

Probabilistic Counterexample Guidance for Safer Reinforcement Learning

Xiaotong Ji and Antonio Filieri

Department of Computing
Imperial College London
London, SW7 2AZ, UK
{xiaotong.ji16, a.filieri}@imperial.ac.uk



Abstract. Safe exploration aims at addressing the limitations of Reinforcement Learning (RL) in safety-critical scenarios, where failures during trial-and-error learning may incur high costs. Several methods exist to incorporate external knowledge or to use proximal sensor data to limit the exploration of unsafe states. However, reducing exploration risks in unknown environments, where an agent must discover safety threats during exploration, remains challenging.

In this paper, we target the problem of safe exploration by guiding the training with counterexamples of the safety requirement. Our method abstracts both continuous and discrete state-space systems into compact abstract models representing the safety-relevant knowledge acquired by the agent during exploration. We then exploit probabilistic counterexample generation to construct minimal simulation submodels eliciting safety requirement violations, where the agent can efficiently train offline to refine its policy towards minimising the risk of safety violations during the subsequent online exploration.

We demonstrate our method’s effectiveness in reducing safety violations during online exploration in preliminary experiments by an average of 40.3% compared with QL and DQN standard algorithms and 29.1% compared with previous related work, while achieving comparable cumulative rewards with respect to unrestricted exploration and alternative approaches.

Keywords: Safe reinforcement learning · Probabilistic model checking · Counterexample guidance.

1 Introduction

A critical limitation of applying Reinforcement Learning (RL) in real-world control systems is its lack of guarantees of avoiding unsafe behaviours. At its core, RL is a trial-and-error process, where the learning agent explores the decision space and receives rewards for the outcome of its decisions. However, in safety-critical scenarios, failing trials may result in high costs or unsafe situations and should be avoided as much as possible.

Several learning methods try to incorporate the advantages of model-driven and data-driven methods to encourage safety during learning [27,16]. One natural approach for encouraging safer learning is to analyse the kinematic model of the learning system with specific safety requirements and to design safe exploration [15,2,38] or safe optimisation [1,45,43] strategies that avoid unsafe states or minimise the expected occurrence of unsafe events during training. However, this approach is not applicable for most control systems with partially-known or unknown dynamics, where not enough information is available to characterise unsafe states or events *a priori*.

To increase the safety of learning in environments with (entirely or partially) unknown dynamics, we propose an online-offline learning scheme where online execution traces collected during exploration are used to construct an abstract representation of the visited state-action space. If during exploration the agent violates a safety requirement with unacceptable frequency, a probabilistic model checker is used to produce from the abstract representation minimal counterexample sub-models, i.e., a minimal subset of the abstract state-action space within which the agent is expected to violate its safety requirement with a probability larger than tolerable. These counterexamples are then used to synthesise small-size offline simulation environments within which the agent’s policy can be conveniently reinforced to reduce the probability of reiterating safety-violating behaviours during subsequent online exploration. As new evidence from online exploration is gathered the abstract representation is incrementally updated and additional offline phases can be enforced when necessary, until an acceptable safety exploration rate is achieved. Overall, our strategy aims at migrating most trial-and-error risks to the offline training phases, while discouraging the repeated exploration of risky behaviours during online learning. As new evidence is collected during online exploration, the abstract representation is incrementally updated and the current *value* the agent expects from each action is used to prioritise the synthesis of more relevant counterexample-guided simulations.

Our main conceptual contribution in this paper is the use of probabilistic counterexamples to automatically synthesise small-scale simulation submodels where the agent can refine its policy to reduce the risk of violating a safety requirement during learning. In particular, we 1) propose a conservative geometric abstraction model representing safety-relevant experience collected by the agent at any time during online exploration, with theoretical convergence and accuracy guarantees, suitable for the representation of both discrete and continuous state spaces and finite action spaces, 2) adapt minimal label set probabilistic counterexample generation [49] to generate small-scale submodels for the synthesis of offline agent training environments aimed at reducing the likelihood of violating safety requirements during online exploration, and 3) a preliminary evaluation of our method to enhance Q-Learning [48] and DQN [35] agents on problems from literature and the OpenAI Gym, demonstrating how it achieves comparable cumulative rewards while increasing the exploration safety rate by an average of 40.3% compared with QL/DQN, and of 29.1% compared with previous related work [20,23].

2 Background

2.1 Problem Framework

Definition 1. Markov Decision Process (MDP). An MDP [4] is a tuple (S, A, s_0, P, R, L) , where S is a set of states, A is a finite set of actions, s_0 is the initial state, $P : S \times A \times S \rightarrow [0, 1]$ is the probability of transitioning from a state $s \in S$ to $s' \in S$ with action $a \in A$, $R : S \times A \rightarrow \mathbb{R}$ is a reward function and $L : S \rightarrow 2^{AP}$ is a labelling function that assigns atomic propositions (AP) to each state.

A state in S is typically represented by a vector of finite length $n_S \geq 1$. A state space is *discrete* if the elements of the vector are countable, where we assume $S \subseteq \mathcal{Z}^{n_S}$, continuous if $S \subseteq \mathbb{R}^{n_S}$, or hybrid if some elements are discrete and others are continuous. When possible, we omit the cardinality n_S for readability.

Definition 2. Trace. A finite trace (also called path or trajectory) through an MDP is a sequence $\sigma = s_0, a_0, s_1, a_1, \dots, s_i, a_i, \dots, s_n$, where s_0 is the initial state and $P(s_i, a_i, s_{i+1}) > 0$.

Definition 3. Policy. A (deterministic) policy $\pi : S \rightarrow A$ selects in every state s the action a to be taken by the agent.

