
A Lyapunov-based Approach to Safe Reinforcement Learning

Yinlam Chow
DeepMind
yinlamchow@google.com

Ofir Nachum
Google Brain
ofirnachum@google.com

Edgar Duenez-Guzman
DeepMind
duenez@google.com

Mohammad Ghavamzadeh
Facebook AI Research
mgh@fb.com

Abstract

In many real-world reinforcement learning (RL) problems, besides optimizing the main objective function, an agent must concurrently avoid violating a number of constraints. In particular, besides optimizing performance, it is crucial to guarantee the *safety* of an agent during training as well as deployment (e.g., a robot should avoid taking actions - exploratory or not - which irrevocably harm its hardware). To incorporate safety in RL, we derive algorithms under the framework of constrained Markov decision processes (CMDPs), an extension of the standard Markov decision processes (MDPs) augmented with constraints on expected cumulative costs. Our approach hinges on a novel *Lyapunov* method. We define and present a method for constructing Lyapunov functions, which provide an effective way to guarantee the global safety of a behavior policy during training via a set of local linear constraints. Leveraging these theoretical underpinnings, we show how to use the Lyapunov approach to systematically transform dynamic programming (DP) and RL algorithms into their safe counterparts. To illustrate their effectiveness, we evaluate these algorithms in several CMDP planning and decision-making tasks on a safety benchmark domain. Our results show that our proposed method significantly outperforms existing baselines in balancing constraint satisfaction and performance.

1 Introduction

Reinforcement learning (RL) has shown exceptional successes in a variety of domains such as video games [25] and recommender systems [40], where the main goal is to optimize a single return. However, in many real-world problems, besides optimizing the main objective (the return), there can exist several conflicting constraints that make RL challenging. In particular, besides optimizing performance it is crucial to guarantee the *safety* of an agent in deployment [5, 32, 33], as well as during training [2]. For example, a robot should avoid taking actions which irrevocably harm its hardware; a recommender system must avoid presenting harmful or offending items to users.

Sequential decision-making in non-deterministic environments has been extensively studied in the literature under the framework of Markov decision processes (MDPs). To incorporate safety into the RL process, we are particularly interested in deriving algorithms under the context of constrained Markov decision processes (CMDPs), which is an extension of MDPs with expected cumulative constraint costs. The additional constraint component of CMDPs increases flexibility in modeling problems with *trajectory-based* constraints, when compared with other approaches that customize immediate costs in MDPs to handle constraints [34]. As shown in numerous applications from robot motion planning [30, 26, 11], resource allocation [24, 18], and financial engineering [1, 41], it is more natural to define safety over the whole trajectory, instead of over particular state and action pairs. Under this framework, we denote an agent's behavior policy to be *safe* if it satisfies the cumulative cost constraints of the CMDP.

Despite the capabilities of CMDPs, they have not been very popular in RL. One main reason is that, although optimal policies of finite CMDPs are Markov and stationary, and with known models the CMDP can be solved using linear programming (LP) [3], it is unclear how to extend this algorithm to handle cases when the model is unknown, or when the state and action spaces are large or continuous. A well-known approach to solve CMDPs is the Lagrangian method [4, 15], which augments the standard expected reward objective with a penalty on constraint violation. With a fixed Lagrange multiplier, one can use standard dynamic programming (DP) or RL algorithms to solve for an optimal policy. With a learnable Lagrange multiplier, one must solve the resulting saddle point problem. However, several studies [21] showed that iteratively solving the saddle point is apt to run into numerical stability issues. More importantly, the Lagrangian policy is only safe *asymptotically* and makes little guarantee with regards to safety of the behavior policy during each training iteration.

Motivated by these observations, several recent works have derived surrogate algorithms for solving CMDPs, which transform the original constraint to a more conservative one that yields an easier problem to solve. A straight-forward approach is to replace the cumulative constraint cost with a conservative *stepwise* surrogate constraint [9] that only depends on the current state-action pair. Since this surrogate constraint can be easily embedded into the admissible control set, this formulation can be modeled by an MDP that has a restricted set of admissible actions. Another surrogate algorithm was proposed by [14] in which the algorithm first computes a uniform *super-martingale* constraint value function surrogate w.r.t. all policies, and then finds a CMDP feasible policy by optimizing the surrogate problem using the *lexicographical ordering* method [39]. These methods are advantageous in the sense that (i) there are RL algorithms available to handle the surrogate problems (for example see [12] for the step-wise surrogate and [27] for the super-martingale surrogate), (ii) the policy returned by this method is safe, even during training. However, the main drawback of these approaches is their conservativeness. Characterizing sub-optimality performance of the corresponding solution policy also remains a challenging task. On the other hand, recently in policy gradient, [2] proposed the constrained policy optimization (CPO) method that extends trust-region policy optimization (TRPO) to handle the CMDP constraints. While this algorithm is scalable and its policy is safe during training, applying this methodology to more general RL algorithms (that are not in the family of proximal PG algorithms) is quite non-trivial.

Lyapunov functions have been extensively used in control theory to analyze the stability of dynamic systems [20, 28]. A Lyapunov function is a type of scalar potential function that keeps track of the *energy* that a system continually dissipates. Besides modeling physical energy, Lyapunov functions can also represent abstract quantities, such as the steady-state performance of a Markov process [16]. In many fields, Lyapunov functions provide a powerful paradigm to translate global properties of a system to local ones and vice-versa. Using Lyapunov functions in RL was first studied by [31], where Lyapunov functions were used to guarantee closed-loop stability of an agent. Recently [6] used Lyapunov functions to guarantee a model-based RL agent’s ability to re-enter an “attraction region” during exploration. However, no previous works have used Lyapunov approaches to explicitly model constraints in a CMDP. Furthermore, one major drawback of these approaches is that the Lyapunov functions are hand-crafted, and there are no principled guidelines on designing Lyapunov functions that can guarantee the agent’s performance.

The contribution of this paper is four-fold. First, we formulate the problem of safe RL as a CMDP and propose a novel *Lyapunov approach* to solve it. While the main challenge of other Lyapunov-based methods is to design a Lyapunov function candidate, we propose an LP-based algorithm to construct Lyapunov functions w.r.t. generic CMDP constraints. We also show that our method is guaranteed to always return a feasible policy, and under certain technical assumptions, it achieves optimality. Second, leveraging the theoretical underpinnings of the Lyapunov approach, we present two safe DP algorithms – safe policy iteration (SPI) and safe value iteration (SVI) – and analyze the feasibility and performance of these algorithms. Third, to handle unknown environments and large state/action spaces, we develop two scalable safe RL algorithms – (i) *safe DQN*, an off-policy fitted Q -iteration method, and (ii) *safe DPI*, an approximate policy iteration method. Fourth, to illustrate the effectiveness of these algorithms, we evaluate them in several tasks on a benchmark 2D planning problem and show that they outperform common baselines in terms of balancing performance and constraint satisfaction.

2 Preliminaries

We consider RL problems in which the agent’s interaction with the system is modeled as a Markov decision process (MDP). A MDP is a tuple $(\mathcal{X}, \mathcal{A}, c, P, x_0)$, where $\mathcal{X} = \mathcal{X}' \cup \{x_{\text{Term}}\}$ is the state space, with transient state space \mathcal{X}' and terminal state x_{Term} ; \mathcal{A} is the action space;

$c(x, a) \in [0, C_{\max}]$ is the immediate cost function (negative reward); $P(\cdot|x, a)$ is the transition probability distribution; and $x_0 \in \mathcal{X}'$ is the initial state. Our results easily generalize to random initial states and random costs, but for simplicity we will focus on the case of deterministic initial state and immediate cost. In a more general setting where cumulative constraints are taken into account, we define a constrained Markov decision process (CMDP), which extends the MDP model by introducing additional costs and associated constraints. A CMDP is defined by $(\mathcal{X}, \mathcal{A}, c, d, P, x_0, d_0)$, where the components $\mathcal{X}, \mathcal{A}, c, P, x_0$ are the same for the unconstrained MDP; $d(x) \in [0, D_{\max}]$ is the immediate constraint cost; and $d_0 \in \mathbb{R}_{\geq 0}$ is an upper-bound on the expected cumulative (through time) constraint cost. To formalize the optimization problem associated with CMDPs, let Δ be the set of Markov stationary policies, i.e., $\Delta(x) = \{\pi(\cdot|x) : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}^s : \sum_a \pi(a|x) = 1\}$, for any state $x \in \mathcal{X}$. Also let T^* be a random variable corresponding to the first-hitting time of the terminal state x_{Term} induced by policy π . In this paper, we follow the standard notion of transient MDPs and assume that the first-hitting time is uniformly bounded by an upper bound \bar{T} for any stationary policies [10]. This assumption implies that every stationary policy is *proper* [7], whose induced Markov chain has an *absorbing* property (see [13] for an example). While this assumption may seem restrictive, it is a standard one in *stochastic shortest path* problems for showing that the Bellman operator is a contraction. Its justification follows from the fact that sample trajectories collected in most RL algorithms consist of a finite stopping time (also known as a time-out); In general this assumption may also be relaxed in cases where a discount factor $\gamma < 1$ is applied on future costs. For notational convenience, at each state $x \in \mathcal{X}'$, we define the generic Bellman operator w.r.t. policy $\pi \in \Delta$ and generic cost function h : $T_{\pi, h}[V](x) = \sum_a \pi(a|x) \left[h(x, a) + \sum_{x' \in \mathcal{X}'} P(x'|x, a)V(x') \right]$.

Given a policy $\pi \in \Delta$, an initial state x_0 , the cost function is defined as $\mathcal{C}_\pi(x_0) := \mathbb{E} \left[\sum_{t=0}^{T^*-1} c(x_t, a_t) \mid x_0, \pi \right]$, and the safety constraint is defined as $\mathcal{D}_\pi(x_0) \leq d_0$, where the safety constraint function is given by $\mathcal{D}_\pi(x_0) := \mathbb{E} \left[\sum_{t=0}^{T^*-1} d(x_t) \mid x_0, \pi \right]$. In general the CMDP problem we wish to solve is given as follows:

Problem OPT: Given an initial state x_0 and a threshold d_0 , solve $\min_{\pi \in \Delta} \{ \mathcal{C}_\pi(x_0) : \mathcal{D}_\pi(x_0) \leq d_0 \}$. If there is a non-empty solution, the optimal policy is denoted by π^* .

Under the transient CMDP assumption, Theorem 8.1 in [3] shows that if the feasibility set is non-empty, then there exists an optimal policy in the class of stationary Markovian policies Δ . To motivate the CMDP formulation studied in this paper, in Appendix A, we include two real-world examples in modeling safety using (i) the reachability constraint, and (ii) the constraint that limits the agent's visits to undesirable states. Recently there has been a number of works on CMDP algorithms; their details can be found in Appendix B.

3 A Lyapunov Approach to Solve CMDPs

In this section, we develop a novel methodology for solving CMDPs using the *Lyapunov approach*. To start with, without loss of generality we assume to have access to a *baseline* feasible policy of the *OPT* problem, namely $\pi_B \in \Delta$.¹ We define a non-empty² set of Lyapunov functions w.r.t. the initial state $x_0 \in \mathcal{X}$ and constraint threshold d_0 as

$$\mathcal{L}_{\pi_B}(x_0, d_0) = \left\{ L : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0} : T_{\pi_B, d}[L](x) \leq L(x), \forall x \in \mathcal{X}'; L(x) = 0, \forall x \in \mathcal{X} \setminus \mathcal{X}'; L(x_0) \leq d_0 \right\}. \quad (1)$$

For any arbitrary Lyapunov function $L \in \mathcal{L}_{\pi_B}(x_0, d_0)$, we denote by $\mathcal{F}_L(x) = \{ \pi(\cdot|x) \in \Delta : T_{\pi, d}[L](x) \leq L(x) \}$ the set of L -induced Markov stationary policies. Since $T_{\pi, d}$ is a contraction mapping [7], any L -induced policy π has the following property: $\mathcal{D}_\pi(x) = \lim_{k \rightarrow \infty} T_{\pi, d}^k[L](x) \leq L(x), \forall x \in \mathcal{X}'$. Together with the property of $L(x_0) \leq d_0$, this further implies any L -induced policy is a feasible policy of the *OPT* problem. However, in general the set $\mathcal{F}_L(x)$ does not necessarily contain any optimal policies of the *OPT* problem, and our main contribution is to design a Lyapunov function (w.r.t. a baseline policy) that provides this guarantee. In other words, our main goal is to construct a Lyapunov function $L \in \mathcal{L}_{\pi_B}(x_0, d_0)$ such that

$$L(x) \geq T_{\pi^*, d}[L](x), \quad L(x_0) \leq d_0. \quad (2)$$

¹One example of π_B is a policy that minimizes the constraint, i.e., $\pi_B(\cdot|x) \in \arg \min_{\pi \in \Delta(x)} \mathcal{D}_\pi(x)$.

²To see this, the constraint cost function $\mathcal{D}_{\pi_B}(x)$ is a valid Lyapunov function, i.e., $\mathcal{D}_{\pi_B}(x_0) \leq d_0$, $\mathcal{D}_{\pi_B}(x) = 0, \forall x \in \mathcal{X} \setminus \mathcal{X}'$, and $\mathcal{D}_{\pi_B}(x) = T_{\pi_B, d}[\mathcal{D}_{\pi_B}](x) = \mathbb{E} \left[\sum_{t=0}^{T^*-1} d(x_t) \mid \pi_B, x \right], \forall x \in \mathcal{X}'$.

Before getting into the main results, we consider the following important technical lemma, which states that with appropriate *cost-shaping*, one can always transform the constraint value function $\mathcal{D}_{\pi^*}(x)$ w.r.t. optimal policy π^* into a Lyapunov function that is induced by π_B , i.e., $L_\epsilon(x) \in \mathcal{L}_{\pi_B}(x_0, d_0)$. The proof of this lemma can be found in Appendix C.1.

Lemma 1. *There exists an auxiliary constraint cost $\epsilon : \mathcal{X}' \rightarrow \mathbb{R}$ such that a Lyapunov function is given by $L_\epsilon(x) = \mathbb{E} \left[\sum_{t=0}^{T^*-1} d(x_t) + \epsilon(x_t) \mid \pi_B, x \right]$, $\forall x \in \mathcal{X}'$, and $L_\epsilon(x) = 0$, $\forall x \in \mathcal{X} \setminus \mathcal{X}'$. Moreover, L_ϵ is equal to the constraint value function w.r.t. π^* , i.e., $L_\epsilon(x) = \mathcal{D}_{\pi^*}(x)$.*

From the structure of L_ϵ , one can see that the auxiliary constraint cost function ϵ is uniformly bounded by $\epsilon^*(x) := 2\bar{T}D_{\max}D_{TV}(\pi^*||\pi_B)(x)$,³ i.e., $\epsilon(x) \in [-\epsilon^*(x), \epsilon^*(x)]$, for any $x \in \mathcal{X}'$. However, in general it is unclear how to construct such a cost-shaping term ϵ without explicitly knowing π^* a-priori. Rather, inspired by this result, we consider the bound ϵ^* to propose a Lyapunov function *candidate* L_{ϵ^*} . Immediately from its definition, this function has the following properties:

$$L_{\epsilon^*}(x) \geq T_{\pi_B, d}[L_{\epsilon^*}](x), \quad L_{\epsilon^*}(x) \geq \max \{ \mathcal{D}_{\pi^*}(x), \mathcal{D}_{\pi_B}(x) \} \geq 0, \quad \forall x \in \mathcal{X}'. \quad (3)$$

The first property is due to the facts that: (i) ϵ^* is a non-negative cost function; (ii) $T_{\pi_B, d+\epsilon^*}$ is a contraction mapping, which by the fixed point theorem [7] implies $L_{\epsilon^*}(x) = T_{\pi_B, d+\epsilon^*}[L_{\epsilon^*}](x) \geq T_{\pi_B, d}[L_{\epsilon^*}](x)$, $\forall x \in \mathcal{X}'$. For the second property, from the above inequality one concludes that the Lyapunov function L_{ϵ^*} is a uniform upper-bound to the constraint cost, i.e., $L_{\epsilon^*}(x) \geq \mathcal{D}_{\pi_B}(x)$, because the constraint cost $\mathcal{D}_{\pi_B}(x)$ w.r.t. policy π_B is the unique solution to the fixed-point equation $T_{\pi_B, d}[V](x) = V(x)$, $x \in \mathcal{X}'$. On the other hand, by construction, $\epsilon^*(x)$ is an upper-bound of the cost-shaping term $\epsilon(x)$. Therefore, Lemma 1 implies that the Lyapunov function L_{ϵ^*} is a uniform upper-bound to the constraint cost w.r.t. optimal policy π^* , i.e., $L_{\epsilon^*}(x) \geq \mathcal{D}_{\pi^*}(x)$.

