158	1377	4608	13	19
RacetrackBig	0.89	12913	101701	568334	1874	3326
SnL100	-17.41	80	480	1894	99	99
wlan_cost	-220	2954	3972	5202	134	22

In the main evaluation section we abstracted time constraints by limiting each learning approach to 50 episodes with a fixed number of runs each. In Table 2 we give the real-time runtime for each algorithm as well as the expected reward of the optimal strategy and model sizes.

The runtimes are given for the LCB sampling method, for both without scoping, and with a scoping parameter $h = 0.05$ and eager scoping. Introducing scoping increases the runtime of up to a factor of 2, as it is necessary to perform an additional value iteration to compute \underline{V}_h and \overline{V}_h as opposed to only \underline{V} and \overline{V} . We found that neither the sampling method nor the scoping method (i.e., eager or conservative, and choice of h) have a measurable impact on the runtime. The overhead from checking the scoping criterion is negligible. In most examples, the reduction of the model does not yield a significantly faster value iteration since its runtime is mostly dependant on the state-space whereas scopes mostly only reduce the action space. However, in cases where a lot of actions are removed from scopes, and certain states become unreachable, we do indeed observe a significant speedup. This is most emphasized in the wlan_cost example.

In the table we can clearly see a correlation between model size and runtime. The most contributing factors here are the state-space and the topological complexity of the model, i.e., the number of transitions, especially in the presence of loops which causes the value iteration to require more iterations. Even though models with the reachability objective (those with optimal solution $\in [0, 1]$) can initialize the value interval to $[0, 1]$ in each state while models with cost-minimizing objectives (those with negative optimal solutions) have to start with more conservative bounds, there is no clear correlation between the reward structure and runtime.

B Further examples

Here we provide results for further examples of MDPs from the formal methods and RL community. Important model parameters as well as the value function the optimistically optimal strategy σ and pessimistically optimal strategy σ after the full 50 episodes for a scoping parameter of $h = 0.05$ are provided in Table 3. Additionally, we provide plots of the value function bounds for each example after each episode. For the experimental setup and a full description of the plotted values, we refer to Section 5.