Q-learning (QL) [48] is a reinforcement learning algorithm where an agent aims at finding an optimal policy π^* for an MDP that maximises the expected cumulative reward. Given a learning rate $\alpha \in (0, 1]$ and a discount factor $\gamma \in (0, 1]$, such that rewards received after n transitions are discounted by the factor γ^n , the agent learns a value function Q based on the following update rule:

$$Q_t(s, a) = (1 - \alpha)Q_{t-1}(s, a) + \alpha(R(s, a) + \gamma \max_{a' \in A} Q_{t-1}(s', a')) \quad (1)$$

The optimal Q-function Q^* satisfies the Bellman optimality equation:

$$Q^*(s, a) = \mathbb{E}[R(s, a) + \gamma \max_{a' \in A} Q^*(s', a') | s, a] \quad (2)$$

For finite-state and finite-action spaces, QL converges to an optimal policy as long as every state action pair is visited infinitely often [48], but it is not suitable for learning in continuous state spaces. For continuous state spaces instead, the Deep Q-Learning method [35] parameterises Q-values with weights θ as a Q-network and the learning process is adapted to minimising a sequence of loss function L_i at each iteration i (cf. Algorithm 1 in [35]):

$$L_i(\theta_i) = \mathbb{E}[(y_i - Q(s, a; \theta_i))^2], \quad (3)$$

where $y_i = \mathbb{E}[R(s, a) + \gamma \max_{a' \in A} Q(s', a'; \theta_{i-1}) | s, a]$.

During learning, the agent selects the next action among those available in the current state at random with probability $\epsilon_{QL} > 0$ while with probability $1 - \epsilon_{QL}$ it will select an action a yielding $Q^*(s, a)$.

Definition 4. Optimal Policy. An optimal policy π^* , in the context of Q learning and DQN, is given by $\pi^* = \arg \max_{a \in A} Q^*(s, a), \forall s \in S$.

2.2 Probabilistic Model Checking and Counterexamples

Probabilistic model checking is an automated verification method that, given a stochastic model – an MDP in our case – and a property expressed in a suitable probabilistic temporal logic, can verify whether the model complies with the property or not [3]. In this work, we use Probabilistic Computational Temporal Logic (PCTL) [19] to specify probabilistic requirements for the safety of the agent. The syntax of PCTL is recursively defined as:

$$\Phi := \text{true} \mid \alpha \mid \Phi \wedge \Phi \mid \neg \Phi \mid P_{\bowtie p} \varphi \quad (4)$$

$$\varphi := X\Phi \mid \Phi U \Phi \quad (5)$$

A PCTL property is defined by a state formula Φ , whose satisfaction can be determined in each state of the model. true is a tautology satisfied in every state, $\alpha \in AP$ is satisfied in any state whose labels include $(\alpha \in L(s))$, and $\wedge$ and $\neg$ are the Boolean conjunction and negation operators. The modal operator $P_{\bowtie p} \varphi$, with $\bowtie \in \{<, \leq, \geq, >\}$ and $p \in [0, 1]$, holds in a state s if the cumulative probability of all the paths originating in s and satisfying the path formula φ is $\bowtie p$ under any possible policy. The Next operator $X\Phi$ is satisfied by any path originating in s such that the next state satisfies Φ . The Until operator $\Phi_1 U \Phi_2$ is satisfied by any path originating in s such that a state s' satisfying Φ_2 is eventually encountered along the path, and all the states between s and s' (if any) satisfy Φ_1 . The formula $\text{true} U \Phi$ is commonly abbreviated as $F\Phi$ and satisfied by any path that eventually reaches a state satisfying Φ . A model M satisfies a PCTL property Φ if Φ holds in the initial state s_0 of M [3].

PCTL allows specifying a variety of safety requirements. For simplicity, in this work, we focus on safety requirements specified as upper-bounds on the probability of eventually reaching a state labelled as **unsafe**:

Definition 5. Safety Requirement. *Given a threshold $\lambda \in (0, 1]$, the safety requirement for a learning agent is formalised by the PCTL property $P_{\leq \lambda} [F \text{unsafe}]$, i.e., the maximum probability of reaching a $s \in S$ such that $\text{unsafe} \in L(s)$ must be less than or equal to λ .*

Definition 6. Counterexamples in Probabilistic Model Checking. *A counterexample is a minimal possible sub-model $M_{cex} = (S_{cex}, A_{cex}, s_0, P_{cex})$ derived from the model M , where S_{cex}, A_{cex} are of subsets S and A , containing violating behaviours of a PCTL property from the initial state s_0 in M .*

When a model M does not satisfy a PCTL property, a counterexample can be computed as evidence of the violation [18]. In this work, we adapt the *minimal critical label set* counterexample generation method of [49].

The computation of a minimal possible sub-model requires the solution of a mixed-integer linear optimisation problem that selects the smallest number of transitions from the state-action space original model that allows the construction of violations. An extensive description of the counterexample generation algorithm, including a heuristic to bias the counterexample generation towards

including actions that a tabular Q-learning agent is more likely to select is included in [26].

Generating multiple Counterexamples. For an MDP violating the safety requirement, there can exist, in general, multiple counterexamples (both with minimal or non-minimal sizes), each potentially highlighting different policies that lead to requirement violations [10].

In this work, we use counterexamples to guide the generation of offline training environments where the agent learns to reduce the value of actions that may eventually lead to the violation of safety requirements. We therefore aim at generating multiple, *diverse counterexamples* (if they exist), while keeping each of them at a small size for faster training. Given a counterexample, a different one can be obtained by adding a blocking clause to the minimisation problem, i.e., forcing the optimiser to exclude one or more previously selected action pairs (by imposing the corresponding selector variables $x_\ell = 0$ in the optimisation problem given in [26]). Hence, we can systematically add (an increasing number of) blocking clauses to obtain multiple diverse counterexamples that jointly provide a more comprehensive representation of the different violating behaviors the agent explored at any time.