To show that L_{ϵ^*} is a Lyapunov function that satisfies (2), we propose the following condition that enforces a baseline policy π_B to be *sufficiently close* to an optimal policy π^* .

Assumption 1. *The feasible baseline policy π_B satisfies the condition $\max_{x \in \mathcal{X}'} \epsilon^*(x) \leq D_{\max} \cdot \min \left\{ \frac{d_0 - \mathcal{D}_{\pi_B}(x_0)}{\bar{T}D_{\max}}, \frac{\bar{T}D_{\max} - \bar{D}}{\bar{T}D_{\max} + \bar{D}} \right\}$, where $\bar{D} = \max_{x \in \mathcal{X}'} \max_{\pi} \mathcal{D}_{\pi}(x)$.*

This condition characterizes the maximum allowable distance between π_B and π^* , such that the set of L_{ϵ^*} -induced policies contains an optimal policy. To formalize this claim, we have the following main result showing that $L_{\epsilon^*} \in \mathcal{L}_{\pi_B}(x_0, d_0)$, and the set of policies $\mathcal{F}_{L_{\epsilon^*}}$ contains an optimal policy.

Theorem 1. *Suppose the baseline policy π_B satisfies Assumption 1, then on top of the properties in (3), the Lyapunov function candidate L_{ϵ^*} also satisfies the properties in (2), and thus, its induced feasible set of policies $\mathcal{F}_{L_{\epsilon^*}}$ contains an optimal policy.*

The proof of this theorem is given in Appendix C.2. Suppose the distance between the baseline and optimal policies can be estimated effectively. Using the above result, one can immediately determine if the set of L_{ϵ^*} -induced policies contain an optimal policy. Equipped with the set of L_{ϵ^*} -induced feasible policies, consider the following *safe* Bellman operator:

$$T[V](x) = \begin{cases} \min_{\pi \in \mathcal{F}_{L_{\epsilon^*}}(x)} T_{\pi, c}[V](x) & \text{if } x \in \mathcal{X}' \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

Using standard analysis of Bellman operators, one can show that T is a monotonic and contraction operator (see Appendix C.3 for the proof). This further implies that the solution of the fixed-point equation $T[V](x) = V(x)$, $\forall x \in \mathcal{X}$ is unique. Let V^* be such a value function. The following theorem shows that under Assumption 1, $V^*(x_0)$ is a solution to the *OPT* problem.

Theorem 2. *Suppose that the baseline policy π_B satisfies Assumption 1. Then, the fixed-point solution at $x = x_0$, i.e., $V^*(x_0)$, is equal to the solution of the *OPT* problem. Furthermore, an optimal policy can be constructed by $\pi^*(\cdot|x) \in \arg \min_{\pi \in \mathcal{F}_{L_{\epsilon^*}}(x)} T_{\pi, c}[V^*](x)$, $\forall x \in \mathcal{X}'$.*

The proof of this theorem can be found in Appendix C.4. This shows that under Assumption 1, an optimal policy of the *OPT* problem can be found using standard DP algorithms. Note that verifying whether π_B satisfies this assumption is still challenging, because one requires a good estimate of $D_{TV}(\pi^*||\pi_B)$. Yet to the best of our knowledge, this is *the first result* that connects the optimality of CMDP to Bellman's principle of optimality. Another key observation is that in

³The definition of total variation distance is given by $D_{TV}(\pi^*||\pi_B)(x) = \frac{1}{2} \sum_{a \in \mathcal{A}} |\pi_B(a|x) - \pi^*(a|x)|$.

practice, we will explore ways of approximating ϵ^* via *bootstrapping* and empirically show that this approach achieves good performance, while guaranteeing safety at each iteration. In particular, in the next section, we will illustrate how to systematically construct a Lyapunov function using an LP in both planning and RL (when the model is unknown and/or we use function approximation) scenarios in order to guarantee safety during learning.

4 Safe Reinforcement Learning Using Lyapunov Functions

Motivated by the challenge of computing a Lyapunov function L_{ϵ^*} such that its induced set of policies contains π^* , in this section, we approximate ϵ^* with an auxiliary constraint cost $\tilde{\epsilon}$, which is the *largest* auxiliary cost that satisfies the Lyapunov condition: $L_{\tilde{\epsilon}}(x) \geq T_{\pi_B, d}[L_{\tilde{\epsilon}}](x)$, $\forall x \in \mathcal{X}'$, and the safety condition $L_{\tilde{\epsilon}}(x_0) \leq d_0$. The larger the $\tilde{\epsilon}$, the larger the set of policies $\mathcal{F}_{L_{\tilde{\epsilon}}}$. Thus, by choosing the largest such auxiliary cost, we hope to have a better chance of including the optimal policy π^* in the set of feasible policies. So, we consider the following LP problem:

$$\tilde{\epsilon} \in \arg \max_{\epsilon: \mathcal{X}' \rightarrow \mathbb{R}_{\geq 0}} \left\{ \sum_{x \in \mathcal{X}'} \epsilon(x) : d_0 - \mathcal{D}_{\pi_B}(x_0) \geq \mathbf{1}(x_0)^\top (I - \{P(x'|x, \pi_B)\}_{x, x' \in \mathcal{X}'})^{-1} \epsilon \right\}. \quad (5)$$

Here $\mathbf{1}(x_0)$ represents a one-hot vector in which the non-zero element is located at $x = x_0$.

On the other hand, whenever π_B is a feasible policy, the problem in (5) always has a non-empty solution.⁴ Furthermore, note that $\mathbf{1}(x_0)^\top (I - \{P(x'|x, \pi_B)\}_{x, x' \in \mathcal{X}'})^{-1} \mathbf{1}(x)$ represents the total visiting probability $\mathbb{E}[\sum_{t=0}^{T^*-1} \mathbf{1}\{x_t = x\} \mid x_0, \pi_B]$ from the initial state x_0 to any state $x \in \mathcal{X}'$, which is a non-negative quantity. Therefore, using the extreme point argument in LP [23], one can simply conclude that the maximizer of problem (5) is an indicator function whose non-zero element locates at state x that corresponds to the minimum total visiting probability from x_0 , i.e., $\tilde{\epsilon}(x) = (d_0 - \mathcal{D}_{\pi_B}(x_0)) \cdot \mathbf{1}\{x = \underline{x}\} / \mathbb{E}[\sum_{t=0}^{T^*-1} \mathbf{1}\{x_t = \underline{x}\} \mid x_0, \pi_B] \geq 0$, $\forall x \in \mathcal{X}'$, where $\underline{x} \in \arg \min_{x \in \mathcal{X}'} \mathbb{E}[\sum_{t=0}^{T^*-1} \mathbf{1}\{x_t = x\} \mid x_0, \pi_B]$. On the other hand, suppose that we further restrict the structure of $\tilde{\epsilon}(x)$ to be a constant function, i.e., $\tilde{\epsilon}(x) = \tilde{\epsilon}$, $\forall x \in \mathcal{X}'$. Then, one can show that the maximizer is given by $\tilde{\epsilon}(x) = (d_0 - \mathcal{D}_{\pi_B}(x_0)) / \mathbb{E}[T^* \mid x_0, \pi_B]$, $\forall x \in \mathcal{X}'$, where $\mathbf{1}(x_0)^\top (I - \{P(x'|x, \pi_B)\}_{x, x' \in \mathcal{X}'})^{-1} [\mathbf{1}, \dots, \mathbf{1}]^\top = \mathbb{E}[T^* \mid x_0, \pi_B]$ is the expected stopping time of the transient MDP. In cases where computing the expected stopping time is expensive, one reasonable approximation is to replace the denominator of $\tilde{\epsilon}$ with the upper-bound \bar{T} .

Using this Lyapunov function $L_{\tilde{\epsilon}}$, we propose the safe policy iteration (SPI) in Algorithm 1, in which the Lyapunov function is updated via *bootstrapping*, i.e., at each iteration $L_{\tilde{\epsilon}}$ is recomputed using (5), w.r.t. the current baseline policy. Properties of SPI are summarized in the following proposition.

Algorithm 1 Safe Policy Iteration (SPI)

Input: Initial feasible policy π_0 ;
for $k = 0, 1, 2, \dots$ **do**
 Step 0: With $\pi_b = \pi_k$, evaluate the Lyapunov function L_{ϵ_k} , where ϵ_k is a solution of (5)
 Step 1: Evaluate the cost value function $V_{\pi_k}(x) = \mathcal{C}_{\pi_k}(x)$
 Step 2: Update the policy by solving the problem $\pi_{k+1}(\cdot|x) \in \arg \min_{\pi \in \mathcal{F}_{L_{\epsilon_k}}(x)} T_{\pi, c}[V_{\pi_k}](x)$, $\forall x \in \mathcal{X}'$
end for
Return Final policy π_{k^*}

Proposition 1. *Algorithm 1 has the following properties: (i) Consistent Feasibility, i.e., suppose that the current policy π_k is feasible, then the updated policy π_{k+1} is also feasible, i.e., $\mathcal{D}_{\pi_k}(x_0) \leq d_0$ implies $\mathcal{D}_{\pi_{k+1}}(x_0) \leq d_0$; (ii) Monotonic Policy Improvement, i.e., the cumulative cost induced by π_{k+1} is lower than or equal to that by π_k , i.e., $\mathcal{C}_{\pi_{k+1}}(x) \leq \mathcal{C}_{\pi_k}(x)$, $\forall x \in \mathcal{X}'$; (iii) Convergence, i.e., if we add a strictly concave regularizer to the optimization problem (5) and a strictly convex regularizer to the policy optimization step, then the policy sequence asymptotically converges.⁵*

The proof of this proposition is given in Appendix C.5, and the sub-optimality performance bound of SPI can be found in Appendix C.6. Analogous to SPI, we also propose a safe value iteration (SVI), in which the Lyapunov function estimate is updated at every iteration via bootstrapping, using the current optimal value estimate. Details of SVI is given in Algorithm 2 and its properties are summarized in the following proposition, whose proof is given in Appendix C.7.

⁴This is due to the fact that $d_0 - \mathcal{D}_{\pi_B}(x_0) \geq 0$, and thus, $\tilde{\epsilon}(x) = 0$ is a feasible solution.

⁵The strict concavity property in the objective function is mainly for the purpose of tie-breaking. One standard example is the entropy regularizer with a small regularization term.

Proposition 2. *Algorithm 2 has: (i) Consistent Feasibility and (ii) Convergence.*

To justify the notion of bootstrapping in both SVI and SPI, the Lyapunov function is updated based on the *best* baseline policy (the policy that is feasible and by far has the lowest cumulative cost). Once the current baseline policy π_k is *sufficiently close* to an optimal policy π^* , then by Theorem 1, one may conclude that the $L_{\tilde{c}}$ -induced set of policies contains an optimal policy. Although these algorithms do not have optimality guarantees, empirically, they often return a near-optimal policy.

At each iteration, the policy optimization step in SPI and SVI requires solving $|\mathcal{X}'|$ LP sub-problems, where each of them has $|\mathcal{A}| + 2$ constraints and has a $|\mathcal{A}|$ -dimensional decision-variable. Collectively, at each iteration its complexity is $O(|\mathcal{X}'||\mathcal{A}|^2(|\mathcal{A}| + 2))$. While in the worst case SVI converges in $K = O(\bar{T})$ steps [7] and SPI converges in $K = O(|\mathcal{X}'||\mathcal{A}|\bar{T} \log \bar{T})$ steps [38], in practice, K is much smaller than $|\mathcal{X}'||\mathcal{A}|$. Therefore, even with the additional complexity of policy evaluation in SPI that is $O(\bar{T}|\mathcal{X}'|^2)$, or the complexity of updating Q -function in SVI that is $O(|\mathcal{A}|^2|\mathcal{X}'|^2)$, the complexity of these methods is $O(K|\mathcal{X}'||\mathcal{A}|^3 + K|\mathcal{X}'|^2|\mathcal{A}|^2)$, which in practice is much lower than that of the dual LP method, whose complexity is $O(|\mathcal{X}'|^3|\mathcal{A}|^3)$ (see Appendix B for more details).

Algorithm 2 Safe Value Iteration (SVI)

Input: Initial Q -function Q_0 ; Initial Lyapunov function L_{ϵ_0} w.r.t. auxiliary cost function $\epsilon_0(x) = 0$;
for $k = 0, 1, 2, \dots$ **do**
 Step 0: Compute Q -function $Q_{k+1}(x, a) = c(x, a) + \sum_{x'} P(x'|x, a) \min_{\pi \in \mathcal{F}_{L_{\epsilon_k}}(x')} \pi(\cdot|x')^\top Q_k(x', \cdot)$
 and policy $\pi_k(\cdot|x) \in \arg \min_{\pi \in \mathcal{F}_{L_{\epsilon_k}}(x)} \pi(\cdot|x)^\top Q_k(x, \cdot)$
 Step 1: With $\pi_B = \pi_k$, construct the Lyapunov function $L_{\epsilon_{k+1}}$, where ϵ_{k+1} is a solution to (5);
end for
Return Final policy π_{k^*}

4.1 Lyapunov-based Safe RL Algorithms

In order to improve scalability of SVI and SPI, we develop two *off-policy* safe RL algorithms, namely safe DQN and safe DPI, which replace the value and policy updates in safe DP with function approximations. Their pseudo-codes can be found in Appendix D. Before going into their details, we first introduce the policy distillation method, which will be later used in the safe RL algorithms.

Policy Distillation: Consider the following LP problem for policy optimization in SVI and SPI:

$$\pi'(\cdot|x) \in \arg \min_{\pi \in \Delta} \left\{ \pi(\cdot|x)^\top Q(x, \cdot) : (\pi(\cdot|x) - \pi_B(\cdot|x))^\top Q_L(x, \cdot) \leq \tilde{c}(x) \right\}, \quad (6)$$

where $Q_L(x, a) = d(x) + \tilde{c}(x) + \sum_{x'} P(x'|x, a) L_{\tilde{c}}(x')$ is the state-action Lyapunov function. When the state-space is large (or continuous), we shall use function approximation. Consider a parameterized policy π_ϕ with weights ϕ . Utilizing the distillation concept [36], after computing the optimal action probabilities w.r.t. a batch of states, the policy π_ϕ is updated by solving $\phi^* \in \arg \min_\phi \frac{1}{m} \sum_{m=1}^M \sum_{t=0}^{\bar{T}-1} D_{\text{JSD}}(\pi_\phi(\cdot|x_{t,m}) \parallel \pi'(\cdot|x_{t,m}))$, where D_{JSD} is the Jensen-Shannon divergence. Pseudo-code of distillation is given in Algorithm 3 in Appendix D.

Safe Q -learning (SDQN): Here we sample an off-policy mini-batch of state, action, cost, and next-state from the replay buffer, and use it to update the value function estimates that minimize the MSE losses of the Bellman residuals. We first construct the state-action Lyapunov function estimate $\hat{Q}_L(x, a; \theta_D, \theta_T) = \hat{Q}_D(x, a; \theta_D) + \tilde{c} \cdot \hat{Q}_T(x, a; \theta_T)$ by learning the constraint value network \hat{Q}_D and stopping time value network \hat{Q}_T . With a current baseline policy π_k , one can use function approximation to approximate the auxiliary constraint cost (which is the solution to (5)) by $\tilde{c}(x) = \tilde{c} = (d_0 - \pi_k(\cdot|x_0)^\top \hat{Q}_D(x_0, \cdot; \theta_D)) / \pi_k(\cdot|x_0)^\top \hat{Q}_T(x_0, \cdot; \theta_T)$. Equipped with the Lyapunov function, at each iteration, one can do a standard DQN update, except that the optimal action probabilities are computed by solving (6). Details of SDQN is given in Algorithm 4 in Appendix D.