Table 3: Optimistic and pessimistic strategy bounds for further examples

model	solution	strategy	UCB			LCB		
			no scoping	conservative	eager	no scoping	conservative	eager
Bandit	0.99	$\bar{\sigma}$	[0.71, 0.98]	[0.75, 0.98]	[0.88, 0.98]	[0.01, 0.99]	[0.04, 0.99]	[0.88, 0.98]
		$\underline{\sigma}$	[0.84, 0.98]	[0.86, 0.98]	[0.92, 0.98]	[0.87, 0.92]	[0.84, 0.89]	[0.92, 0.97]
Bandit25-75	0.75	$\bar{\sigma}$	[0.28, 0.96]	[0.31, 0.95]	[0.71, 0.78]	[0.01, 0.99]	[0.01, 0.99]	[0.61, 0.90]
		$\underline{\sigma}$	[0.63, 0.92]	[0.64, 0.91]	[0.71, 0.78]	[0.61, 0.67]	[0.65, 0.72]	[0.69, 0.75]
BanditGauss	0.7	$\bar{\sigma}$	[0.21, 0.94]	[0.18, 0.93]	[0.66, 0.72]	[0.01, 0.99]	[0.01, 0.99]	[0.48, 0.92]
		$\underline{\sigma}$	[0.57, 0.89]	[0.57, 0.87]	[0.66, 0.72]	[0.54, 0.60]	[0.57, 0.64]	[0.61, 0.68]
BettingFav	17.55	$\bar{\sigma}$	[11.84, 23.67]	[12.13, 23.59]	[12.60, 22.40]	[0.12, 57.53]	[0.16, 57.09]	[0.54, 55.50]
		$\underline{\sigma}$	[13.17, 20.97]	[12.86, 22.01]	[13.18, 21.26]	[11.46, 12.54]	[11.73, 13.09]	[11.49, 13.03]
BettingUnfav	10	$\bar{\sigma}$	[4.47, 11.63]	[4.79, 11.30]	[10.00, 10.00]	[0.00, 53.37]	[0.00, 53.70]	[0.39, 50.07]
		$\underline{\sigma}$	[10.00, 10.00]	[10.00, 10.00]	[10.00, 10.00]	[10.00, 10.00]	[10.00, 10.00]	[10.00, 10.00]
consensus	0.1	$\bar{\sigma}$	[0.03, 0.25]	[0.04, 0.22]	[0.05, 0.19]	[0.00, 0.98]	[0.00, 0.90]	[0.01, 0.67]
		$\underline{\sigma}$	[0.03, 0.25]	[0.04, 0.22]	[0.05, 0.19]	[0.02, 0.14]	[0.03, 0.13]	[0.03, 0.12]
csma	0.5	$\bar{\sigma}$	[0.44, 0.56]	[0.44, 0.56]	[0.44, 0.56]	[0.30, 0.69]	[0.32, 0.70]	[0.40, 0.61]
		$\underline{\sigma}$	[0.44, 0.56]	[0.44, 0.56]	[0.45, 0.56]	[0.46, 0.55]	[0.45, 0.55]	[0.46, 0.55]
Gridworld	-5.48	$\bar{\sigma}$	[-6.26, -5.11]	[-6.23, -5.13]	[-6.21, -5.12]	[-67.53, -4.35]	[-65.07, -4.40]	[-45.05, -4.48]
		$\underline{\sigma}$	[-6.14, -5.13]	[-6.15, -5.13]	[-6.13, -5.13]	[-6.35, -5.17]	[-6.25, -5.15]	[-6.21, -5.11]
pacman	0.45	$\bar{\sigma}$	[0.36, 0.53]	[0.37, 0.53]	[0.36, 0.53]	[0.37, 0.53]	[0.37, 0.53]	[0.36, 0.53]
		$\underline{\sigma}$	[0.36, 0.53]	[0.37, 0.53]	[0.36, 0.53]	[0.37, 0.53]	[0.37, 0.53]	[0.36, 0.53]
RacetrackBig	0.89	$\bar{\sigma}$	[0.00, 0.99]	[0.00, 0.99]	[0.00, 1.00]	[0.00, 1.00]	[0.00, 1.00]	[0.00, 1.00]
		$\underline{\sigma}$	[0.28, 0.98]	[0.61, 0.95]	[0.51, 0.97]	[0.41, 0.87]	[0.43, 0.86]	[0.44, 0.86]
RacetrackSmall	0.49	$\bar{\sigma}$	[0.11, 0.75]	[0.06, 0.75]	[0.18, 0.70]	[0.00, 0.95]	[0.00, 0.95]	[0.00, 0.95]
		$\underline{\sigma}$	[0.29, 0.74]	[0.30, 0.73]	[0.31, 0.69]	[0.33, 0.53]	[0.35, 0.52]	[0.34, 0.52]
SnL100	-17.41	$\bar{\sigma}$	[-20.24, -15.96]	[-19.66, -15.83]	[-19.62, -15.86]	[-60.05, -13.60]	[-62.04, -13.72]	[-49.32, -13.83]
		$\underline{\sigma}$	[-19.41, -15.93]	[-19.42, -15.89]	[-19.38, -15.93]	[-20.03, -16.81]	[-19.84, -16.63]	[-19.86, -16.68]
wlan.cost	-220	$\bar{\sigma}$	[-226.27, -213.56]	[-225.00, -214.79]	[-224.99, -214.77]	[-226.49, -213.34]	[-224.98, -214.72]	[-225.03, -214.78]
		$\underline{\sigma}$	[-226.26, -213.58]	[-224.99, -214.80]	[-224.99, -214.77]	[-226.49, -213.34]	[-224.96, -214.75]	[-225.02, -214.78]

B.1 Bandit




Fig. 8: Bounds of the subsystem obtained by scoping (UCB left, LCB right)