3 Counterexample-guided Reinforcement Learning

We assume that the agent does not have prior knowledge about the environment. In particular, it will only discover unsafe states upon visiting them during exploration. During the learning process the agent will iteratively interact with either the actual environment (*online* exploration) or with a counterexample-guided *offline* simulation. The online phases aim at exploring the actual environment and improving the agent’s policy expected reward, while acquiring information to build and continuously refine an abstract, compact representation of the control problem. The offline phases expose the agent to simulated, small-size, environments, within which the agent can revise its policy to penalise decisions that may lead to safety violations during the subsequent online phases. In the remaining of the section, we first show how to construct and update an abstract finite MDP that compactly represents safety-relevant aspects of the (parts of) environment explored by the agent at any time (sec. 3.1). Then, in sec. 3.2, we introduce the main learning algorithm with the online-offline alternation scheme, and discuss the main challenges of the offline learning phases.

3.1 Safety-relevant State-space Abstraction

For simplicity, let us assume the agent has no information about the topology of the state space S at the beginning of the exploration. Each online interaction with the environment (*episode*) can be described by the trace of the states visited by the agent and the actions it took. Besides the reward associated with each state-action pair, we assume states in the trace can be assigned a set of labels. Labels represent properties of the state related specific to the learning

problem, e.g., a goal has been reached. We assume a special label **unsafe** labels the occurrence of unsafe situations the agent should aim to avoid. W.l.o.g., we assume an episode terminates when the agent enters an **unsafe** state. We will refer to the states of the online environment as *concrete* states.

In this section, we propose an abstraction procedure to construct a finite, abstract MDP that retains sufficient information about the explored concrete environment to enable the synthesis of abstract counterexamples to the safety requirement. Each counterexample will therefore be an abstract representative of a set of possible safety violating behaviors that can happen in the concrete environment. To maintain the size of the abstract MDP tractable – especially in the presence of continuous concrete state spaces – the abstraction will retain only (approximate) safety-relevant information.

We assume that any state not labeled as **unsafe** is **safe** to explore and that the **unsafe** label is time-invariant. Furthermore, the abstraction must preserve at any time a **safety invariant**: *every explored unsafe concrete state should be mapped to an unsafe abstract state*. The finite abstract state space must therefore separate safe and unsafe regions of the concrete space, with only safe concrete states possibly misclassified as unsafe but not vice versa.

Inspired by the idea of casting the learning of geometric concepts as a set cover problem in [8,40], we frame the separation task as a minimal red-blue set cover problem to abstract the explored concrete state space as a finite set of disjoint boxes or polyhedra, each expressed as a set of logical constraints and defined as the intersection of a finite set of hyperplanes.

To formalise the construction of the abstract state-space, we first introduce the notion of coverage of a concrete state by a polyhedra predicate.

Definition 7. Coverage of a polyhedra predicate Let $\bar{S} \subseteq S = \{\mathbf{s}_0, \mathbf{s}_1, \dots, \mathbf{s}_n\}$ be the set of all explored concrete states, a particular state $\mathbf{s} \in \bar{S}$ is covered by a (polyhedra) predicate C_i if $\mathbf{s} \in C_i$, where $C_i = \{\mathbf{s} \mid \boldsymbol{\omega}\mathbf{s} + \mathbf{b} \leq 0\}$, in which $\boldsymbol{\omega}$ represents the vector of slopes and $\mathbf{b}$ represents the vector of biases corresponding to half-spaces enclosing the predicate.

Remark 1. The general affine form of the predicate C_i accounts for a variety of common numerical abstract domains, including, e.g., boxes (intervals), octagons, zonotopes, or polyhedra. In this work, we fix $\omega = 1$, i.e., restrict to hyper-boxes. We allow the user to specify a minimum size $d > 0$, which ensures that no dimension of the box will be reduced to a length smaller than d . This coarsening of the abstract domain struck a convenient trade-off between computational cost and accuracy of the abstraction in our preliminary experiments (see Appendix in [26] for additional discussion); the restriction can be lifted for applications requiring different trade-offs [42].

The identification of a finite set of predicates that allow separating the concrete state space preserving the safety invariant can thus be reduced to the following: **Minimal Red-Blue Set Cover Problem.** Let $\bar{S} \subseteq S = \{\mathbf{s}_0, \mathbf{s}_1, \dots, \mathbf{s}_n\}$ be the set of all explored concrete states and $U = \{\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_m\}$ be the set of explored states assigned the **unsafe** label, find the minimal set $C = \{\cup_{i=1} C_i\}$

s.t. every element $\mathbf{u} \in U$ is covered by some predicate C_i , with an overall false positive rate $\text{fpr} \leq f \in (0, 1]$ for safe concrete states ($\mathbf{s} \in \bar{S} \setminus U$) covered by C .

Remark 2. In general, f cannot be zero, since the concrete state space, whether discrete or continuous, may not be perfectly partitioned by a finite set of polyhedra predicates, with smaller values of f possibly resulting in a larger number of predicates $|C|$.

To solve this optimisation problem, we employ a branch and bound method [30] to systematically enumerate possible combinations of predicates. The solution set guarantees all the unsafe concrete states are covered by a Boolean combination of predicates in C , while safe concrete states may also be covered by some predicate C_i , with a prescribed maximum tolerable rate f .

Definition 8. *Safety-relevant Abstraction MDP.* A safety-relevant abstraction MDP M_a is a tuple $(S_a, A_a, s_{a0}, R_a, P_a, L_a)$, where S_a is the abstract state space, which is the partition of the concrete state space S induced by the boundaries of $C_i \in C$ from the solution of the minimal set cover above, A_a is the set of applicable actions, s_{a0} is the initial abstract state, $P_a : S_a \times A_a \times S_a \rightarrow [0, 1]$ is a probability transition function, R_a is the abstract reward function (which will be defined later), and $L_a : S \rightarrow \{\text{safe}, \text{unsafe}\}$ is a labelling function.

M_a is constructed from the concrete traces collected during online learning, with the satisfaction of the predicates C_i determining the abstraction of concrete states and the abstract transition function is estimated accordingly from the frequencies observed in the concrete traces. The abstraction must preserve the *safety invariant*, therefore it may overapproximate explored unsafe regions, but not underapproximate them. Initially, the entire state space is assumed safe to explore. As new traces are collected during online exploration, the abstract model is incrementally updated: when a concrete state is found to be unsafe, the hyperbox containing its numerical vector representation is split to correct the classification (after the concrete state is wrapped around with a hyperbox of minimal size d , if $d > 0$).