Safe Policy Improvement (SDPI): Similar to SDQN, in this algorithm, we first sample an off-policy mini-batch from the replay buffer and use it to update the value function estimates (w.r.t. objective, constraint, and stopping-time estimate) that minimize MSE losses. Different from SDQN, in SDPI the value estimation is done using policy evaluation, which means that the objective Q -function is trained to minimize the Bellman residual w.r.t. actions generated by the current policy π_k , instead of the greedy actions. Using the same construction as in SDQN for auxiliary cost \tilde{c} and state-action Lyapunov function \hat{Q}_L , we then perform a policy improvement step by computing




Figure 1: Results of various planning algorithms on the grid-world environment with obstacles, with x-axis showing the obstacle density. From the leftmost column, the first figure illustrates the 2D planning domain example ($\rho = 0.25$). The second and the third figures show the average return and the average cumulative constraint cost of the CMDP methods, respectively. The fourth figure displays all the methods used in the experiment. The shaded regions indicate the 80% confidence intervals. Clearly the safe DP algorithms compute policies that are safe and have good performance.

a set of greedy action probabilities from (6) and constructing an updated policy π_{k+1} using policy distillation. Assuming both value and policy approximations have low error, SDPI resembles several interesting properties of SPI, such as maintaining safety during training and monotonically improving the policy. To improve learning stability, instead of the full policy update, one can further consider a partial update $\pi_{k+1} = (1-\alpha)\pi_k + \alpha\pi'$, where $\alpha \in (0, 1)$ is a *mixing constant* that controls safety and exploration [2, 19]. Details of SDPI is summarized in Algorithm 5 in Appendix D.

5 Experiments

Motivated by the safety issues of RL in [22], we validate our safe RL algorithms using a stochastic 2D grid-world motion planning problem. In this domain, an agent (e.g., a robotic vehicle) starts in a safe region and its objective is to travel to a given destination. At each time step, the agent can move to any of its four neighboring states. Due to sensing and control noise, however, with probability δ a move to a random neighboring state occurs. To account for fuel usage, the stage-wise cost of each move until reaching the destination is 1, while the reward achieved for reaching the destination is 1000. Thus, we would like the agent to reach the destination in the shortest possible number of moves. In between the starting and destination points, there are number of obstacles that the agent may pass through but should avoid for safety; each time the agent hits an obstacle it incurs a constraint cost of 1. Thus, in the CMDP setting, the agent’s goal is to reach the destination in the shortest possible number of moves, while hitting the obstacles at most d_0 times or less. For demonstration purposes, we choose a 25×25 grid-world (see Figure 1) with a total of 625 states. We also have a density ratio $\rho \in (0, 1)$ that sets the obstacle-to-terrain ratio. When ρ is close to 0, the problem is obstacle-free, and if ρ is close to 1, then the problem becomes more challenging. In the normal problem setting, we choose a density $\rho = 0.3$, an error probability $\delta = 0.05$, a constraint threshold $d_0 = 5$, and a maximum horizon of 200 steps. The initial state is located in $(24, 24)$ and the goal is placed in $(0, \alpha)$, where $\alpha \in [0, 24]$ is a uniform random variable. To account for statistical significance, the results of each experiment are averaged over 20 trials.

CMDP Planning: In this task, we have explicit knowledge of the reward function and transition probability. The main goal is to compare our safe DP algorithms (SPI and SVI) with the following common CMDP baseline methods: (i) *Step-wise Surrogate*, (ii) *Super-martingale Surrogate*, (iii) *Lagrangian*, and (iv) *Dual LP*. Since the methods in (i) and (ii) are surrogate algorithms, we will also evaluate these methods with both value iteration and policy iteration. To illustrate the level of sub-optimality, we will also compare the returns and constraint costs of these methods with baselines that are generated by maximizing return or minimizing constraint cost of two separate MDPs. The main objective here is to illustrate that safe DP algorithms are less conservative than other surrogate methods, are more numerically stable than the Lagrangian method, and are more computationally efficient than the Dual LP method (see Appendix F), without using function approximations.

Figure 1 presents the results on returns and cumulative constraint costs of the aforementioned CMDP methods over a spectrum of ρ values, ranging from 0 to 0.5. In each method, the initial policy is a conservative baseline policy π_B that minimizes the constraint cost. The empirical results indicate that although the policies generated by the four surrogate algorithms are feasible, they do not have significant policy improvement, i.e., return values are close to that of the initial baseline policy. Over all density settings, the SPI algorithm consistently computes a solution that is feasible and has good performance. The policy returned by SVI is always feasible and has near-optimal performance

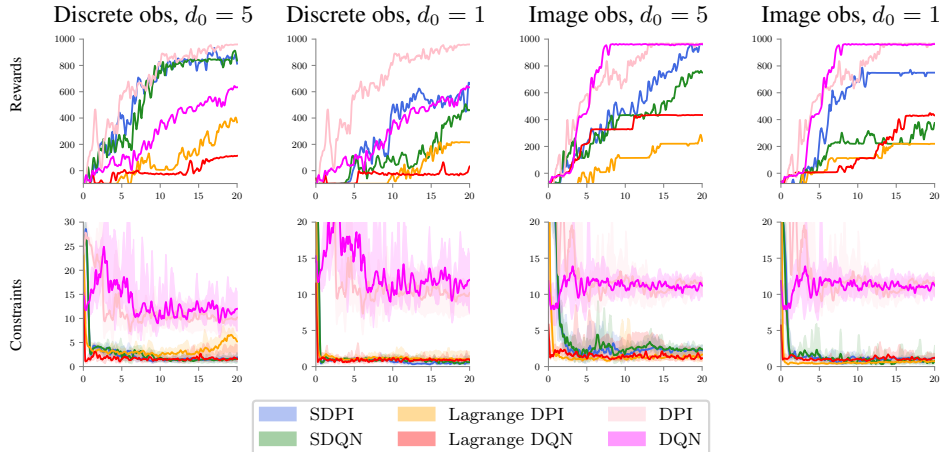


Figure 2: Results of various RL algorithms on the grid-world environment with obstacles, with x-axis in thousands of episodes. We include runs using discrete observations (a one-hot encoding of the agent’s position) and image observations (showing the entire RGB 2D map of the world). We discover that the Lyapunov-based approaches can perform safe learning, despite the fact that the model of the environment is not known and that deep function approximation is necessary.

when the obstacle density is low. However, due to numerical instability, its performance degrades as ρ grows. Similarly, the Lagrangian methods return a near-optimal solution over most settings, but due to numerical issues their solutions start to violate constraint as ρ grows.

Safe Reinforcement Learning: Here we present the results of RL algorithms on this safety task. We evaluate their learning performance on two variants: one in which the observation is a one-hot encoding of the agent’s location, and the other in which the observation is the 2D image representation of the grid map. In each of these, we evaluate performance when $d_0 = 1$ and $d_0 = 5$. We compare our proposed safe RL algorithms, SDPI and SDQN, with their unconstrained counterparts, DPI and DQN, as well as the Lagrangian approach to safe RL, in which the Lagrange multiplier is optimized via extensive grid search. Details of the experimental setup are given in Appendix F. To make the tasks more challenging, we initialize the RL algorithms with a randomized baseline policy.

Figure 2 shows the results of these methods across all task variants. We observe that SDPI and SDQN can adequately solve the tasks and compute good return performance (similar to that of DQN and DPI in some cases), while guaranteeing safety. Another interesting observation in the SDQN and SDPI algorithms is that, once the algorithm finds a safe policy, then all updated policies *remain* safe throughout training. On the contrary, the Lagrangian approach often achieves worse rewards and is more apt to violate the constraints during training,⁶ and the performance is very sensitive to the initial conditions. Furthermore, in some cases (in experiment with $d_0 = 5$ and with discrete observations) the Lagrangian method cannot guarantee safety throughout training.

6 Conclusion

In this paper, we formulated the problem of safe RL as a CMDP and proposed a *novel* Lyapunov approach to solve CMDPs. We also derived an effective LP-based method to generate Lyapunov functions, such that the corresponding algorithm guarantees feasibility and optimality under certain conditions. Leveraging these theoretical underpinnings, we showed how Lyapunov approaches can be used to transform DP (and RL) algorithms into their safe counterparts that only require straightforward modifications in the algorithm implementations. We empirically validated our theoretical findings in using the Lyapunov approach to guarantee safety and robust learning in RL. In general, our work represents a step forward in deploying RL to real-world problems in which guaranteeing safety is of paramount importance. Future research will focus on two directions. On the algorithmic perspective, one major extension is to apply the Lyapunov approach to policy gradient algorithms and compare its performance with CPO in continuous action problems. On the practical aspect, future work includes evaluating the Lyapunov-based RL algorithms on several real-world testbeds.

⁶In Appendix F, we also report the results from the Lagrangian method in which the Lagrange multiplier is learned using gradient ascent method [10] and we observe similar (or even worse) behaviors.

References

- [1] N. Abe, P. Melville, C. Pendus, C. Reddy, D. Jensen, V. Thomas, J. Bennett, G. Anderson, B. Cooley, M. Kowalczyk, et al. Optimizing debt collections using constrained reinforcement learning. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 75–84, 2010.
- [2] J. Achiam, D. Held, A. Tamar, and P. Abbeel. Constrained policy optimization. In *International Conference of Machine Learning*, 2017.
- [3] E. Altman. *Constrained Markov decision processes*, volume 7. CRC Press, 1999.
- [4] Eitan Altman. Constrained Markov decision processes with total cost criteria: Lagrangian approach and dual linear program. *Mathematical methods of operations research*, 48(3):387–417, 1998.
- [5] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, and D. Mané. Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565*, 2016.
- [6] F. Berkenkamp, M. Turchetta, A. Schoellig, and A. Krause. Safe model-based reinforcement learning with stability guarantees. In *Advances in Neural Information Processing Systems*, pages 908–919, 2017.
- [7] D. Bertsekas. *Dynamic programming and optimal control*. Athena scientific Belmont, MA, 1995.
- [8] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [9] M El Chamie, Y. Yu, and B. Açıkmeşe. Convex synthesis of randomized policies for controlled Markov chains with density safety upper bound constraints. In *American Control Conference*, pages 6290–6295, 2016.
- [10] Y. Chow, M. Ghavamzadeh, L. Janson, and M. Pavone. Risk-constrained reinforcement learning with percentile risk criteria. *arXiv preprint arXiv:1512.01629*, 2015.
- [11] Y. Chow, M. Pavone, B. Sadler, and S. Carpin. Trading safety versus performance: Rapid deployment of robotic swarms with robust performance constraints. *Journal of Dynamic Systems, Measurement, and Control*, 137(3), 2015.
- [12] G. Dalal, K. Dvijotham, M. Vecerik, T. Hester, C. Paduraru, and Y. Tassa. Safe exploration in continuous action spaces. *arXiv preprint arXiv:1801.08757*, 2018.
- [13] R. Fruit and A. Lazaric. Exploration–exploitation in mdps with options. In *AISTATS*, 2017.
- [14] Z. Gábor and Z. Kalmár. Multi-criteria reinforcement learning. In *International Conference of Machine Learning*, 1998.
- [15] P. Geibel and F. Wyszotzki. Risk-sensitive reinforcement learning applied to control under constraints. *Journal of Artificial Intelligence Research*, 24:81–108, 2005.
- [16] P. Glynn, A. Zeevi, et al. Bounding stationary expectations of markov processes. In *Markov processes and related topics: a Festschrift for Thomas G. Kurtz*, pages 195–214. Institute of Mathematical Statistics, 2008.
- [17] S. Gu, T. Lillicrap, I. Sutskever, and S. Levine. Continuous deep Q-learning with model-based acceleration. In *International Conference on Machine Learning*, pages 2829–2838, 2016.
- [18] S. Junges, N. Jansen, C. Dehnert, U. Topcu, and J. Katoen. Safety-constrained reinforcement learning for MDPs. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 130–146, 2016.
- [19] S. Kakade and J. Langford. Approximately optimal approximate reinforcement learning. In *International Conference on International Conference on Machine Learning*, pages 267–274, 2002.
- [20] Hassan K Khalil. Nonlinear systems. *Prentice-Hall, New Jersey*, 2(5):5–1, 1996.

- [21] J. Lee, I. Panageas, G. Piliouras, M. Simchowitz, M. Jordan, and B. Recht. First-order methods almost always avoid saddle points. *arXiv preprint arXiv:1710.07406*, 2017.
- [22] J. Leike, M. Martic, V. Krakovna, P. Ortega, T. Everitt, A. Lefrancq, L. Orseau, and S. Legg. Ai safety gridworlds. *arXiv preprint arXiv:1711.09883*, 2017.
- [23] D. Luenberger, Y. Ye, et al. *Linear and nonlinear programming*, volume 2. Springer, 1984.
- [24] N. Mastronarde and M. van der Schaar. Fast reinforcement learning for energy-efficient wireless communication. *IEEE Transactions on Signal Processing*, 59(12):6262–6266, 2011.
- [25] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller. Playing Atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- [26] T. Moldovan and P. Abbeel. Safe exploration in Markov decision processes. *arXiv preprint arXiv:1205.4810*, 2012.
- [27] H. Mossalam, Y. Assael, D. Roijers, and S. Whiteson. Multi-objective deep reinforcement learning. *arXiv preprint arXiv:1610.02707*, 2016.
- [28] M. Neely. Stochastic network optimization with application to communication and queueing systems. *Synthesis Lectures on Communication Networks*, 3(1):1–211, 2010.
- [29] G. Neu, A. Jonsson, and V. Gómez. A unified view of entropy-regularized markov decision processes. *arXiv preprint arXiv:1705.07798*, 2017.
- [30] M. Ono, M. Pavone, Y. Kuwata, and J. Balaram. Chance-constrained dynamic programming with application to risk-aware robotic space exploration. *Autonomous Robots*, 39(4):555–571, 2015.
- [31] T. Perkins and A. Barto. Lyapunov design for safe reinforcement learning. *Journal of Machine Learning Research*, 3:803–832, 2002.
- [32] M. Pirodda, M. Restelli, and L. Bascetta. Adaptive step-size for policy gradient methods. In *Advances in Neural Information Processing Systems*, pages 1394–1402, 2013.
- [33] M. Pirodda, M. Restelli, A. Pecorino, and D. Calandriello. Safe policy iteration. In *International Conference on Machine Learning*, pages 307–315, 2013.
- [34] K. Regan and C. Boutilier. Regret-based reward elicitation for Markov decision processes. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pages 444–451, 2009.
- [35] D. Roijers, P. Vamplew, S. Whiteson, and R. Dazeley. A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research*, 48:67–113, 2013.
- [36] A. Rusu, S. Colmenarejo, C. Gulcehre, G. Desjardins, J. Kirkpatrick, R. Pascanu, V. Mnih, K. Kavukcuoglu, and R. Hadsell. Policy distillation. *arXiv preprint arXiv:1511.06295*, 2015.
- [37] T. Schaul, J. Quan, I. Antonoglou, and D. Silver. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*, 2015.
- [38] Bruno Scherrer. Performance bounds for λ policy iteration and application to the game of Tetris. *Journal of Machine Learning Research*, 14(Apr):1181–1227, 2013.
- [39] M. Schmitt and L. Martignon. On the complexity of learning lexicographic strategies. *Journal of Machine Learning Research*, 7:55–83, 2006.
- [40] G. Shani, D. Heckerman, and R. Brafman. An MDP-based recommender system. *Journal of Machine Learning Research*, 6:1265–1295, 2005.
- [41] A. Tamar, D. Di Castro, and S. Mannor. Policy gradients with variance related risk criteria. In *International Conference of Machine Learning*, 2012.
- [42] H. van Hasselt. Double Q-learning. In *Advances in Neural Information Processing Systems*, pages 2613–2621, 2010.

A Safety Constraints in Planning Problems

To motivate the CMDP formulation studied in this paper, in this section we include two real-life examples of modeling safety using the reachability constraint, and the constraint that limits the agent's visits to undesirable states.