Fig. 9: Bounds for pessimistically optimal strategy (UCB left, LCB right)




Fig. 10: Expected total reward in \mathcal{M} w.r.t. policy σ (UCB left, LCB right)

B.2 Bandit25-75




Fig. 11: Bounds of the subsystem obtained by scoping (UCB left, LCB right)




Fig. 12: Bounds for pessimistically optimal strategy (UCB left, LCB right)




Fig. 13: Expected total reward in \mathcal{M} w.r.t. policy σ (UCB left, LCB right)

B.3 BanditGauss




Fig. 14: Bounds of the subsystem obtained by scoping (UCB left, LCB right)




Fig. 15: Bounds for pessimistically optimal strategy (UCB left, LCB right)




Fig. 16: Expected total reward in \mathcal{M} w.r.t. policy σ (UCB left, LCB right)

B.4 BettingFav




Fig. 17: Bounds of the subsystem obtained by scoping (UCB left, LCB right)




Fig. 18: Bounds for pessimistically optimal strategy (UCB left, LCB right)




Fig. 19: Expected total reward in \mathcal{M} w.r.t. policy σ (UCB left, LCB right)

B.5 BettingUnfav




Fig. 20: Bounds of the subsystem obtained by scoping (UCB left, LCB right)




Fig. 21: Bounds for pessimistically optimal strategy (UCB left, LCB right)




Fig. 22: Expected total reward in \mathcal{M} w.r.t. policy σ (UCB left, LCB right)

B.6 consensus




Fig. 23: Bounds of the subsystem obtained by scoping (UCB left, LCB right)




Fig. 24: Bounds for pessimistically optimal strategy (UCB left, LCB right)




Fig. 25: Expected total reward in \mathcal{M} w.r.t. policy σ (UCB left, LCB right)

B.7 csma




Fig. 26: Bounds of the subsystem obtained by scoping (UCB left, LCB right)




Fig. 27: Bounds for pessimistically optimal strategy (UCB left, LCB right)




Fig. 28: Expected total reward in \mathcal{M} w.r.t. policy σ (UCB left, LCB right)

B.8 Gridworld




Fig. 29: Bounds of the subsystem obtained by scoping (UCB left, LCB right)




Fig. 30: Bounds for pessimistically optimal strategy (UCB left, LCB right)




Fig. 31: Expected total reward in \mathcal{M} w.r.t. policy σ (UCB left, LCB right)

B.9 pacman




Fig. 32: Bounds of the subsystem obtained by scoping (UCB left, LCB right)




Fig. 33: Bounds for pessimistically optimal strategy (UCB left, LCB right)




Fig. 34: Expected total reward in \mathcal{M} w.r.t. policy σ (UCB left, LCB right)

B.10 RacetrackBig




Fig. 35: Bounds of the subsystem obtained by scoping (UCB left, LCB right)




Fig. 36: Bounds for pessimistically optimal strategy (UCB left, LCB right)




Fig. 37: Expected total reward in \mathcal{M} w.r.t. policy σ (UCB left, LCB right)

B.11 RacetrackSmall




Fig. 38: Bounds of the subsystem obtained by scoping (UCB left, LCB right)




Fig. 39: Bounds for pessimistically optimal strategy (UCB left, LCB right)




Fig. 40: Expected total reward in \mathcal{M} w.r.t. policy σ (UCB left, LCB right)

B.12 SnL100




Fig. 41: Bounds of the subsystem obtained by scoping (UCB left, LCB right)




Fig. 42: Bounds for pessimistically optimal strategy (UCB left, LCB right)




Fig. 43: Expected total reward in \mathcal{M} w.r.t. policy σ (UCB left, LCB right)

B.13 wlan_cost




Fig. 44: Bounds of the subsystem obtained by scoping (UCB left, LCB right)




Fig. 45: Bounds for pessimistically optimal strategy (UCB left, LCB right)




Fig. 46: Expected total reward in \mathcal{M} w.r.t. policy σ (UCB left, LCB right)