The incremental branch-and-bound refinement of the abstraction could lead to excessive fragmentation of the abstract state space, making it intractably large, particularly for the purpose of counterexample generation. To mitigate this issue, we merge adjacent states, i.e., abstract states sharing at least one separating hyperplane, into a single abstract state, adapting the general notion of probabilistic approximate ϵ -simulation [12] as in the following definition:

Definition 9. *Adjacent ϵ -simulation.* Let S_l be the partitions of S_a induced by the equivalence relation $s \sim s'$ iff $L_a(s) = L_a(s')$. Then, for a given $\epsilon \in [0, 1]$, two adjacent states $s \in S_a$ and $s' \in S_a$ are ϵ -similar if $(\exists s_l \in S_l)(s \in s_l, s' \in s_l)$ and $(\forall s_l \in S_l)(|P_a(s, a, s_l) - P_a(s', a, s_l)| \leq \epsilon, \forall a \in A_a)$.

The ϵ -simulation in def. 9 induces a hierarchical merging scheme. Let level l_0 contain the initial abstract states, which partition the explored concrete state space into a finite set of boxes – one per abstract state – which are labeled as

either safe or unsafe. ϵ -similar adjacent states from level l_i are merged into a single state at level l_{i+1} , until no further merge is possible. Besides reducing the number of abstract states, and in turn the cost of generating counterexamples, this hierarchical merging scheme brings the indirect benefit of more aggressively merging abstract states corresponding to safe regions of the concrete state space, while preserving a finer-grained approximation of concrete state space regions in proximity of explored unsafe states, as discussed in [26].

Counterexample-guided Simulation. If a probabilistic safety requirement can be violated, one or more counterexamples M_{cex} can be generated from the abstract model M_a , where each counterexample includes a (near-)minimal subset of the abstract state-action space. We then use each counterexample as a guide to build an offline, simulation environment where the agent can update its Q-values towards avoiding eventually reaching an unsafe state.

Starting from the initial concrete state s_0 , the abstract state s_{cex} in M_{cex} corresponding to the current concrete state is computed. By construction, each counterexample selects one action a from state s_{cex} . The abstract transition is randomly simulated according to the transition function P_{cex} from (s_{cex}, a) and an abstract destination state s'_{cex} is identified. Such abstract state is concretised by sampling from the past concrete traces a transition (s, a, s') where $s \in s_{cex}$ and $s' \in s'_{cex}$. If s'_{cex} is an unsafe state, a penalty (negative reward; the impact of its magnitude is further discussed in [26]) is given to the agent, which has the transitive effect of re-weighting also the Q-value of the actions that led the agent to the current state. The simulation traces can be used by both Q-learning and DQN agents. The simulation terminates when an unsafe state is reached (penalty) or when we fail to concretise an abstract transition, which may happen when concrete safe states are misclassified as unsafe in the abstraction, but there is no actual transition to unsafe states from them. In the latter case, the simulation trace is discarded (no reward). While every simulation within the counterexample is designed to eventually reach an unsafe state with probability 1 by construction [49], to avoid excessive length of a simulation, it can be practical to set an arbitrary, large bound on the maximum number of steps per run as additional termination criterion.

Multiple counterexamples, and corresponding simulations, can be generated up to a maximum simulation budget allowed by the user, adding blocking clauses in random order as described previously. Each counterexample is typically of small-size, which results in short simulation traces, thus reducing the overall cost of each offline learning experience.

3.2 Online-Offline Learning with Counterexample Guidance

Algorithm 1 summarises the main steps of our online-offline learning method with counterexample guidance. We initially assume no knowledge about the environment is given to the agent: both the abstract model M_a and the set of explored paths D are empty (line 3). If prior knowledge was available, either in the form of an initial abstraction or of previously explored paths, M_a and D can be initialised accordingly.

Online learning. The `onlineQLearning` procedure (line 7) lets the agent operate in the concrete environment with either tabular Q-Learning in discrete state space or DQN in continuous state space. We augment the exploration with a sequential Bayesian hypothesis testing (line 6) that monitors the frequency of violations of the safety requirement, by incrementally updating after each online episode a Beta distribution that estimates the probability of violation [13]. If the odds of such probability exceeding λ is larger than a prescribed Bayes factor β , the online learning phase is interrupted. The updated Q-values/Q-network of the agent are stored and the set of explored traces D is updated (line 8).

Offline learning. If the online learning phases has been interrupted because by the Bayesian test (line 9), an offline learning phase is triggered to reinforce the avoidance of discovered unsafe behaviors in future online exploration.

First, the abstraction M_a is updated with the current set of online traces D (line 10) to support the generation of current counterexamples. While there are theoretically a finite number of counterexample submodels [49], for large M_a it could be computationally too expensive; instead, up to a maximum number N_{cex} of counterexamples is generated at each iteration. The addition of random blocking clauses (as described in sec. 2.2) will increase the diversity of the counterexamples within and across different offline learning phases.

The offline simulation traces synthesised from each M_{cex} (line 13) as described in the previous section are used by the agent to update its Q-values/Q-network (line 14), thus penalising the selection of eventually unsafe actions before the next online learning phase begins. Notice that the Bayesian hypothesis test is re-initialised before the next online learning phase (line 6) since the agent is expected to behave differently after an offline learning phase.

Discussion. The interleaving of offline and online learning phases aims at reducing the frequency of unsafe events during the exploration of the environment. This goal is pursued by synthesising simulation environments from counterexamples to the safety requirement computed from an abstraction of the state space explored online. Notably, the offline phases never preclude the exploration of an action during the following online phases, rather they reduce the likelihood of selecting actions that may eventually lead to reaching unsafe states by lowering their Q-values. Due to space limitations, we report here the two main results related to the convergence of our abstraction method and of the online-offline learning process, and defer a more extensive discussion to [26].