A.1 Reachability Constraint

Reachability is a common concept in motion-planning and engineering applications, where for any given policy π and initial state x_0 , the following the constraint function is considered:

$$\mathbb{P}(\exists t \in \{0, 1, \dots, T^* - 1\}, x_t \in \mathcal{S}_H \mid x_0, \pi).$$

Here \mathcal{S}_H represents the real subset of hazardous regions for the states and actions. Therefore, the constraint cost represents the probability of reaching an unsafe region at any time before the state reaches the terminal state. To further analyze this constraint function, one notices that

$$\begin{aligned} \mathbb{P}(\exists t \in \{0, 1, \dots, T^* - 1\}, x_t \in \mathcal{S}_H \mid x_0, \pi) &= \mathbb{P} \left(\bigcup_{t=0}^{T^*-1} \bigcap_{j=0}^{t-1} \{x_j \notin \mathcal{S}_H\} \cap \{x_t \in \mathcal{S}_H\} \mid x_0, \pi \right) \\ &= \mathbb{E} \left[\sum_{t=0}^{T^*-1} \prod_{j=0}^{t-1} \mathbf{1}\{x_j \notin \mathcal{S}_H\} \cdot \mathbf{1}\{x_t \in \mathcal{S}_H\} \mid x_0, \pi \right]. \end{aligned}$$

In this case, a policy π is deemed *safe* if the reachability probability to the unsafe region is bounded by threshold $d_0 \in (0, 1)$, i.e.,

$$\mathbb{P}(\exists t \in \{0, 1, \dots, T^* - 1\}, x_t \in \mathcal{S}_H \mid x_0, \pi) \leq d_0. \quad (7)$$

To transform the reachability constraint into a standard CMDP constraint, we define an additional state $s \in \{0, 1\}$ that keeps track of the reachability status at time t . Here $s_t = 1$ indicates the system has never visited a hazardous region up till time t , and otherwise $s_t = 0$. Let $s_0 = 1$, we can easily see that by defining the following deterministic transition

$$s_t = s_{t-1} \cdot \mathbf{1}\{x_t \notin \mathcal{S}_H\}, \quad \forall t \geq 1,$$

s_t has the following formulation: $s_t = \prod_{j=0}^{t-1} \mathbf{1}\{x_j \notin \mathcal{S}_H\}$.

Collectively, with the state augmentation $\hat{x} = (x, s)$, one defines the augmented CMDP $(\hat{\mathcal{X}}, \mathcal{A}, C, \hat{D}, \hat{P}, \hat{x}_0, d_0)$, where $\hat{\mathcal{X}} = \mathcal{X} \times \{0, 1\}$ is the augmented state space, $\hat{d}(\hat{x}) = s \cdot d(x)$ is the augmented constraint cost, $\hat{P}(\hat{x}' | \hat{x}, a) = P(x' | x, a) \cdot \mathbf{1}\{s' = s \cdot \mathbf{1}\{x \notin \mathcal{S}_H\}\}$ is the augmented transition probability, and $\hat{x}_0 = (x_0, 1)$ is the initial (augmented) state. By using this augmented CMDP, immediately the reachability constraint is equivalent to $\mathbb{E}[\sum_{t=0}^{T^*-1} \hat{d}(\hat{x}_t) \mid x_0, \pi] \leq d_0$.

A.2 Constraint w.r.t. Undesirable Regions of States

Consider the notion of safety where one restricts the total visiting frequency of an agent to an undesirable region (of states). This notion of safety appears in applications such as system maintenance, in which the system can only tolerate its state to visit (in expectation) a hazardous region, namely \mathcal{S}_H , for a fixed number of times. Specifically, for given initial state x_0 , consider the following constraint that bounds the total frequency of visiting \mathcal{S}_H with a pre-defined threshold d_0 , i.e., $\mathbb{E}[\sum_{t=0}^{T^*-1} d(x_t) \mid x_0, \pi] \leq d_0$, where $d(x_t) = \mathbf{1}\{x_t \in \mathcal{S}_H\}$. To model this notion of safety using a CMDP, one can rewrite the above constraint using the constraint immediate cost $d(x) = \mathbf{1}\{x \in \mathcal{S}_H\}$, and the constraint threshold d_0 . To study the connection between the reachability constraint, and the above constraint w.r.t. undesirable region, notice that

$$\mathbb{E} \left[\sum_{t=0}^{T^*-1} \prod_{j=0}^{t-1} \mathbf{1}\{x_j \notin \mathcal{S}_H\} \cdot \mathbf{1}\{x_t \in \mathcal{S}_H\} \mid x_0, \pi \right] \leq \mathbb{E} \left[\sum_{t=0}^{T^*-1} \mathbf{1}\{x_t \in \mathcal{S}_H\} \mid x_0, \pi \right].$$

This clearly indicates that any policies which satisfies the constraint w.r.t. undesirable region, also satisfies the reachability constraint.

B Existing Approaches for Solving CMDPs

Before going to the main result, we first revisit several existing CMDP algorithms in the literature, which later serve as the baselines for comparing with our safe CMDP algorithms. For the sake of brevity, we will only provide an overview of these approaches here and defer their details to Appendix B.1.

The Lagrangian Based Algorithm: The standard way of solving problem \mathcal{OPT} is by applying the Lagrangian method. To start with, consider the following minimax problem: $\min_{\pi \in \Delta} \max_{\lambda \geq 0} \mathcal{C}_\pi(x_0) + \lambda(\mathcal{D}_\pi(x_0) - d_0)$, where λ is the Lagrange multiplier w.r.t. the CMDP constraint. According to Theorem 9.9 and Theorem 9.10 in [3], the optimal policy π^* of problem \mathcal{OPT} can be calculated by solving the following Lagrangian function $\pi^* \in \arg \min_{\pi \in \Delta} \mathcal{C}_\pi(x_0) + \lambda_{d_0}^* \mathcal{D}_\pi(x_0)$, where $\lambda_{d_0}^*$ is the optimal Lagrange multiplier. Utilizing this result, one can compute the saddle point pair (π^*, λ^*) using primal-dual iteration. Specifically, for a given $\lambda \geq 0$, solve the policy minimization problem using standard dynamic programming with λ -parametrized Bellman operator $T_\lambda[V](x) = \min_{\pi \in \Delta(x)} T_{\pi, c + \lambda d}[V](x)$ if $x \in \mathcal{X}$; For a given policy π , solve for the following linear optimization problem: $\max_{\lambda \geq 0} \mathcal{C}_\pi(x_0) + \lambda(\mathcal{D}_\pi(x_0) - d_0)$. Based on Theorem 9.10 in [3], this procedure will asymptotically converge to the saddle point solution. However, this algorithm presents several major challenges. (i) In general there is no known convergence rate guarantees, several studies [21] also showed that using primal-dual first-order iterative method to find saddle point may run into numerical instability issues; (ii) Choosing a good initial estimate of the Lagrange multiplier is not intuitive; (iii) Following the same arguments from [2], during iteration the policy may be infeasible w.r.t. problem \mathcal{OPT} , and feasibility is guaranteed after the algorithm converges. This is hazardous in RL when one needs to execute the intermediate policy (which may be unsafe) during training.

The Dual LP Based Algorithm: Another method of solving problem \mathcal{OPT} is based on computing its *occupation measures* w.r.t. the optimal policy. In transient MDPs, for any given policy π and initial state x_0 , the state-action occupation measure is $\rho_\pi(x, a) = \mathbb{E} \left[\sum_{t=0}^{T^*-1} \mathbf{1}\{x_t = x, a_t = a\} \mid x_0, \pi \right]$, which characterizes the total visiting probability of state-action pair $(x, a) \in \mathcal{X} \times \mathcal{A}$, induced by policy π and initial state x_0 . Utilizing this quantity, Theorem 9.13 in [3], has shown that problem \mathcal{OPT} can be reformulated as a linear programming (LP) problem (see Equation (8) to (9) in Appendix B.1), whose decision variable is of dimension $|\mathcal{X}'| |\mathcal{A}|$, and it has $2|\mathcal{X}'| |\mathcal{A}| + 1$ constraints. Let ρ^* be the solution of this LP, the optimal Markov stationary policy is given by $\pi^*(a|x) = \rho^*(x, a) / \sum_{a \in \mathcal{A}} \rho^*(x, a)$. To solve this problem, one can apply the standard algorithm such as interior point method, which is a strong polynomial time algorithm with complexity $O(|\mathcal{X}'|^2 |\mathcal{A}|^2 (2|\mathcal{X}'| |\mathcal{A}| + 1))$ [8]. While this is a straight-forward methodology, it can only handle CMDPs with finite state and action spaces. Furthermore, this approach is computationally expensive when the size of these spaces are large. To the best of our knowledge, it is also unclear how to extend this approach to RL, when transition probability and immediate reward/constraint reward functions are unknown.

Step-wise Constraint Surrogate Approach: This approach transforms the multi-stage CMDP constraint into a sequence of step-wise constraints, where each step-wise constraint can be directly embedded into set of admissible actions in the Bellman operator. To start with, for any state $x \in \mathcal{X}'$, consider the following feasible set of policies: $\mathcal{F}^{\text{SW}}(d_0, x) = \{ \pi \in \Delta : \sum_{a \in \mathcal{A}} \sum_{x' \in \mathcal{X}'} \pi(a|x) P(x'|x, a) d(x') \leq \frac{d_0}{T} \}$, where \bar{T} is the upper-bound of the MDP stopping time. Based on (10) in Appendix B.1, one deduces that every policy π in $\bigcap_{x \in \mathcal{X}'} \mathcal{F}^{\text{SW}}(d_0, x)$ is a feasible policy w.r.t. problem \mathcal{OPT} . Motivated by this observation, a solution policy can be solved by $\min_{\pi \in \bigcap_{x \in \mathcal{X}'} \mathcal{F}^{\text{SW}}(d_0, x)} \mathbb{E} \left[\sum_{t=0}^{T^*-1} c(x_t, a_t) \mid x_0, \pi \right]$. One benefit of studying this surrogate problem is that its solution satisfies the Bellman optimality condition w.r.t. the step-wise Bellman operator as $T^{\text{SW}}[V](x) = \min_{\pi \in \mathcal{F}^{\text{SW}}(d_0, x)} \sum_a \pi(a|x) [c(x, a) + \sum_{x' \in \mathcal{X}'} P(x'|x, a) V(x')]$ for any $x \in \mathcal{X}'$. In particular T^{SW} is a contraction operator, which implies that there exists a unique solution $V^{*, \text{SW}}$ to fixed point equation $T^{\text{SW}}[V](x) = V(x)$ for $x \in \mathcal{X}'$ such that $V^{*, \text{SW}}(x_0)$ is a solution to the surrogate problem. Therefore this problem can be solved by standard DP methods such as value iteration or policy iteration. Furthermore, based on the structure of $\mathcal{F}^{\text{SW}}(d_0, x)$, any surrogate policy

is feasible w.r.t. problem \mathcal{OPT} . However, the major drawback is that the step-wise constraint in $\mathcal{F}^{\text{SW}}(d_0, x)$ can be much more stringent than the original safety constraint in problem \mathcal{OPT} .

Super-martingale Constraint Surrogate Approach: This surrogate algorithm is originally proposed by [14], where the CMDP constraint is reformulated as the surrogate value function $\mathcal{DS}_\pi(x) = \max\{d_0, \mathcal{D}_\pi(x)\}$ at initial state $x \in \mathcal{X}'$. It has been shown that an arbitrary policy π is a feasible policy of the CMDP if and only if $\mathcal{DS}_\pi(x_0) = d_0$. Notice that \mathcal{DS}_π is known as a *super-martingale surrogate*, due to the inequality $\mathcal{DS}_\pi(x) \leq T_\pi^{\text{DS}}[\mathcal{DS}_\pi](x)$ with respect to the contraction Bellman operator $T_\pi^{\text{DS}}[V](x) = \sum_{a \in \mathcal{A}} \pi(a|x) \max\{d_0, d(x) + \sum_{x' \in \mathcal{X}'} P(x'|x, a)V(x')\}$ of the constraint value function. However, for arbitrary policy π , in general it is non-trivial to compute the value function $\mathcal{DS}_\pi(x)$, and instead one can easily compute its upper-bound value function $\overline{\mathcal{DS}}_\pi(x)$ which is the solution of the fixed-point equation $V(x) = T_\pi^{\text{DS}}[V](x)$, $\forall x \in \mathcal{X}'$, using standard dynamic programming techniques. To better understand how this surrogate value function guarantees feasibility in problem \mathcal{OPT} , at each state $x \in \mathcal{X}'$ consider the optimal value function of the minimization problem $\overline{\mathcal{DS}}(x) = \min_{\pi \in \Delta} \overline{\mathcal{DS}}_\pi(x)$. Then whenever $\overline{\mathcal{DS}}(x_0) \leq d_0$, the corresponding solution policy π is a feasible policy of problem \mathcal{OPT} , i.e., $\mathcal{D}_\pi(x_0) \leq \overline{\mathcal{DS}}_\pi(x_0) = d_0$. Now define $\mathcal{F}^{\text{DS}}(x) := \{\pi \in \Delta : T_\pi^{\text{DS}}[\overline{\mathcal{DS}}](x) = \overline{\mathcal{DS}}(x)\}$ as the set of refined feasible policies induced by $\overline{\mathcal{DS}}$. If the condition $\overline{\mathcal{DS}}(x_0) \leq d_0$ holds, then all the policies in $\mathcal{F}^{\text{DS}}(x)$ are feasible w.r.t. problem \mathcal{OPT} . Utilizing this observation, a surrogate solution policy of problem \mathcal{OPT} can be found by computing the solution policy of the fixed-point equation $T_{\mathcal{F}^{\text{DS}}}[V](x) = V(x)$, for $x \in \mathcal{X}'$, where $T_{\mathcal{F}^{\text{DS}}}[V](x) = \min_{\pi \in \mathcal{F}^{\text{DS}}(d_0, x)} T_{\pi, c}[V](x)$. Notice that $T_{\mathcal{F}^{\text{DS}}}$ is a contraction operator, this procedure can also be solved using standard DP methods. The major benefit of this 2-step approach is that the computation of the feasibility set is decoupled from solving the optimization problem. This allows us to apply approaches such as the *lexicographical ordering* method from multi-objective stochastic optimal control methods [35] to solve the CMDP, for which the constraint value function has a higher lexicographical order than the objective value function. However, since the refined set of feasible policies is constructed prior to policy optimization, it might still be overly conservative. Furthermore, even if there exists a non-trivial solution policy to the surrogate problem, characterizing its sub-optimality performance bound remains a challenging task.

B.1 Details of Existing Solution Algorithms

In this section, we provide the details of the existing algorithms for solving CMDPs.

The Lagrangian Based Algorithm: The standard way of solving problem \mathcal{OPT} is by applying the Lagrangian method. To start with, consider the following minimax problem:

$$\min_{\pi \in \Delta} \max_{\lambda \geq 0} \mathcal{C}_\pi(x_0) + \lambda(\mathcal{D}_\pi(x_0) - d_0),$$

where λ is the Lagrange multiplier of the CMDP constraint, and the Lagrangian function is given by

$$\mathcal{L}_{x_0, d_0}(\pi, \lambda) = \mathcal{C}_\pi(x_0) + \lambda(\mathcal{D}_\pi(x_0) - d_0) = \mathbb{E} \left[\sum_{t=0}^{T^*-1} c(x_t, a_t) + \lambda d(x_t) \mid \pi, x_0 \right] - \lambda d_0.$$

A solution pair (π^*, λ^*) is considered as a saddle point of Lagrangian function $\mathcal{L}_{x_0, d_0}(\pi, \lambda)$ if the following condition holds:

$$\mathcal{L}_{x_0, d_0}(\pi^*, \lambda) \leq \mathcal{L}_{x_0, d_0}(\pi, \lambda) \leq \mathcal{L}_{x_0, d_0}(\pi, \lambda^*), \forall \pi, \lambda \geq 0.$$

According to Theorem 9.10 in [3], suppose the interior set of feasible set of problem \mathcal{OPT} is non-empty, then there exists a solution pair (π^*, λ^*) to the minimax problem that is a saddle point of Lagrangian function $\mathcal{L}_{x_0, d_0}(\pi, \lambda)$. Furthermore, Theorem 9.9 in [3] shows that strong duality holds:

$$\min_{\pi \in \Delta} \max_{\lambda \geq 0} \mathcal{L}_{x_0, d_0}(\pi, \lambda) = \max_{\lambda \geq 0} \min_{\pi \in \Delta} \mathcal{L}_{x_0, d_0}(\pi, \lambda).$$

This implies that the optimal policy $\pi^* \in \Delta$ can be calculated by solving the following Lagrangian function $\pi^* \in \arg \min_{\pi \in \Delta} \mathcal{L}_{x_0, d_0}(\pi, \lambda^*)$, with optimal Lagrange multiplier λ^* .