Counterexample guidance relies on the abstraction constructed from online exploration phases, which classifies every region of the explored state space as either safe or unsafe. While by construction the abstraction preserves the safety invariant (every explored concrete unsafe state is mapped to an abstract unsafe state), the quality of the offline guidance relies also on controlling the misclassification error of safe concrete regions, which may unduly penalise the exploration of safe states.

Proposition 1. *The maximum misclassification error of a concrete safe state into an abstract unsafe state can eventually be reduced below an arbitrary bound $0 < \bar{u} \leq 1$ with probability at least $1 - \delta$ ($0 < \delta < 1$) throughout the exploration.*

Further empirical analysis of the convergence of the abstraction and the impact of the abstraction parameters is provided in [26].

Finally, the following proposition states that the introduction of counter-example guidance does not preclude the convergence of the overall learning process to a maximal reward policy that satisfies the safety requirement, if such policy exists.

Proposition 2. *If there exist maximal-reward policies that satisfy the safety requirement, then the online-offline learning process eventually converges to one of them.*

Further discussion of the convergence properties of the offline-online learning process are included in [26], including elaborating on the validity of the two propositions above. In the next section, we instead report on our preliminary experimental evaluation of the performance of our counterexample-guided learning process.

Algorithm 1: CEX-Guided Reinforcement Learning

```

1 Input: Safety requirement bound  $\lambda > 0$ , max number of CEX  $N_{cex}$ , num
  simulations per CEX  $NS_{cex}$ , Bayes factor  $\beta$ , penalty for offline learning
  traces penalty
2 Output: Synthesised Policy  $\pi$ 
3 Initialise  $M_a, D \leftarrow \emptyset$ ;
4 Initialise  $Q$  to random values;
5 while not converged do
6    $\delta_{bht} \leftarrow \text{BayesianHypothesisTester}(\lambda, \beta)$ 
7    $Q, D_{online} \leftarrow \text{onlineQLearning}(Q, \delta_{bht})$ ;
8    $D \leftarrow D \cup D_{online}$ ;
9   if  $\delta_{bht} > \beta$  then
10    //cex-guided offline learning
11     $M_a \leftarrow \text{refine}(M_a, D)$ ; // update the abstract MDP
12    for up to  $N_{cex}$  do
13       $M_{cex} \leftarrow \text{cexGeneration}(M_a, \lambda)$ 
14       $\text{simTraces} \leftarrow \text{collectSimulationTraces}(M_{cex}, Q, NS_{cex})$ ;
15       $Q \leftarrow \text{Q-value/Q-network update}(\text{simTraces}, \text{penalty})$ ;
16    end
17  end
18  $\pi \leftarrow \arg \max_a Q$ ;

```

4 Evaluation

In this section, we present a preliminary experimental evaluation of the performance of our method from two perspectives: 1) the improvement in the exploration safety rate, and 2) the impact on the cumulative reward achieved by the

agent. Finally, we briefly discuss the overhead of counterexample guidance and make some observations on the policies it synthesises. Additional experimental results and discussion, including on abstraction effectiveness and sensitivity to hyperparameters can be found in [26].

Environments. We consider four environments: **DiscreteGrid** from the implementation of [23], the slippery **FrozenLake8x8** from OpenAI Gym [6], **HybridGrid** – where we change the state space from discrete to continuous with the same layout in [23], and **MarsRover** [22] (in particular, the exploration of the melas chasma in the Coprates quadrangle [34]). In all the environments, the agent decides a direction of move between **up**, **down**, **left**, and **right**. We define the objective of the agent as finding a policy with maximum Q-value while avoiding unsafe behaviours with intolerably high probability λ during exploration. Specifically, in the **FrozenLake8x8**, the agent aims to find a walkable path in a 8x8 grid environment with slippery actions while avoiding entering states labelled with *H*. With the slippery setting, the agent will move in the intended direction with a probability of only 1/3 else will move in either perpendicular directions with equal probabilities of 1/3 respectively. In the **DiscreteGrid** and **HybridGrid**, the agent aims to reach the states labelled with **goal1** and then states labelled with **goal2** in a fixed order while avoiding entering any states labelled with **unsafe** along the path, with 15% probability of moving to random directions at every action. In **HybridGrid**, the distance covered in a move is also randomly sampled from a Gaussian $\mathcal{N}(2, 0.5)$, thus choosing the same direction from a state may reach different states. In the **MarsRover** environment, the agent aims to find one of the target states labelled with **goal** while avoiding reaching the unsafe regions, which in this scenario cannot be perfectly abstracted by boxes or other affine abstract domains. Following [21], the distance covered in each move is sampled uniformly in the range $(0, 10)$, with the addition of further uniform noise from $\mathcal{U}(-0.1, 0.5)$.

Baselines. We compare the learning performance of our method with classical Q-Learning and DQN [35] as the baseline for discrete and continuous MDPs, respectively. For discrete MDPs, we further compare our method with [20] (referred to as QL-LCRL in the following), using the same set of hyper-parameters as provided in their implementation and the associated tool paper [21,23]. Given an automata corresponding to an LTL property, QL-LCRL guides the agent’s exploration of an initially unknown MDP by reshaping on-the-fly the reward function to encourage the exploration of behaviors that satisfy such property. In this application QL-LCRL will encourage a safe exploration by discouraging reaching **unsafe** states.

Implementation and parameters. We implemented a standard tabular Q-learning in Python and used the DQN implementation from OpenAI Gym Baselines [37]. We parameterise the abstraction process with a learning rate α and a discount factor γ for the Q-value/Q-Network updates, and ϵ for adjacent ϵ -bisimulation. The agents move within Cartesian planes of sizes $|S|$ and the minimisation of the abstract models reduces the state space to $|S_a|$, where the minimum size of each box is set to 1 and 0.01 for the discrete and the continuous

environments, respectively. We set the safety specification parameter λ according to the intrinsic uncertainty in the respective environments. A summary of the parameters used for each environment is reported in the left side of tab. 1 (additional parameters are discussed in [26] to ease reproducibility). We require at least 50 samples to be collected by the Bayesian hypothesis test before it can trigger offline training to reduce false positive triggers.