Utilizing the structure of the Lagrangian function $\mathcal{L}_{x_0, d_0}(\pi, \lambda)$, for any fixed Lagrange multiplier $\lambda \geq 0$, consider the λ -Bellman operator $T_\lambda[V]$, where

$$T_\lambda[V](x) = \begin{cases} \min_{\pi \in \Delta(x)} T_{\pi, c + \lambda d}[V](x) & \text{if } x \in \mathcal{X}', \\ 0 & \text{otherwise.} \end{cases}$$

Since $T_\lambda[V]$ is a γ -contraction operator, by the Bellman principle of optimality, there is a unique solution V^* to the fixed point equation $T_\lambda[V](x) = V(x)$ for $x \in \mathcal{X}$, which can be solved by dynamic programming algorithms, such as value iteration or policy iteration. Furthermore, the λ -optimal policy π_λ^* has the form of

$$\pi_\lambda^*(\cdot|x) \in \arg \min_{\pi \in \Delta(x)} T_{\pi, c + \lambda d}[V^*](x) \quad \text{if } x \in \mathcal{X}',$$

and $\pi_\lambda^*(\cdot|x)$ is an arbitrary probability distribution function if $x \notin \mathcal{X}'$.

The Dual LP Based Algorithm: The other commonly-used method for solving problem \mathcal{OPT} is based on computing its *occupation measures* w.r.t. the optimal policy. In a transient MDP, for any given policy π and initial state $x_0 \in \mathcal{X}'$ the state-action occupation measure is defined as

$$\rho_{\pi, x_0}(x, a) = \mathbb{E} \left[\sum_{t=0}^{T^*-1} \mathbf{1}\{x_t = x, a_t = a\} \mid x_0, \pi \right], \quad \forall x \in \mathcal{X}, \forall a \in \mathcal{A},$$

and the occupation measure at state $x \in \mathcal{X}$ is defined as $\rho_{\pi, x_0}(x) = \sum_{a \in \mathcal{A}} \rho_{\pi, x_0}(x, a)$. Clearly these two occupation measures are related by the following property: $\rho_{\pi, x_0}(x, a) = \rho_{\pi, x_0}(x) \cdot \pi(a|x)$. Furthermore, using the fact that a occupation measure $\rho_{\pi, x_0}(x, a)$ is indeed the sum of visiting distribution of the transient MDP induced by policy π , one clearly sees that it satisfies the following set of constraints:

$$\mathcal{Q}(x_0) = \left\{ \rho : \mathcal{X}' \times \mathcal{A} \rightarrow \mathbb{R} : \rho(x, a) \geq 0, \forall x \in \mathcal{X}', a \in \mathcal{A}, \right. \\ \left. \sum_{x_p \in \mathcal{X}', a \in \mathcal{A}} \rho(x_p, a) (\mathbf{1}\{x_p = x\} - P(x|x_p, a)) = \mathbf{1}\{x = x_0\} \right\}.$$

Therefore, by Theorem 9.13 in [3], equivalently problem \mathcal{OPT} can be solved by the LP optimization problem with $2|\mathcal{X}'||\mathcal{A}| + 1$ constraints:

$$\min_{\rho \in \mathcal{Q}(x_0)} \sum_{x \in \mathcal{X}', a \in \mathcal{A}} \rho(x, a) c(x, a) \quad (8)$$

$$\text{subject to} \quad \sum_{x \in \mathcal{X}', a \in \mathcal{A}} \rho(x, a) d(x) \leq d_0, \quad (9)$$

and equipped with the minimizer $\rho^* \in \mathcal{Q}(x_0)$, the (non-uniform) optimal Markovian stationary policy is given by the following form:

$$\pi^*(a|x) = \frac{\rho^*(x, a)}{\sum_{a \in \mathcal{A}} \rho^*(x, a)}, \quad \forall x \in \mathcal{X}', \forall a \in \mathcal{A}.$$

Step-wise Constraint Surrogate Approach: To start with, without loss of generality assume that the agent is safe at the initial phase, i.e., $d(x_0) \leq 0$. For any state $x \in \mathcal{X}'$, consider the following feasible set of policies:

$$\mathcal{F}^{\text{SW}}(d_0, x) = \left\{ \pi \in \Delta(x) : \sum_{a \in \mathcal{A}} \sum_{x' \in \mathcal{X}'} \pi(a|x) P(x'|x, a) d(x') \leq \frac{d_0}{\bar{T}} \right\},$$

where \bar{T} is the uniform upper-bound of the random stopping time in the transient MDP. Immediately, for any policy $\pi \in \bigcap_{x \in \mathcal{X}'} \mathcal{F}^{\text{SW}}(d_0, x)$, one has the following inequality:

$$\mathbb{E} \left[\sum_{t=0}^{T^*-1} d(x_t) \mid \pi, x_0 \right] = d(x_0) + \sum_{x \in \mathcal{X}', a \in \mathcal{A}} \rho_{\pi, x_0}(x, a) \sum_{x' \in \mathcal{X}'} P(x'|x, a) d(x') \\ \leq \mathbb{E} \left[\sum_{t=0}^{T^*-1} \frac{d_0}{\bar{T}} \mid \pi, x_0 \right] \leq d_0, \quad (10)$$

where $\rho_{\pi, x_0}(x, a) = \mathbb{E} \left[\sum_{t=0}^{T^*-1} \mathbf{1}\{x_t = x\} | \pi, x_0 \right]$ is the state-action occupation measure with initial state x_0 and policy π , which implies that the policy π is safe, i.e., $\mathcal{D}_\pi(x_0) \leq d_0$. Equipped with this property, we propose the following surrogate problem for problem \mathcal{OPT} , whose solution (if exists) is guaranteed to be safe:

Problem $\mathcal{OPT}^{\text{SW}}$: Given an initial state x_0 , a threshold d_0 , solve
$$\min_{\pi \in \cap_{x \in \mathcal{X}'} \mathcal{F}^{\text{SW}}(d_0, x)} \mathbb{E} \left[\sum_{t=0}^{T^*-1} c(x_t, a_t) | x_0, \pi \right].$$

To solve problem $\mathcal{OPT}^{\text{SW}}$, for each state $x \in \mathcal{X}'$ define the step-wise Bellman operator as

$$T^{\text{SW}}[V](x) = \min_{\pi \in \mathcal{F}^{\text{SW}}(d_0, x)} T_{\pi, c}[V](x).$$

Based on the standard arguments in [7], the Bellman operator T^{SW} is a contraction operator, which implies that there exists a unique solution $V^{*, \text{SW}}$ to the fixed-point equation $T^{\text{SW}}[V](x) = V(x)$, for any $x \in \mathcal{X}'$, such that $V^{*, \text{SW}}(x_0)$ is a solution to problem $\mathcal{OPT}^{\text{SW}}$.

Super-martingale Constraint Surrogate Approach: The following surrogate algorithm is proposed by [14]. Before going to the main algorithm, first consider the following surrogate constraint value function w.r.t. policy $\pi \in \Delta$ and state $x \in \mathcal{X}'$:

$$\begin{aligned} \mathcal{DS}_\pi(x) &= \max \{ d_0, \mathcal{D}_\pi(x) \} = \max \left\{ d_0, d(x) + \sum_{a \in \mathcal{A}} \pi(a|x) \sum_{x' \in \mathcal{X}'} P(x'|x, a) \mathcal{D}_\pi(x') \right\} \\ &\leq \sum_{a \in \mathcal{A}} \pi(a|x) \max \left\{ d_0, d(x) + \sum_{x' \in \mathcal{X}'} P(x'|x, a) \mathcal{D}_\pi(x') \right\}. \end{aligned}$$

The last inequality is due to the fact that the max operator is convex. Clearly, by definition one has

$$\mathcal{DS}_\pi(x) \geq \mathcal{D}_\pi(x) \text{ for each } x \in \mathcal{X}'.$$

On the other hand, one also has the following property: $\mathcal{DS}_\pi(x_0) = d_0$ if and only if the constraint of problem \mathcal{OPT} is satisfied, i.e., $\mathcal{D}_\pi(x_0) \leq d_0$. Now, by utilizing the contraction operator

$$T_\pi^{\text{DS}}[V](x) = \sum_{a \in \mathcal{A}} \pi(a|x) \max \left\{ d_0, d(x) + \sum_{x' \in \mathcal{X}'} P(x'|x, a) V(x') \right\},$$

w.r.t. policy π , and by utilizing the definition of the constraint value function \mathcal{DS}_π , one immediately has the chain of inequalities:

$$\mathcal{DS}_\pi(x) \leq T_\pi^{\text{DS}}[\mathcal{D}_\pi](x) \leq T_\pi^{\text{DS}}[\mathcal{DS}_\pi](x),$$

which implies that the value function $\{\mathcal{DS}_\pi(x)\}_{x \in \mathcal{X}'}$ is a *super-martingale*, i.e., $\mathcal{DS}_\pi(x) \leq T_\pi^{\text{DS}}[\mathcal{DS}_\pi](x)$, for $x \in \mathcal{X}'$.

However in general the constraint value function of interest, i.e., \mathcal{DS}_π , cannot be directly obtained as the solution of fixed point equation. Thus in what follows, we will work with its approximation $\overline{\mathcal{DS}}_\pi(x)$, which is the fixed-point solution of

$$V(x) = T_\pi^{\text{DS}}[V](x), \quad \forall x \in \mathcal{X}'.$$

By definition, the following properties always hold:

$$\overline{\mathcal{DS}}_\pi(x) \geq d_0, \quad \overline{\mathcal{DS}}_\pi(x) \geq \mathcal{DS}_\pi(x) \geq \mathcal{D}_\pi(x), \quad \text{for any } x \in \mathcal{X}'.$$

To understand how this surrogate value function guarantees feasibility in problem \mathcal{OPT} , consider the optimal value function

$$\overline{\mathcal{DS}}(x) = \min_{\pi \in \Delta(x)} \overline{\mathcal{DS}}_\pi(x), \quad \forall x \in \mathcal{X}'.$$

which is also the unique solution w.r.t. the fixed-point equation $\min_{\pi \in \Delta(x)} T_{\pi}^{DS}[V](x) = V(x)$. Now suppose at state x_0 , the following condition holds: $\overline{DS}(x_0) \leq d_0$. Then there exists a policy π that is safe w.r.t. problem \mathcal{OPT} , i.e., $\mathcal{D}_{\pi}(x_0) \leq \overline{DS}_{\pi}(x_0) = d_0$.

Motivated by the above observation, we first check if the following condition holds:

$$\overline{DS}(x_0) \leq d_0.$$

If that is the case, define the set of feasible policies that is induced by the super-martingale \overline{DS} as

$$\mathcal{F}^{DS}(x) := \{\pi \in \Delta : T_{\pi}^{DS}[\overline{DS}](x) = \overline{DS}(x)\},$$

and solve the following problem, whose solution (if exists) is guaranteed to be safe.

Problem \mathcal{OPT}^{DS} : Assume $\overline{DS}(x_0) \leq d_0$. Then given an initial state $x_0 \in \mathcal{X}'$, and a threshold $d_0 \in \mathbb{R}_{\geq 0}$, solve $\min_{\pi \in \bigcup_{x \in \mathcal{X}'} \mathcal{F}^{DS}(x)} \mathbb{E} \left[\sum_{t=0}^{T^*-1} c(x_t, a_t) \mid x_0, \pi \right]$.

Similar to the step-wise approach, clearly the \overline{DS} -induced Bellman operator

$$T_{\mathcal{F}^{DS}}[V](x) = \min_{\pi \in \mathcal{F}^{DS}(d_0, x)} T_{\pi, c}[V](x)$$

is a contraction operator. This implies that there exists a unique solution $V_{\mathcal{F}^{DS}}^*$ to the fixed-point equation $T_{\mathcal{F}^{DS}}[V](x) = V(x)$, for $x \in \mathcal{X}'$, such that $V_{\mathcal{F}^{DS}}^*(x_0)$ is a solution to problem \mathcal{OPT}^{DS} .

C Proofs of the Technical Results in Section 3

C.1 Proof of Lemma 1

In the following part of the analysis we will use shorthand notation P^* to denote the transition probability for $x \in \mathcal{X}'$ induced by the optimal policy, and P_B to denote the transition probability for $x \in \mathcal{X}'$ induced by the baseline policy. These matrices are sub-stochastic because we exclude the terms in the recurrent states. This means that both spectral radii $\rho(P^*)$ and $\rho(P_B)$ are less than 1, and thus both $(I - P^*)$ and $(I - P_B)$ are invertible. By the Neumann series expansion, one can also show that

$$(I - P^*)^{-1} = \left\{ \sum_{t=0}^{T^*-1} \mathbb{P}(x_t = x' | x_0 = x, \pi^*) \right\}_{x, x' \in \mathcal{X}'},$$

and

$$(I - P_B)^{-1} = \left\{ \sum_{t=0}^{T^*-1} \mathbb{P}(x_t = x' | x_0 = x, \pi_B) \right\}_{x, x' \in \mathcal{X}'}. .$$

We also define $\Delta(a|x) = \pi_B(a|x) - \pi^*(a|x)$ for any $x \in \mathcal{X}'$ and $a \in \mathcal{A}$, and $P_\Delta = \{\sum_{a \in \mathcal{A}} P(x'|x, a)\Delta(a|x)\}_{x, x' \in \mathcal{X}'}$. Therefore, one can easily see that

$$(I - P^*)(I - P_B + P_\Delta)^{-1} = I_{|\mathcal{X}'| \times |\mathcal{X}'|}.$$

Therefore, by the Woodbury Sherman Morrison identity, we have that

$$(I - P^*)^{-1} = (I - P_B)^{-1}(I_{|\mathcal{X}'| \times |\mathcal{X}'|} + P_\Delta(I - P^*)^{-1}). \quad (11)$$

By multiplying the constraint cost function vector $d(x)$ on both sides of the above equality, This further implies that for each $x \in \mathcal{X}'$, one has

$$\mathcal{D}_{\pi^*}(x) = \mathbb{E} \left[\sum_{t=0}^{T^*-1} d(x_t) + \epsilon(x_t) \mid \pi_B, x \right] = L_\epsilon(x), \quad (12)$$

such that

$$T_{\pi^*}[L_\epsilon](x) = L_\epsilon(x), \quad \forall x \in \mathcal{X}'.$$

Here, the auxiliary constraint cost is given by

$$\epsilon(x) = \sum_{a \in \mathcal{A}} \Delta(a|x) \sum_{x' \in \mathcal{X}'} P(x'|x, a) \mathcal{D}_{\pi^*}(x'). \quad (13)$$

By construction, equation (11) immediately implies that L_ϵ is a fixed point solution of $T_{\pi^*}[V](x) = V(x)$ for $x \in \mathcal{X}'$. Furthermore, equation (13) further implies that the upper bound of the constraint cost ϵ is given by:

$$-2\bar{T}D_{\max}D_{TV}(\pi^*||\pi_B)(x) \leq \epsilon(x) \leq 2\bar{T}D_{\max}D_{TV}(\pi^*||\pi_B)(x), \quad x \in \mathcal{X}',$$

where \bar{T} is the uniform upper-bound of the MDP stopping time.

Since π^* is also a feasible policy of problem \mathcal{OPT} , this further implies that $L_\epsilon(x_0) \leq d_0$.

C.2 Proof of Theorem 1

First, under Assumption 1 the following inequality holds:

$$D_{TV}(\pi^*||\pi_B)(x) \leq \frac{d_0 - \mathcal{D}^{\pi_B}(x_0)}{2\bar{T}^2 D_{\max}}, \quad \forall x \in \mathcal{X}'.$$

Recall that $\epsilon^*(x) = 2\bar{T}D_{\max}D_{TV}(\pi^*||\pi_B)(x)$. The above expression implies that

$$\mathbb{E} \left[\sum_{t=0}^{T^*-1} \epsilon^*(x_t) \mid \pi_B, x_0 \right] \leq 2\bar{T}^2 D_{\max} \max_{x \in \mathcal{X}'} D_{TV}(\pi^*||\pi_B)(x) \leq d_0 - \mathcal{D}^{\pi_B}(x_0),$$

which further implies

$$L_{\epsilon^*}(x_0) = \mathbb{E} \left[\sum_{t=0}^{T^*-1} \epsilon^*(x_t) | \pi_B, x_0 \right] + \mathcal{D}^{\pi_B}(x_0) \leq d_0,$$

i.e., the second property in (2) holds.

Second, recall the following equality from (12):

$$(I - P^*)(I - P_B)^{-1}(\{d(x)\}_{x \in \mathcal{X}'} + \{\epsilon(x)\}_{x \in \mathcal{X}'}) = \{d(x)\}_{x \in \mathcal{X}'}$$

with the definition of the auxiliary constraint cost ϵ given by (13). We want to show that the first condition in (2) holds. By adding the term $(I - P^*)(I - P_B)^{-1}\epsilon^*$ to both sides of the above equality, it implies that:

$$(I - P^*)(I - P_B)^{-1}(D + \epsilon^*) = D + (I - P^*)(I - P_B)^{-1}(\epsilon^* - \epsilon),$$

where $\epsilon^*(x) - \epsilon(x) \geq 0$, for $x \in \mathcal{X}'$. Therefore, the proof is completed if we can show that for any $x \in \mathcal{X}'$:

$$\{(I - P^*)(I - P_B)^{-1}(\epsilon^* - \epsilon)\}(x) \geq 0. \quad (14)$$

Now consider the following inequalities derived from Assumption 1:

$$\max_x D_{TV}(\pi^* | \pi_B)(x) \leq \frac{1}{2\bar{T}} \frac{\bar{T}D_{\max} - \bar{\mathcal{D}}}{\bar{T}D_{\max} + \bar{\mathcal{D}}} \leq \frac{1}{2\bar{T}} \frac{\bar{T}D_{\max} - \bar{\mathcal{D}}^*}{\bar{T}D_{\max} + \bar{\mathcal{D}}^*},$$

the last inequality is due to the fact that $\bar{\mathcal{D}} \geq \bar{\mathcal{D}}^*$, where $\bar{\mathcal{D}}^* = \max_x \mathcal{D}_{\pi^*}(x)$ is the constraint upper-bound w.r.t. optimal policy π^* . Multiplying the ratio $D_{TV}(\pi^* | \pi_B)(x) / \max_x D_{TV}(\pi^* | \pi_B)(x) \geq 0$ on both sides of the above inequality, for each $x \in \mathcal{X}'$ one obtains the following inequality:

$$\begin{aligned} D_{TV}(\pi^* | \pi_B)(x) &\leq \frac{(\bar{T}D_{\max} - \bar{\mathcal{D}}^*)D_{TV}(\pi^* | \pi_B)(x)}{2\bar{T}(\bar{T}D_{\max} + \bar{\mathcal{D}}^*) \max_x D_{TV}(\pi^* | \pi_B)(x)} \\ &\leq \frac{\epsilon^*(x) - \epsilon(x)}{2\bar{T} \max_x \{\epsilon^*(x) - \epsilon(x)\}}, \end{aligned}$$

the last inequality holds due to the fact that for any $x \in \mathcal{X}'$,

$$2(\bar{T}D_{\max} - \bar{\mathcal{D}}^*)D_{TV}(\pi^* | \pi_B)(x) \leq \epsilon^*(x) - \epsilon(x) \leq 2(\bar{T}D_{\max} + \bar{\mathcal{D}}^*)D_{TV}(\pi^* | \pi_B)(x).$$

Multiplying $2\bar{T} \max_x \{\epsilon^*(x) - \epsilon(x)\}$ on both sides, it further implies that for each $x \in \mathcal{X}'$, one has the following inequality:

$$(\epsilon^*(x) - \epsilon(x)) - 2\bar{T} \max_{x'} \{\epsilon^*(x') - \epsilon(x')\} D_{TV}(\pi^* | \pi_B)(x) \geq 0.$$

Now recall that $P^* = P_B - P_\Delta$, where Δ is equal to the matrix that characterizes the difference between the baseline and the optimal policy for each state in \mathcal{X}' and action in \mathcal{A} , i.e., $\Delta(a|x) = \pi_B(a|x) - \pi^*(a|x)$, $\forall x \in \mathcal{X}'$, $\forall a \in \mathcal{A}$ and $P_\Delta = \{\sum_{a \in \mathcal{A}} P(x'|x, a)\Delta(a|x)\}_{x, x' \in \mathcal{X}'}$, the above condition guarantees that

$$\{(I - P_B + P_\Delta)(I - P_B)^{-1}(\epsilon^* - \epsilon)\}(x) \geq 0, \quad \forall x \in \mathcal{X}'.$$

This finally comes to the conclusion that under Condition 1, the inequality in (14) holds, which further implies that

$$(I - P^*)(I - P_B)^{-1}(d(x) + \epsilon^*(x)) \geq d(x), \quad \forall x \in \mathcal{X},$$

i.e., the first property in (2) holds with $L_{\epsilon^*}(x) = \mathbb{E} \left[\sum_{t=0}^{T^*-1} d(x) + \epsilon^*(x) | \pi_B, x \right]$.

By combining the above results, one shows that L_{ϵ^*} is a Lyapunov function that satisfies the properties in (2) and (3), which concludes the proof.

C.3 Properties of Safe Bellman Operator

Proposition 3. *The safe Bellman operator has the following properties.*

- **Contraction:** *There exists a vector with positive components, i.e., $\rho : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$, and a discounting factor $0 < \gamma < 1$ such that*

$$\|T[V] - T[W]\|_{\rho} \leq \gamma \|V - W\|_{\rho},$$

where the weighted norm is defined as $\|V\|_{\rho} = \max_{x \in \mathcal{X}} \frac{V(x)}{\rho(x)}$.

- **Monotonicity:** *For any value functions $V, W : \mathcal{X} \rightarrow \mathbb{R}$ such that $V(x) \leq W(x)$, one has the following inequality: $T[V](x) \leq T[W](x)$, for any state $x \in \mathcal{X}$.*

Proof. First, we show the monotonicity property. For the case of $x \in \mathcal{X} \setminus \mathcal{X}'$, the property trivially holds. For the case of $x \in \mathcal{X}'$, given value functions $W, V : \mathcal{X}' \rightarrow \mathbb{R}$ such that $V(x) \leq W(x)$ for any $x \in \mathcal{X}'$, by the definition of Bellman operator T , one can show that for any $x \in \mathcal{X}'$ and any $a \in \mathcal{A}$,

$$c(x, a) + \sum_{x' \in \mathcal{X}'} P(x'|x, a)V(x') \leq c(x, a) + \sum_{x' \in \mathcal{X}'} P(x'|x, a)W(x').$$

Therefore, by multiplying $\pi(a|x)$ on both sides, summing the above expression over $a \in \mathcal{A}$, and taking the minimum of π over the feasible set $\mathcal{F}_{L_{\epsilon^*}}(x)$, one can show that $T[V](x) \leq T[W](x)$ for any $x \in \mathcal{X}'$.

Second we show that the contraction property holds. For the case of $x \in \mathcal{X} \setminus \mathcal{X}'$, the property trivially holds. For the case of $x \in \mathcal{X}'$, following the construction in Proposition 3.3.1 of [7], consider a stochastic shortest path problem where the transition probabilities and the constraint cost function are the same as the one in problem \mathcal{OPT} , but the cost are all equal to -1 . Then, there exists a fixed point value function \hat{V} , such that

$$\hat{V}(x) = -1 + \min_{\pi \in \mathcal{F}_{L_{\epsilon^*}}(x)} \sum_a \pi(a|x) \sum_{x' \in \mathcal{X}'} P(x'|x, a)\hat{V}(x'), \quad \forall x \in \mathcal{X}',$$

such that the following inequality holds for given feasible Markovian policy π' :

$$\hat{V}(x) \leq -1 + \sum_a \pi'(a|x) \sum_{x' \in \mathcal{X}'} P(x'|x, a)\hat{V}(x'), \quad \forall x \in \mathcal{X}'.$$

Notice that $\hat{V}(x) \leq -1$ for all $x \in \mathcal{X}'$. By defining $\rho(x) = -\hat{V}(x)$, and by constructing $\gamma = \max_{x \in \mathcal{X}'} (\rho(x) - 1)/\rho(x)$, one immediately has $0 < \gamma < 1$, and

$$\sum_a \pi'(a|x) \sum_{x' \in \mathcal{X}'} P(x'|x, a)\rho(x') \leq \rho(x) - 1 \leq \gamma\rho(x), \quad \forall x \in \mathcal{X}'.$$

Then by using Proposition 1.5.2 of [7], one can show that T is a contraction operator. \square

C.4 Proof of Theorem 2

Let $V_{\mathcal{OPT}}(x_0)$ be the optimal value function of problem \mathcal{OPT} , and let V^* be a fixed point solution: $V(x) = T[V](x)$, for any $x \in \mathcal{X}$. For the case when $x_0 \in \mathcal{X} \setminus \mathcal{X}'$, the following result trivially holds: $V_{\mathcal{OPT}}(x_0) = T[V_{\mathcal{OPT}}](x_0) = V^*(x_0) = 0$. Below, we show the equality for the case of $x_0 \in \mathcal{X}'$.

First, we want to show that $V_{\mathcal{OPT}}(x_0) \leq V^*(x_0)$. Consider the greedy policy $\bar{\pi}^*$ constructed from the fixed point equation. Immediately, one has that $\bar{\pi}^*(\cdot|x) \in \mathcal{F}_{L_{\epsilon^*}}(x)$. This implies

$$T_{\bar{\pi}^*, d}[L_{\epsilon^*}](x) \leq L_{\epsilon^*}(x), \quad \forall x \in \mathcal{X}'. \quad (15)$$

Thus by recursively applying $T_{\bar{\pi}^*, d}$ on both sides of the above inequality, the contraction property of Bellman operator $T_{\bar{\pi}^*, d}$ implies that one has the following expression:

$$\lim_{n \rightarrow \infty} T_{\bar{\pi}^*, d}^n[L_{\epsilon^*}](x_0) = \mathbb{E} \left[\sum_{t=0}^{T^*-1} d(x_t) + L_{\epsilon^*}(x_{T^*}) \mid x_0, \bar{\pi}^* \right] \leq L_{\epsilon^*}(x_0) \leq d_0.$$

Since the state enters the terminal set at time $t = T^*$, we have that $L_{\epsilon^*}(x_{T^*}) = 0$ almost surely. Then the above inequality implies $\mathbb{E} \left[\sum_{t=0}^{T^*-1} d(x_t) \mid x_0, \bar{\pi}^* \right] \leq d_0$, which further shows that $\bar{\pi}^*$ is a feasible policy to problem \mathcal{OPT} . On the other hand, recall that $V^*(x)$ is a fixed point solution to $V(x) = T[V](x)$, for any $x \in \mathcal{X}'$. Then for any bounded initial value function V_0 , the contraction property of Bellman operator $T_{\bar{\pi}^*, c}$ implies that

$$V^*(x) = \lim_{n \rightarrow \infty} T_{\bar{\pi}^*, c}^n[V_0](x) = \lim_{n \rightarrow \infty} \mathbb{E} \left[\sum_{t=0}^{n-1} c(x_t, a_t) + V_0(x_n) \mid x_0 = x, \bar{\pi}^* \right],$$

for which the transient assumption of stopping MDPs further implies that

$$V^*(x) = \mathbb{E} \left[\sum_{t=0}^{T^*-1} c(x_t, a_t) \mid x_0 = x, \bar{\pi}^* \right].$$

Since $\bar{\pi}^*$ is a feasible solution to problem \mathcal{OPT} . This further implies that $V_{\mathcal{OPT}}(x_0) \leq V^*(x_0)$.

Second, we want to show that $V_{\mathcal{OPT}}(x_0) \geq V^*(x_0)$. Consider the optimal policy π^* of problem \mathcal{OPT} that is used to construct Lyapunov function L_{ϵ^*} . Since the Lyapunov function satisfies the following Bellman inequality:

$$T_{\pi^*, d}[L_{\epsilon^*}](x) \leq L_{\epsilon^*}(x), \quad \forall x \in \mathcal{X}'$$

it implies that the optimal policy π^* is a feasible solution to the optimization problem in Bellman operator $T[V^*](x)$. Note that V^* is a fixed point solution to equation: $V^*(x) = T[V^*](x)$, for any $x \in \mathcal{X}'$. Immediately the above result yields the following inequality:

$$V^*(x) = T_{\bar{\pi}^*, c}[V^*](x) \leq T_{\pi^*, c}[V^*](x), \quad \forall x \in \mathcal{X}'$$

the first equality holds because $\bar{\pi}^*(\cdot|x)$ is the minimizer of the optimization problem in $T[V^*](x)$, $x \in \mathcal{X}'$. By recursively applying Bellman operator $T_{\bar{\pi}^*, c}$ to this inequality, one has the following result:

$$V^*(x) \leq \lim_{n \rightarrow \infty} T_{\bar{\pi}^*, c}^n[V^*](x) = \mathcal{C}_{\pi^*}(x) = V_{\mathcal{OPT}}(x), \quad \forall x \in \mathcal{X}'.$$

One thus concludes that $V_{\mathcal{OPT}}(x_0) \geq V^*(x_0)$.

Combining the above analysis, we prove the claim of $V_{\mathcal{OPT}}(x_0) = V^*(x_0)$, and the greedy policy of the fixed-point equation, i.e., $\bar{\pi}^*$, is an optimal policy to problem \mathcal{OPT} .

C.5 Proof of Proposition 1

For the derivations of consistent feasibility and policy improvement, without loss of generality we only consider the case of $k = 0$.

To show the property of consistent feasibility, consider an arbitrary feasible policy π_0 of problem \mathcal{OPT} . By definition, one has $\mathcal{D}_{\pi_0}(x_0) \leq d_0$, and the value function \mathcal{D}_{π_0} has the following property:

$$\sum_a \pi_0(a|x) \left[d(x) + \sum_{x'} P(x'|x, a) \mathcal{D}_{\pi_0}(x') \right] = \mathcal{D}_{\pi_0}(x), \quad \forall x \in \mathcal{X}'.$$

Immediately, since \mathcal{D}_{π_0} satisfies the constraint in (5), one can treat it as a Lyapunov function, this shows that the set of Lyapunov functions $\mathcal{L}_{\pi_0}(x_0, d_0)$ is non-empty. Therefore, there exists a bounded Lyapunov function $\{L_{\epsilon_0}(x)\}_{x \in \mathcal{X}}$ as the solution of the optimization problem in Step 0. Now consider the policy optimization problem in Step 1. Based on the construction of $\{L_{\epsilon_0}(x)\}_{x \in \mathcal{X}}$, the current policy π_0 is a feasible solution to this problem, therefore the feasibility set is non-empty. Furthermore, by recursively applying the inequality constraint on the updated policy π_1 for $T^* - 1$ times, one has the following inequality:

$$\mathbb{E} \left[\sum_{t=0}^{T^*-1} d(x_t) + L_{\epsilon_0}(x_{T^*}) \mid x_0, \pi_1 \right] \leq L_{\epsilon_0}(x_0) \leq d_0.$$

This shows that π_1 is a feasible policy to problem \mathcal{OPT} .