Environment	$ S $	$ S_A $	λ	α	γ	ϵ	fpr	QL/DQN safe. prob	QL/DQN reward	QL-LCRL safe. prob	QL-LCRL reward	Q-CEX safe. prob	Q-CEX reward
DiscreteGrid	3200	30	0.2	0.9	0.9	0.01	0.05	0.299	0.924	0.387	0.99	0.572	0.99
FrozenLake8x8	64	50	0.35	0.1	0.9	0.01	0.05	0.339	0.55	0.401	0.54	0.443	0.59
HybridGrid	∞	41	0.2	5e-4	0.9	0.01	0.05	0.548	0.753	-	-	0.687	0.78
MarsRover	∞	149	0.2	5e-4	0.9	0.05	0.15	0.681	0.947	-	-	0.775	0.86

Table 1. Learning results with corresponding hyperparamters in CEX-guided Learning

Experimental Results. Fig. 1 shows the accumulated safety rates (bottom) and the rolling average of accumulated rewards, indicating the real-time learning performance of the agent. The line is the average across 10 runs, while the shaded region around it is the standard deviation. We do not provide any prior information to the agent. The dashed vertical lines in the figure indicate the average episode number where an offline learning phase in QL/DQN-CEX (our method) is triggered, and the solid horizontal line indicates the target safety rate in corresponding safety specifications. The cumulative rewards of different methods converge to similar values, demonstrating that the performance under guidance (QL-LCRL and Q-CEX) achieves cumulative rewards comparable to the baseline QL and DQN method. As expected, providing additional guidance to discourage actions that may lead to reaching unsafe states with QL/DQN-CEX or QL-LCRL improves the safety rate of the exploration, with QL/DQN-CEX achieving on average higher safety rates.

In **DiscreteGrid**, the safety rate of online learning in our method exceeds the threshold much faster than other methods. This is due to the rapid convergence of the abstraction thanks to the grid layout that can be accurately and efficiently abstracted (and minimised). In turn, no further offline phases were required after the first 1000 episodes.

In **FrozenLake8x8**, more online exploration is required to support comprehensive counterexample guidance, partly due to the high uncertainty in the outcome of the actions. Another phenomenon due to high uncertainty is that the agent takes a longer time to reach a stable performance, therefore the offline learning phase is triggered more frequently compared with other environments and also more episodes were required to stably satisfy the safety requirement.

In **HybridGrid** and **MarsRover**, while the baseline eventually achieves a marginally higher cumulative reward in MarsRover, DQN-CEX achieves a higher safety rate – the number of failures experienced by the agent is inversely proportional to the integral of the safety rate curve, which results in significantly




Fig. 1. Average over 10 runs of the accumulated safety rates and rolling rewards in QL/DQN, QL-LCRL [23] and QL/DQN-CEX (our method). The shaded region shows $\pm$ standard deviation.

fewer failure events. The frequency of offline training phases also decreases over time. This is not surprising due to a safer exploration being possibly less speculative and slower in exploring the optimal policy. QL-LCRL is not applicable to HybridGrid and MarsRover. Although LCRL [23] applies NFQ-based method for continuous MDPs, NFQ-LCRL trains the agent in a completely offline man-

ner based on randomly sampled data, thus it is not suitable for comparison with our method from the perspective of safe exploration.

Overhead. Counterexamples generation requires solving a MILP problem on the abstract state space. On a Macbook Air with M1 CPU and 8Gb of memory, the average $\pm$ stdev time to solve the optimization problems were 0.71 ± 0.98 s, 0.08 ± 0.06 s, 2.31 ± 3.94 s, and 3.87 ± 4.07 s for **FrozenLake8x8**, **DiscreteGrid**, **HybridGrid**, and **MarsRover**, respectively. We used Gurobi Optimiser v9.1.0 [17] off-the-shelf. The numbers of counterexamples generated for each offline learning phase were, on average: 11, 6, 19, and 21, respectively. Notice that both counterexample generation and the simulations can be parallelised. While the cost of solving the MILP may be higher, we notice that offline learning is triggered only when the recent online exploration resulted in an unacceptable failure rate. As shown in fig. 1, thanks to counterexample-guidance, QL/DQN-CEX can achieve the required safety exploration rate much faster, which helps amortizing the initial MILP solution cost. Finally, the optimisation problem might be relaxed to sacrifice the optimality of the solution (i.e., size of the counterexamples) for computation time.

5 Related Work

Safe Exploration. [16,27,7] provide surveys and taxonomies of recent safe reinforcement learning methods with different emphases. Most existing safe exploration methods assume prior knowledge about the environment or the agent’s dynamics [36,11,14], known safety constraints [33,25,20,38], or utilise expert guidance [15,52,24,39] to provide guarantees of safety constraints satisfaction during exploration. A different class of methods [44,47,32] utilises surrogate Gaussian process models to characterise the unknown dynamics and optimise the unknown function with a *prior* model structure. There are only a few methods [22,5] tackling safe exploration without known model structure. [22] focuses on solving continuous, unknown MDPs into sub-tasks using an online RL framework under LTL specification with less emphasis on safety rate but blocking actions unsafe according to the specification, and [5] trains a safety layer/critic used for filtering out probabilistic unsafe actions with an offline dataset. While motivated by the same idea of safer exploration without prior knowledge, our method can be initialised with none or any amount of previously explored paths, meanwhile it converged to cumulative rewards comparably or better than baseline methods.