To show the property of policy improvement, consider the policy optimization in Step 1. Notice that the current policy π_0 is a feasible solution of this problem (with Lyapunov function L_0), and the updated policy π_1 is a minimizer of this problem. Then, one immediately has the following chain of inequalities:

$$\begin{aligned} T_{\pi_1, c}[V_0](x) &= \sum_{a \in \mathcal{A}} \pi_1(a|x) \left[c(x, a) + \sum_{x' \in \mathcal{X}'} P(x'|x, a) V_0(x') \right] \\ &\leq \sum_{a \in \mathcal{A}} \pi_0(a|x) \left[c(x, a) + \sum_{x' \in \mathcal{X}'} P(x'|x, a) V_0(x') \right] = V_0(x), \quad \forall x \in \mathcal{X}', \end{aligned}$$

where the last equality is due to the fact that $V_0(x) = \mathcal{C}_{\pi_0}(x)$, for any $x \in \mathcal{X}$. By the contraction property of Bellman operator T_{π_1} , the above condition further implies

$$\mathcal{C}_{\pi_1}(x) = \lim_{n \rightarrow \infty} T_{\pi_1, c}^n[V_0](x) \leq V_0(x) = \mathcal{C}_{\pi_0}(x), \quad \forall x \in \mathcal{X}',$$

which proves the claim about policy improvement.

To show the property of asymptotic convergence, notice that the value function sequence $\{\mathcal{C}_{\pi_k}(\cdot)\}_{k \geq 0}$ is uniformly monotonic, and each element is uniformly lower bounded by the unique solution of fixed point equation: $V(x) = \min_{a \in \mathcal{A}} c(x, a) + \sum_{x' \in \mathcal{X}'} P(x'|x, a) V(x')$, $\forall x \in \mathcal{X}'$. Therefore, this sequence of value function converges (point-wise) as soon as in the limit the policy improvement stops. Whenever this happens, i.e., there exists $K \geq 0$ such that $\max_{x \in \mathcal{X}'} |\mathcal{C}_{\pi_{K+1}}(x) - \mathcal{C}_{\pi_K}(x)| \leq \epsilon$ for any $\epsilon > 0$, then this value function is the fixed point of $\min_{\pi \in \mathcal{F}_{L_K}(x)} T_{\pi, c}[V](x) = V(x)$, $\forall x \in \mathcal{X}'$, whose solution policy is unique (due to the strict convexity of the objective function in the policy optimization problem after adding a convex regularizer). Furthermore, due to the strict concavity of the objective function in problem in (5) (after adding a concave regularizer), the solution pair of this problem is unique, which means the update of $\{(L_{\epsilon_k}, \epsilon_k)\}$ stops at step K . Together, this also means that policy update $\{\pi_k\}$ converges.

C.6 Analysis on Performance Improvement in Safe Policy Iteration

Similar to the analysis in [2], the following lemma provides a bound in policy improvement.

Lemma 2. *For any policies π' and π , define the following error function:*

$$\begin{aligned} \Lambda(\pi, \pi') &= \mathbb{E} \left[\sum_{t=0}^{T^*-1} \left(\frac{\pi'(a_t|x_t)}{\pi(a_t|x_t)} - 1 \right) \underbrace{(C(x_t, a_t) + V^\pi(x_{t+1}) - V^\pi(x_t))}_{\delta^\pi(x_t, a_t, x_{t+1})} \mid x_0, \pi \right] \\ &= \mathbb{E} \left[\sum_{t=0}^{T^*-1} \mathbb{E}_{a \sim \pi'(\cdot|x)} [Q^\pi(x_t, a)] - V^\pi(x_t) \mid x_0, \pi \right] \end{aligned}$$

and $\Delta^{\pi'} = \max_{x \in \mathcal{X}'} |\mathbb{E}_{a \sim \pi'(\cdot|x)} [Q^\pi(x, a)] - V^\pi(x)|$. Then, the following error bound on the performance difference between π and π' holds:

$$\Lambda(\pi, \pi') - \mathcal{E}_{TV}^{\pi, \pi'} \leq \mathcal{C}_{\pi'}(x_0) - \mathcal{C}_\pi(x_0) \leq \Lambda(\pi, \pi') + \mathcal{E}_{TV}^{\pi, \pi'}.$$

where $\mathcal{E}_{TV}^{\pi, \pi'} = 2\Delta^{\pi'} \cdot (\max_{x_0 \in \mathcal{X}'} \mathbb{E}[T^*|x_0, \pi']) \cdot \mathbb{E} \left[\sum_{t=0}^{T^*-1} D_{TV}(\pi'|\pi)(x_t) \mid x_0, \pi \right]$.

Proof. First, it is clear from the property of telescopic sum that

$$\begin{aligned} \mathcal{C}_{\pi'}(x_0) - \mathcal{C}_\pi(x_0) &= \mathbb{E} \left[\sum_{t=0}^{T^*-1} \delta^{\pi'}(x_t, a_t, x_{t+1}) \mid x_0, \pi' \right] - \mathbb{E} \left[\sum_{t=0}^{T^*-1} \delta^\pi(x_t, a_t, x_{t+1}) \mid x_0, \pi \right] \\ &\leq \mathbb{E} \left[\sum_{t=0}^{T^*-1} \delta^{\pi'}(x_t, a_t, x_{t+1}) - \delta^\pi(x_t, a_t, x_{t+1}) \mid x_0, \pi \right] + \Delta^{\pi'} \cdot \sum_{y \in \mathcal{X}'} \left| \sum_{t=0}^{T^*-1} \mathbb{P}(x_t = y | x_0, \pi') - \sum_{t=0}^{T^*-1} \mathbb{P}(x_t = y | x_0, \pi) \right|, \end{aligned}$$

where the inequality is based on the Holder inequality $\mathbb{E}[|xy|] \leq (\mathbb{E}[|x|^p])^{1/p} (\mathbb{E}[|y|^q])^{1/q}$ with $p = 1$ and $q = \infty$.

Immediately, the first part of the above expression is re-written as: $\mathbb{E} \left[\sum_{t=0}^{T^*-1} \delta \pi'(x_t, a_t, x_{t+1}) - \delta \pi(x_t, a_t, x_{t+1}) \mid x_0, \pi \right] = \Lambda(\pi, \pi')$. Recall that shorthand notation P_π to denote the transition probability for $x \in \mathcal{X}'$ induced by the policy π . For the second part of the above expression, notice that the following chain of inequalities holds:

$$\begin{aligned}
& \sum_{y \in \mathcal{X}'} \left| \sum_{t=0}^{T^*-1} \mathbb{P}(x_t = y \mid x_0, \pi) - \sum_{t=0}^{T^*-1} \mathbb{P}(x_t = y \mid x_0, \pi') \right| \\
&= \sum_{y \in \mathcal{X}'} \left| \mathbf{1}(x_0)^\top \left((I - P_\pi)^{-1} - (I - P_{\pi'})^{-1} \right) \mathbf{1}(y) \right| \\
&= \sum_{y \in \mathcal{X}'} \left| \mathbf{1}(x_0)^\top (I - P_\pi)^{-1} (P_\pi - P_{\pi'}) (I - P_{\pi'})^{-1} \mathbf{1}(y) \right| \\
&\leq \|(\mathbf{1}(x_0)^\top (I - P_\pi)^{-1} (P_\pi - P_{\pi'})\|_1 \cdot \sum_{y \in \mathcal{X}'} \max_{x_0 \in \mathcal{X}'} \mathbf{1}(x_0) (I - P_{\pi'})^{-1} \mathbf{1}(y) \\
&\leq \sum_{y \in \mathcal{X}'} \bar{T} \cdot \mathbf{1}(y) \cdot \mathbb{E} \left[\sum_{t=0}^{T^*-1} \sum_{x' \in \mathcal{X}'} P(x' \mid x_t, a_t) \sum_{a \in \mathcal{A}} |\pi(a_t \mid x_t) - \pi'(a_t \mid x_t)| \mid x_0, \pi \right] \\
&= 2\bar{T} \cdot \mathbb{E} \left[\sum_{t=0}^{T^*-1} D_{TV}(\pi' \mid \pi)(x_t) \mid x_0, \pi \right],
\end{aligned}$$

the first, is based on the Holder inequality with $p = 1$ and $q = \infty$ and on the fact that all entries in $(I - P_{\pi'})^{-1}$ is non-negative, the second inequality is due to the fact that starting at any initial state x_0 , it almost takes \bar{T} steps to the set of recurrent states $\mathcal{X} \setminus \mathcal{X}'$. In other words, one has the following inequality:

$$\mathbf{1}(x_0) (I - P_{\pi'})^{-1} \mathbf{1}(y) = \mathbb{E} \left[\sum_{t=0}^{T^*-1} \mathbf{1}\{x_t = y\} \mid x_0, \pi' \right] \leq \bar{T}, \quad \forall x_0, y \in \mathcal{X}'.$$

Therefore, combining with these properties the proof of the above error bound is completed. \square

Using this result, the sub-optimality performance bound of policy π_{k^*} from SPI is $\bar{T}C_{\max} - \left(\sum_{k=0}^{k^*-1} \max\{0, \Lambda^{\pi_k, \pi_{k+1}} - \mathcal{E}_{TV}^{\pi_k, \pi_{k+1}}\} + C_{\pi_0}(x_0) \right)$.

C.7 Proof of Proposition 2

For the derivations of consistent feasibility and monotonic improvement on value estimation, without loss of generality we only consider the case of $t = 0$.

To show the property of consistent feasibility, notice that with the definitions of the initial Q -function Q_0 , the initial Lyapunov function L_{ϵ_0} w.r.t. the initial auxiliary cost ϵ_0 , the corresponding induced policy π_0 is feasible to problem \mathcal{OPT} , i.e., $\mathcal{D}_{\pi_0}(x_0) \leq d_0$. Consider the optimization problem in Step 1. Immediately, since \mathcal{D}_{π_0} satisfies the constraint in (5), one can treat it as a Lyapunov function, and the set of Lyapunov functions $\mathcal{L}_{\pi_0}(x_0, d_0)$ is non-empty. Therefore, there exists a bounded Lyapunov function $\{L_{\epsilon_1}(x)\}_{x \in \mathcal{X}}$ and auxiliary cost $\{\epsilon_1(x)\}_{x \in \mathcal{X}}$ as the solution of this optimization problem. Now consider the policy update in Step 0. Since $\pi_1(\cdot \mid \cdot)$ belongs to the set of feasible policies $\mathcal{F}_{L_{\epsilon_1}}(\cdot)$, by recursively applying the inequality constraint on the updated policy π_1 for $T^* - 1$ times, one has the following inequality:

$$\mathbb{E} \left[\sum_{t=0}^{T^*-1} d(x_t) + L_{\epsilon_1}(x_{T^*}) \mid x_0, \pi_1 \right] \leq L_{\epsilon_1}(x_0) \leq d_0.$$

This shows that π_1 is a feasible policy to problem \mathcal{OPT} .

To show the asymptotic convergence property, for every initial state x_0 and any time step K , with the policy $\pi = \{\pi_0, \dots, \pi_{K-1}\}$ generated by the value iteration procedure, the cumulative cost can be broken down into the following two portions, which consists of the cost over the first K stages and the remaining cost. Specifically,

$$\bar{V}(x_0) = \mathbb{E} \left[\sum_{t=0}^{T^*-1} c(x_t, a_t) \mid x_0, \pi \right] = \mathbb{E} \left[\sum_{t=0}^{K-1} c(x_t, a_t) \mid x_0, \pi \right] + \mathbb{E} \left[\sum_{t=K}^{T^*-1} c(x_t, a_t) \mid x_0, \pi \right],$$

where the second term is bounded $\mathbb{E}[T^* - K \mid \pi, x_0] C_{\max}$, which is bounded by $\sum_{t=K}^{\infty} \mathbb{P}(x_t \in \mathcal{X}' \mid x_0, \pi) \cdot C_{\max} > 0$. Since the value function $V_0(x) = \min_{\pi \in \mathcal{F}_{L\epsilon_0}(x)} \pi(\cdot \mid x)^\top Q_0(x, \cdot)$ is also bounded, one can further show the following inequality:

$$\begin{aligned} -\mathbb{P}(x_K \in \mathcal{X}' \mid x_0, \pi) \cdot \|V_0\|_\infty &\leq T_{K-1}[\dots [T_0[V_0]] \dots](x_0) - \mathbb{E} \left[\sum_{t=0}^{K-1} c(x_t, a_t) \mid x_0, \pi \right] \\ &\leq \mathbb{P}(x_K \in \mathcal{X}' \mid x_0, \pi) \cdot \|V_0\|_\infty. \end{aligned}$$

Recall from our problem setting that all policies are proper (see Assumption 3.1.1 and Assumption 3.1.2 in [7]). Then by the property of a transient MDP (see Definition 7.1 in [3]), the sum of probabilities of the state trajectory after step K that is in the transient set \mathcal{X}' , i.e., $\sum_{t=K}^{\infty} \mathbb{P}(x_t \in \mathcal{X}' \mid x_0, \pi)$, is bounded by $M_\pi \cdot \epsilon$. Therefore, as K goes to ∞ , ϵ approaches 0. Using the result that $\sum_{t=K}^{\infty} \mathbb{P}(x_t \in \mathcal{X}' \mid x_0, \pi)$ vanishes as K goes to ∞ , one concludes that

$$\lim_{K \rightarrow \infty} T_{K-1}[\dots [T_0[V_0]] \dots](x_0) = \bar{V}(x_0), \quad \forall x_0 \in \mathcal{X}',$$

which completes the proof of this property.

D Pseudo-code of Safe Reinforcement Learning Algorithms

Algorithm 3 Policy Distillation

Input: Policy parameterization π_ϕ with parameter ϕ ; A batch of state trajectories $\{x_{0,m}, \dots, x_{\bar{T}-1,m}\}_{m=1}^M$ generated by following the baseline policy π_B
 Compute the action probabilities $\{\pi'(\cdot|x_{0,m}), \dots, \pi'(\cdot|x_{\bar{T}-1,m})\}_{m=1}^M$ by solving problem in (16).
 Compute the policy parameter by supervised learning:

$$\phi^* \in \arg \min_{\phi} \frac{1}{m} \sum_{m=1}^M \sum_{t=0}^{\bar{T}-1} D_{\text{JSD}}(\pi_\phi(\cdot|x_{t,m}) \parallel \pi'(\cdot|x_{t,m}))$$

where $D_{\text{JSD}}(P||Q) = \frac{1}{2}D(P \parallel \frac{1}{2}(P + Q)) + \frac{1}{2}D(Q \parallel \frac{1}{2}(P + Q))$ is the Jensen-Shannon divergence

Return Distilled policy π_{ϕ^*}

Algorithm 4 Safe DQN

Input: Initial prioritized replay buffer $M = \{\emptyset\}$; Initial importance weights $w_0 = 1$, $w_{D,0} = 1$, $w_{T,0} = 1$; Mini-batch size $|B|$; Network parameters $\theta^-, \theta_D^-, \theta_T^-$; Initial feasible policy π_0 ;
for $k \in \{0, 1, \dots\}$ **do**
 for $t = 0$ **to** $\bar{T} - 1$ **do**
 Sample action $a_t \sim \pi_k(\cdot|x_t)$, execute a_t and observe next state x_{t+1} , and costs (c_t, d_t)
 Add experiences to replay buffer, i.e., $M \leftarrow (x_t, a_t, c_t, d_t, x_{t+1}, w_0, w_{D,0}, w_{T,0}) \cup M$
 From the buffer M , sample a mini-batch $B = \{(x_j, a_j, c_j, d_j, x'_j, w_j, w_{D,j}, w_{T,j})\}_{j=1}^{|B|}$ and set the targets $y_{D,j}$, $y_{T,j}$, and y_j
 Update the θ parameters as follows:

$$\begin{aligned} \theta_D &\leftarrow \theta_D^- - \kappa_j \cdot w_{D,j} \cdot (y_{D,j} - \widehat{Q}_D(x_j, a_j; \theta_D^-)) \cdot \nabla_{\theta} \widehat{Q}_D(x_j, a_j; \theta) |_{\theta=\theta_D^-}, \\ \theta_T &\leftarrow \theta_T^- - \kappa_j \cdot w_{T,j} \cdot (y_{T,j} - \widehat{Q}_T(x_j, a_j; \theta_T^-)) \cdot \nabla_{\theta} \widehat{Q}_T(x_j, a_j; \theta) |_{\theta=\theta_T^-}, \\ \theta &\leftarrow \theta^- - \eta_j \cdot w_j \cdot (y_j - \widehat{Q}(x_j, a_j; \theta^-)) \cdot \nabla_{\theta} \widehat{Q}(x_j, a_j; \theta) |_{\theta=\theta^-} \end{aligned}$$

where the target values are respectively $y_{D,j} = d(x_j) + \pi_k(\cdot|x'_j)^\top \widehat{Q}_D(x'_j, \cdot; \theta_D)$, $y_{T,j} = \mathbf{1}\{x_j \in \mathcal{X}'\} + \pi_k(\cdot|x'_j)^\top \widehat{Q}_T(x'_j, \cdot; \theta_T)$, and $y_j = c(x_j, a_j) + \pi'(\cdot|x'_j)^\top \widehat{Q}(x'_j, \cdot; \theta)$, and $\pi'(\cdot|x'_j)$ is the greedy action probability w.r.t. x'_j , which is a solution to (6)

Prioritized Sweep: Update importance weights w_j , $w_{D,j}$, and $w_{T,j}$ of the samples in mini-batch B , based on TD errors $|y_j - \widehat{Q}(x_j, a_j; \theta)|$, $|y_{D,j} - \widehat{Q}_D(x_j, a_j; \theta_D)|$ and $|y_{T,j} - \widehat{Q}_T(x_j, a_j; \theta_T)|$

Distillation: Update the policy to π_{k+1} using Algorithm 3 w.r.t. data $\{x'_{0,j}, \dots, x'_{\bar{T}-1,j}\}_{j=1}^{|B|}$ and $\{\pi'(\cdot|x'_{0,j}), \dots, \pi'(\cdot|x'_{\bar{T}-1,j})\}_{j=1}^{|B|}$

end for

Double Q-learning: Update $\theta^- = \theta$, $\theta_D^- = \theta_D$ and $\theta_T^- = \theta_T$ after c iterations

end for

Return

Algorithm 5 Safe DPI

Input: Initial prioritized replay buffer $M = \{\emptyset\}$; Initial importance weights $w_0 = 1$, $w_{D,0} = 1$, $w_{T,0} = 1$; Mini-batch size $|B|$; Initial feasible policy π_0 ;

for $k \in \{0, 1, \dots\}$ **do**

Sample action $a_t \sim \pi_k(\cdot|x_t)$, execute a_t and observe next state x_{t+1} , and costs (c_t, d_t)

Add experiences to replay buffer, i.e., $M \leftarrow (x_t, a_t, c_t, d_t, x_{t+1}, w_0, w_{D,0}, w_{T,0}) \cup M$

From the buffer M , sample a mini-batch $B = \{(x_j, a_j, c_j, d_j, x'_j, w_j, w_{D,j}, w_{T,j})\}_{j=1}^{|B|}$ and set the targets $y_{D,j}$, $y_{T,j}$, and y_j

Update the θ parameters as follows:

$$\theta_{\pi_k}^* \in \arg \min_{\theta} \sum_{t,j} (c_{t,j} - \widehat{Q}(x_{t,j}, a_{t,j}; \theta) + \pi_k(\cdot|x_{t+1,j})^\top \widehat{Q}(x_{t+1,j}, \cdot; \theta))^2,$$

$$\theta_{D, \pi_k}^* \in \arg \min_{\theta_D} \sum_{t,j} (d_{t,j} - \widehat{Q}_D(x_{t,j}, a_{t,j}; \theta_D) + \pi_k(\cdot|x_{t+1,j})^\top \widehat{Q}_D(x_{t+1,j}, \cdot; \theta_D))^2,$$

$$\theta_{T, \pi_k}^* \in \arg \min_{\theta_T} \sum_{t,j} (\mathbf{1}\{x_{t,j} \in \mathcal{X}'\} - Q_T(x_{t,j}, a_{t,j}; \theta_T) + \pi_k(\cdot|x_{t+1,j})^\top Q_T(x_{t+1,j}, \cdot; \theta_T))^2$$

and construct the Q -functions $\widehat{Q}(x, a; \theta_{\pi_k}^*)$ and $\widehat{Q}_L(x, a; \theta_{D, \pi_k}^*, \theta_{T, \pi_k}^*) = \widehat{Q}_D(x, a; \theta_D) + \epsilon' \cdot \widehat{Q}_T(x, a; \theta_T)$

Calculate greedy action probabilities $\{\pi'(\cdot|x_{0,j}), \dots, \pi'(\cdot|x_{\overline{T}-1,j})\}_{j=1}^{|B|}$ by solving problem (6), with respect to batch of states drawn from the replay buffer $\{x_{0,j}, \dots, x_{\overline{T}-1,j}\}_{j=1}^{|B|}$

Distillation: Update the policy to π_{k+1} using Algorithm 3 w.r.t. data $\{x_{0,j}, \dots, x_{\overline{T}-1,j}\}_{j=1}^{|B|}$ and $\{\pi'(\cdot|x_{0,j}), \dots, \pi'(\cdot|x_{\overline{T}-1,j})\}_{j=1}^{|B|}$

end for

Return Final policy π_{k^*}

E Practical Implementations

There are several techniques that improve training and scalability of the safe RL algorithms. To improve stability in training Q networks, one may apply double Q -learning [42] to separate the target values and the value function parameters and to slowly update the target Q values at every predetermined iterations. On the other hand, to incentivize learning at state-action pairs that have high temporal difference (TD) errors, one can use a prioritized sweep in replay buffers [37] to add an importance weight to relevant experience. To extend the safe RL algorithms to handle continuous actions, one may adopt the normalized advantage functions (NAFs) [17] parameterization for Q -functions. Finally, instead of exactly solving the LP problem for policy optimization in (6), one may approximate this solution by solving its entropy regularized counterpart [29]. This approximation has an elegant closed-form solution that is parameterized by a Lagrange multiplier, which can be effectively computed by binary search methods (see Appendix E.1 and Appendix E.2 for details).

E.1 Case 1: Discrete Action Space

In this case, problem (6) is cast as finite dimensional linear programming (LP). In order to effectively approximate the solution policy especially when the action space becomes large, instead of exactly solving this inner optimization problem, one considers its Shannon entropy-regularized variant:

$$\begin{aligned} \min_{\pi \in \Delta} & \pi(\cdot|x)^\top (Q(x, \cdot) + \tau \log \pi(\cdot|x)) \\ \text{s.t.} & (\pi(\cdot|x) - \pi_B(\cdot|x))^\top Q_L(x, \cdot) \leq \tilde{\epsilon}'(x) \end{aligned} \quad (16)$$

where $\tau > 0$ is the regularization constant. When $\tau \rightarrow 0$, then π_τ^* converges to the original solution policy π^* .

We will hereby illustrate how to effectively solve π_τ^* for any given $\tau > 0$ without explicitly solving the LP. Consider the Lagrangian problem for entropy-regularized optimization:

$$\min_{\pi \in \Delta} \max_{\lambda \geq 0} \Gamma_x(\pi, \lambda),$$

where $\Gamma_x(\pi, \lambda) = \pi(\cdot|x)^\top(Q(x, \cdot) + \lambda Q_L(x, \cdot) + \tau \log \pi(\cdot|x)) - \lambda(\pi_B(\cdot|x)^\top Q_L(x, \cdot) + \tilde{\epsilon}(x))$ is the Lagrangian function. Notice that the set of stationary Markovian policies Δ is a convex set, and the objective function is a convex function in π and concave in λ . By strong duality, there exists a saddle-point to the Lagrangian problem where solution policy is equal to π_τ^* , and it can be solved by the maximin problem:

$$\max_{\lambda \geq 0} \min_{\pi \in \Delta} \Gamma_x(\pi, \lambda).$$

For the inner minimization problem, it has been shown that the λ -solution policy has the following closed form:

$$\pi_{\tau, \lambda}^*(\cdot|x) \propto \exp\left(-\frac{Q(x, \cdot) + \lambda Q_L(x, \cdot)}{\tau}\right).$$

Equipped with this formulation, we now solve the problem for the optimal Lagrange multiplier $\lambda^*(x)$ at state $x \in \mathcal{X}'$:

$$\max_{\lambda \geq 0} -\tau \cdot \text{logsumexp}\left(-\frac{Q(x, \cdot) + \lambda Q_L(x, \cdot)}{\tau}\right) - \lambda(\pi_B(\cdot|x)^\top Q_L(x, \cdot) + \tilde{\epsilon}(x)),$$

where $\text{logsumexp}(y) = \log \sum_a \exp(y_a)$ is a strictly convex function in y , and the objective function is a concave function of λ . Notice that this problem has a unique optimal Lagrange multiplier that is the solution of the following KKT condition:

$$\pi_B(\cdot|x)^\top Q_L(x, \cdot) + \tilde{\epsilon}(x) + \frac{\sum_a Q_L(x, a) \exp\left(-\frac{Q(x, a) + \lambda Q_L(x, a)}{\tau}\right)}{\sum_a \exp\left(-\frac{Q(x, a) + \lambda Q_L(x, a)}{\tau}\right)} = 0.$$

Using the parameterization $z = \exp(-\lambda)$, this condition can be written as the following polynomial equation in z :

$$\sum_a \left(Q_L(x, a) + \pi_B(\cdot|x)^\top Q_L(x, \cdot) + \tilde{\epsilon}(x)\right) \cdot \exp\left(-\frac{Q(x, a)}{\tau}\right) \cdot z^{\frac{Q_L(x, a)}{\tau}} = 0. \quad (17)$$

Therefore, the solution $0 \leq z^*(x) \leq 1$ can be solved by computing the root solution of the above polynomial and the optimal Lagrange multiplier is given by $\lambda^*(x) = -\log(z^*(x)) \geq 0$.

Combining the above results, the optimal policy of the entropy-regularized problem is therefore given by

$$\pi_\tau^*(\cdot|x) \propto \exp\left(-\frac{Q(x, \cdot) + \lambda^*(x) \cdot Q_L(x, \cdot)}{\tau}\right). \quad (18)$$

E.2 Case 2: Continuous Action Space

In order to effectively solve the inner optimization problem in (6) when the action space is continuous, on top of the using the entropy-regularized inner optimization problem in (16), we adopt the idea from the normalized advantage functions (NAF) approach for function approximation, where we express the Q -function and the state-action Lyapunov function with their second order Taylor-series expansions at an arbitrary action $\mu(x)$ as follows:

$$\begin{aligned} Q(x, a) &\approx Q(x, \mu(x)) + \nabla_a Q(x, a)|_{a=\mu(x)} \cdot (a - \mu(x)) \\ &\quad + \frac{1}{2}(a - \mu(x))^\top \nabla_a^2 Q(x, a)|_{a=\mu(x)} (a - \mu(x)) + o(\|a - \mu(x)\|^3), \\ Q_L(x, a) &\approx Q_L(x, \mu(x)) + \nabla_a Q_L(x, a)|_{a=\mu(x)} \cdot (a - \mu(x)) \\ &\quad + \frac{1}{2}(a - \mu(x))^\top \nabla_a^2 Q_L(x, a)|_{a=\mu(x)} (a - \mu(x)) + o(\|a - \mu(x)\|^3). \end{aligned}$$

While these representations are more restrictive than the general function approximations, they provide a recipe to determine the policy, which is a minimizer of problem (6) analytically for the updates in the safe AVI and safe DQN algorithms. In particular, using the above parameterizations, notice that

$$\begin{aligned} &\text{logsumexp}\left(-\frac{Q(x, \cdot) + \lambda^*(x) \cdot Q_L(x, \cdot)}{\tau}\right) \\ &= -\frac{1}{2} \log |A_{\lambda^*}(x)| - \frac{1}{2} \psi_{\lambda^*}(x)^\top A_{\lambda^*}^{-1}(x) \psi_{\lambda^*}(x) + K_{\lambda^*}(x) + \frac{n}{2} \log(2\pi), \end{aligned}$$

where n is the dimension of actions,

$$A_{\lambda^*}(x) = \frac{1}{\tau} \left(\nabla_a^2 Q(x, a)|_{a=\mu(x)} + \lambda^*(x) \nabla_a^2 Q_L(x, a)|_{a=\mu(x)} \right),$$

$$\psi_{\lambda^*}(x) = -\frac{1}{\tau} \left(\nabla_a Q(x, a)|_{a=\mu(x)} + \lambda^*(x) \nabla_a Q_L(x, a)|_{a=\mu(x)} \right) - A_{\lambda^*}(x) \mu(x),$$

and

$$K_{\lambda^*}(x) = -\frac{Q(x, \mu(x)) + \lambda^*(x) Q_L(x, \mu(x))}{\tau}$$

$$+ \left(\frac{1}{\tau} \left(\nabla_a Q(x, a)|_{a=\mu(x)} + \lambda^*(x) \nabla_a Q_L(x, a)|_{a=\mu(x)} \right)^\top - \frac{1}{2} \mu(x)^\top A_{\lambda^*}(x) \right) \mu(x)$$

is a normalizing constant (that is independent of a). Therefore, according to the closed-form solution of the policy in (18), the optimal policy of problem (16) follows a Gaussian distribution, which is given by

$$\pi_\tau^*(\cdot|x) \sim \mathcal{N}(A_{\lambda^*}(x)^{-1} \psi_{\lambda^*}(x), A_{\lambda^*}(x)^{-1}).$$

In order to completely characterize the solution policy, it is still required to compute the Lagrange multiplier $\lambda^*(x)$, which is a polynomial root solution of (17). Since the action space is continuous, one can only approximate the integral (over actions) in this expression with numerical integration techniques, such as Gaussian quadrature, Simpson's method, or Trapezoidal rule etc. (Notice that if π_B is a Gaussian policy, there is a tractable closed form expression for $\pi_B(\cdot|x)^\top Q_L(x, \cdot)$.)

F Experimental Setup

In the CMDP planning experiment, in order to demonstrate the numerical efficiency of the safe DP algorithms, we run a larger example that has a grid size of 60×60 . To compute the LP policy optimization step, we use the open-source SciPy linprog solver. In terms of the computation time, on average every policy optimization iteration (over all states) in SPI and SVI takes approximately 25.0 seconds, and for this problem SVI takes around 200 iterations to converge, while SPI takes 60 iterations. On the other hand the Dual LP method computes an optimal solution, its computation time is over 9500 seconds.

In the RL experiments, we use the Adam optimizer with learning rate 0.0001. At each iteration, we collect an episode of experience (100 steps) and perform 10 training steps on batches of size 128 sampled uniformly from the replay buffer. We update the target Q networks every 10 iterations and the baseline policy every 50 iterations.

For discrete observations, we use a feed-forward neural network with hidden layers of size 16, 64, 32, and relu activations.

For image observations, we use a convolutional neural network with filters of size $3 \times 3 \times 3 \times 32$, $32 \times 3 \times 3 \times 64$, and $64 \times 3 \times 3 \times 128$, with 2×2 max-pooling and relu activations after each. We then pass the result through a 2-hidden layer network with sizes 512 and 128.




Figure 3: Results of various planning algorithms on the grid-world environment with obstacles (zoomed), with x-axis showing the obstacle density. From the leftmost column, the first figure illustrates the 2D planning domain example. The second and the third figure show the average return and the average cumulative constraint cost of the CMDP methods respectively. The fourth figure displays all the methods used in the experiment. The shaded regions indicate the 80% confidence intervals. Clearly the safe DP algorithms compute policies that are safe and have good performance.

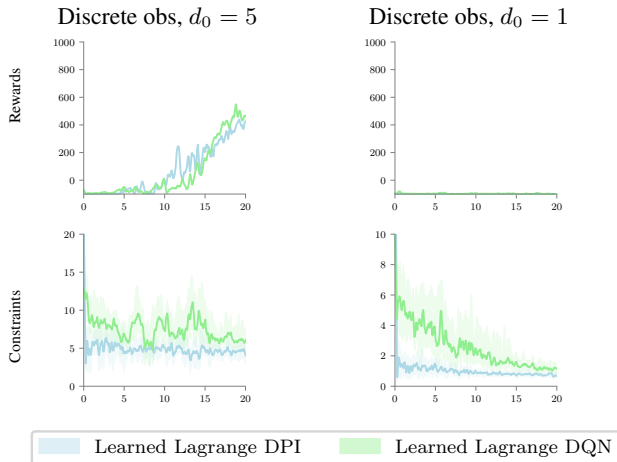


Figure 4: Results of using a saddle-point Lagrangian optimization for solving the grid-world environment with obstacles, with x-axis in thousands of episodes.