Offline RL. Offline RL can be seen as a data-driven formulation of RL, where the agent collects transitions using the behaviour policy instead of interacting with the environment [31]. The biggest challenge in offline RL is the bootstrapping error: the Q-value is evaluated with little or no prior knowledge and propagated through Bellman equation [28]. [50,41] regularise the behaviour policy while optimising to address this issue. [9,29] alternatively update the Q values in more conservative ways to learn a lower bound of the Q function using uncertainty estimated with the sampled information. From the safe RL perspective,

[46] optimises a risk-averse criterion using data previously collected by a *safe* policy offline, and [51] learns a constrained policy maximizing the long-term reward based on offline data without interaction in the concrete environment using a constrained penalised Q-Learning method. These methods have similar motivation of combining offline learning with risk-averse RL as ours, while focusing on continuous control setting instead. Besides utilising offline learning to reduce unsafe online exploration, we alternate online and offline learning to keep the abstract knowledge about the environment up to date and increase the risk aversion of the agent based on the most current evidence it collected online.

6 Conclusion

We presented our investigation of a safer model-free reinforcement learning method using counterexample-guided offline training. We proposed an abstraction strategy to represent the knowledge acquired during online exploration in a succinct, finite MDP model that can consistently and accurately describe safety-relevant dynamics of the explored environment. Counterexample generation methods from probabilistic model checking are then adapted to synthesise small-scale simulation environments capturing scenarios in which the decisions of the agent may lead to the violation of safety requirements. The agent can then train offline within this minimal submodels by replaying concrete transitions recorded during past online exploration consistent with the counterexample, using a reward scheme focused on reducing the likelihood of selecting actions that may eventually lead to visiting again explored unsafe concrete states. The q-values penalized during the offline training phases implicitly reduce the risk of repeating unsafe behaviors during subsequent online exploration, while newly explored paths feedback information to the next offline learning phase.

The alternation of online exploration – and abstraction refinement – and counterexample guided learning can ultimately lead to higher safety rates during exploration, without significant reduction in the achieved cumulative reward, as demonstrated in our preliminary evaluation on problems from previous literature and the OpenAI Gym. While this paper focused on improving Q-Learning (and the related DQN algorithm), the fundamental framework is not specific to Q-Learning, and we plan to explore its impact on other learning algorithms in future work.

Data availability. A prototype Python implementation of our method is available at Github: <https://github.com/xtji/CEX-guided-RL>.

References

1. Achiam, J., Held, D., Tamar, A., Abbeel, P.: Constrained policy optimization. In: International Conference on Machine Learning. pp. 22–31. PMLR (2017)
2. Alshiekh, M., Bloem, R., Ehlers, R., Könighofer, B., Niekum, S., Topcu, U.: Safe reinforcement learning via shielding. In: Thirty-Second AAAI Conference on Artificial Intelligence (2018)
3. Baier, C., Katoen, J.P.: Principles of model checking. MIT press (2008)
4. Bellman, R.: A markovian decision process. *Journal of mathematics and mechanics* pp. 679–684 (1957)
5. Bharadhwaj, H., Kumar, A., Rhinehart, N., Levine, S., Shkurti, F., Garg, A.: Conservative safety critics for exploration. *arXiv preprint arXiv:2010.14497* (2020)
6. Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., Zaremba, W.: Openai gym. *arXiv preprint arXiv:1606.01540* (2016)
7. Brunke, L., Greeff, M., Hall, A.W., Yuan, Z., Zhou, S., Panerati, J., Schoellig, A.P.: Safe learning in robotics: From learning-based control to safe reinforcement learning. *Annual Review of Control, Robotics, and Autonomous Systems* **5**, 411–444 (2022)
8. Bshouty, N.H., Goldman, S.A., Mathias, H.D., Suri, S., Tamaki, H.: Noise-tolerant distribution-free learning of general geometric concepts. *Journal of the ACM (JACM)* **45**(5), 863–890 (1998)
9. Buckman, J., Gelada, C., Bellemare, M.G.: The importance of pessimism in fixed-dataset policy optimization. *arXiv preprint arXiv:2009.06799* (2020)
10. Česka, M., Hensel, C., Junges, S., Katoen, J.P.: Counterexample-driven synthesis for probabilistic program sketches. In: International Symposium on Formal Methods. pp. 101–120. Springer (2019)
11. Dalal, G., Dvijotham, K., Vecerik, M., Hester, T., Paduraru, C., Tassa, Y.: Safe exploration in continuous action spaces. *arXiv preprint arXiv:1801.08757* (2018)
12. Desharnais, J., Laviolette, F., Tracol, M.: Approximate analysis of probabilistic processes: Logic, simulation and games. In: 2008 Fifth International Conference on Quantitative Evaluation of Systems. pp. 264–273. IEEE (2008)
13. Downey, A.: Think Bayes. O’Reilly Media (2021)
14. Fulton, N., Platzer, A.: Safe reinforcement learning via formal methods: Toward safe control through proof and learning. *Proceedings of the AAAI Conference on Artificial Intelligence* **32**(1) (Apr 2018)
15. Garcia, J., Fernández, F.: Safe exploration of state and action spaces in reinforcement learning. *Journal of Artificial Intelligence Research* **45**, 515–564 (2012)
16. Garcia, J., Fernández, F.: A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research* **16**(1), 1437–1480 (2015)
17. Gurobi Optimization, LLC: Gurobi Optimizer Reference Manual (2022), <https://www.gurobi.com>
18. Han, T., Katoen, J.P., Berteun, D.: Counterexample generation in probabilistic model checking. *IEEE Transactions on Software Engineering* **35**(2), 241–257 (2009)
19. Hansson, H., Jonsson, B.: A logic for reasoning about time and reliability. *Formal aspects of computing* **6**(5), 512–535 (1994)
20. Hasanbeig, M., Abate, A., Kroening, D.: Logically-constrained reinforcement learning. *arXiv preprint arXiv:1801.08099* (2018)
21. Hasanbeig, M., Kroening, D., Abate, A.: LCRL: Certified policy synthesis via logically-constrained reinforcement learning - implementation, <https://github.com/grockious/lcrl>

22. Hasanbeig, M., Kroening, D., Abate, A.: Deep reinforcement learning with temporal logics. In: *Formal Modeling and Analysis of Timed Systems*. pp. 1–22. Springer (2020)
23. Hasanbeig, M., Kroening, D., Abate, A.: LCRL: certified policy synthesis via logically-constrained reinforcement learning. In: *Quantitative Evaluation of Systems - 19th International Conference, QEST. LNCS*, vol. 13479, pp. 217–231. Springer (2022)
24. Huang, J., Wu, F., Precup, D., Cai, Y.: Learning safe policies with expert guidance. In: *Proceedings of the 32nd International Conference on Neural Information Processing Systems*. pp. 9123–9132. NIPS’18, Curran Associates Inc. (2018)
25. Jansen, N., Könighofer, B., Junges, S., Bloem, R.: Shielded decision-making in mdps. *arXiv preprint arXiv:1807.06096* (2018)
26. Ji, X., Filieri, A.: Probabilistic counterexample guidance for safer reinforcement learning (extended version). *arXiv preprint arXiv:2307.04927* (2023)
27. Kim, Y., Allmendinger, R., López-Ibáñez, M.: Safe learning and optimization techniques: Towards a survey of the state of the art. In: *International Workshop on the Foundations of Trustworthy AI Integrating Learning, Optimization and Reasoning*. pp. 123–139. Springer (2020)
28. Kumar, A., Fu, J., Tucker, G., Levine, S.: Stabilizing Off-Policy Q-Learning via Bootstrapping Error Reduction. Curran Associates Inc. (2019)
29. Kumar, A., Zhou, A., Tucker, G., Levine, S.: Conservative q-learning for offline reinforcement learning. *Advances in Neural Information Processing Systems* **33**, 1179–1191 (2020)
30. Lawler, E.L., Wood, D.E.: Branch-and-bound methods: A survey. *Operations research* **14**(4), 699–719 (1966)
31. Levine, S., Kumar, A., Tucker, G., Fu, J.: Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643* (2020)
32. Liu, A., Shi, G., Chung, S.J., Anandkumar, A., Yue, Y.: Robust regression for safe exploration in control. In: *Learning for Dynamics and Control*. pp. 608–619. PMLR (2020)
33. Mason, G.R., Calinescu, R.C., Kudenko, D., Banks, A.: Assured reinforcement learning with formally verified abstract policies. In: *9th International Conference on Agents and Artificial Intelligence (ICAART)*. York (2017)
34. McEwen, A.S., Dundas, C.M., Mattson, S.S., Toigo, A.D., Ojha, L., Wray, J.J., Chojnacki, M., Byrne, S., Murchie, S.L., Thomas, N.: Recurring slope lineae in equatorial regions of Mars. *Nature Geoscience* **7**(1), 53–58 (2014)
35. Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., Riedmiller, M.: Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602* (2013)
36. Moldovan, T.M., Abbeel, P.: Safe exploration in markov decision processes. *arXiv preprint arXiv:1205.4810* (2012)
37. OpenAI: Stable baselines version 3 - dqn, <https://stable-baselines3.readthedocs.io/en/master/modules/dqn.html>
38. Pham, T.H., De Magistris, G., Tachibana, R.: Optlayer-practical constrained optimization for deep reinforcement learning in the real world. In: *2018 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 6236–6243. IEEE (2018)
39. Prakash, B., Khatwani, M., Waytowich, N., Mohsenin, T.: Improving safety in reinforcement learning using model-based architectures and human intervention. *arXiv preprint arXiv:1903.09328* (2019)

40. Sharma, R., Gupta, S., Hariharan, B., Aiken, A., Nori, A.V.: Verification as learning geometric concepts. In: International Static Analysis Symposium. pp. 388–411. Springer (2013)
41. Siegel, N.Y., Springenberg, J.T., Berkenkamp, F., Abdolmaleki, A., Neunert, M., Lampe, T., Hafner, R., Heess, N., Riedmiller, M.: Keep doing what worked: Behavioral modelling priors for offline reinforcement learning. arXiv preprint arXiv:2002.08396 (2020)
42. Singh, G., Püschel, M., Vechev, M.: A practical construction for decomposing numerical abstract domains. *Proc. ACM Program. Lang.* **2**(POPL) (dec 2017)
43. Stooke, A., Achiam, J., Abbeel, P.: Responsive safety in reinforcement learning by pid lagrangian methods. In: International Conference on Machine Learning. pp. 9133–9143. PMLR (2020)
44. Sui, Y., Gotovos, A., Burdick, J., Krause, A.: Safe exploration for optimization with gaussian processes. In: International conference on machine learning. pp. 997–1005. PMLR (2015)
45. Tessler, C., Mankowitz, D.J., Mannor, S.: Reward constrained policy optimization. arXiv preprint arXiv:1805.11074 (2018)
46. Urpí, N.A., Curi, S., Krause, A.: Risk-averse offline reinforcement learning. arXiv preprint arXiv:2102.05371 (2021)
47. Wachi, A., Sui, Y., Yue, Y., Ono, M.: Safe exploration and optimization of constrained mdps using gaussian processes. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 32 (2018)
48. Watkins, C.J., Dayan, P.: Q-learning. *Machine learning* **8**(3), 279–292 (1992)
49. Wimmer, R., Jansen, N., Vorpahl, A., Ábrahám, E., Katoen, J.P., Becker, B.: High-level counterexamples for probabilistic automata. In: International Conference on Quantitative Evaluation of Systems. pp. 39–54. Springer (2013)
50. Wu, Y., Tucker, G., Nachum, O.: Behavior regularized offline reinforcement learning. arXiv preprint arXiv:1911.11361 (2019)
51. Xu, H., Zhan, X., Zhu, X.: Constraints penalized q-learning for safe offline reinforcement learning. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 36, pp. 8753–8760 (2022)
52. Zhou, W., Li, W.: Safety-aware apprenticeship learning. In: International Conference on Computer Aided Verification. pp. 662–680. Springer (2018)