



Explicit Explore, Exploit, or Escape (E^4): near-optimal safety-constrained reinforcement learning in polynomial time

David M. Bossens¹ · Nicholas Bishop¹

Received: 17 November 2021 / Revised: 8 April 2022 / Accepted: 26 May 2022
© The Author(s) 2022

Abstract

In reinforcement learning (RL), an agent must explore an initially unknown environment in order to learn a desired behaviour. When RL agents are deployed in real world environments, safety is of primary concern. Constrained Markov decision processes (CMDPs) can provide long-term safety constraints; however, the agent may violate the constraints in an effort to explore its environment. This paper proposes a model-based RL algorithm called Explicit Explore, Exploit, or Escape (E^4), which extends the Explicit Explore or Exploit (E^3) algorithm to a robust CMDP setting. E^4 explicitly separates exploitation, exploration, and escape CMDPs, allowing targeted policies for policy improvement across known states, discovery of unknown states, as well as safe return to known states. E^4 robustly optimises these policies on the worst-case CMDP from a set of CMDP models consistent with the empirical observations of the deployment environment. Theoretical results show that E^4 finds a near-optimal constraint-satisfying policy in polynomial time whilst satisfying safety constraints throughout the learning process. We then discuss E^4 as a practical algorithmic framework, including robust-constrained offline optimisation algorithms, the design of uncertainty sets for the transition dynamics of unknown states, and how to further leverage empirical observations and prior knowledge to relax some of the worst-case assumptions underlying the theory.

Keywords Safe artificial intelligence · Safe exploration · Model-based reinforcement learning · Constrained Markov decision processes · Robust Markov decision processes

Editors: Dana Drachler Cohen, Javier Garcia, Mohammad Ghavamzadeh, Marek Petrik, Philip S. Thomas.

✉ David M. Bossens
d.m.bossens@soton.ac.uk

Nicholas Bishop
nb8g13@soton.ac.uk

¹ School of Electronics and Computer Science, University of Southampton, Southampton SO17 1BJ, UK

1 Introduction

As machine learning methods are increasingly deployed in real-world applications, their safety is of foremost importance. Reinforcement learning (RL) is a particularly relevant example since in RL, an agent explores an initially unknown environment through observation, action, and reward. Among the various challenges to safe artificial intelligence [(see (Everitt et al., 2018; Amodei et al., 2016))], safe exploration and robustness to mismatches between training and application environment are of particular importance when applying RL algorithms in safety-critical missions. Unsafe exploration may occur when an RL system explores a new region of the state space, and in doing so performs actions with disastrous consequences. A mismatch between training and application environment may happen when the model within a model-based RL system is incorrect on some states, which can result in the system learning behaviours that are safe in simulation but dangerous in the real world. More generally, desirable behaviour in a complex world comes with various *safety constraints* such as avoiding damage or long-term wear-and-tear or following legal and social norms to avoid harm.

RL research has traditionally been studied from two angles, namely model-free and model-based RL, within the framework of Markov decision processes (MDPs). In model-free RL [(e.g. (Watkins and Dayan, 1992; Rummery and Niranjan, 1994))], the RL agent has no model of the environment and only learns the long-term cumulative reward associated with the action taken in a given state. Deep model-free RL [e.g. (Mnih et al., 2015; Schulman et al., 2017)] additionally incorporates the expressive power of deep neural networks for high performance across a variety of simulation environments. Such methods do not consider safety throughout the exploration process, and therefore are likely to violate safety criteria. In model-based RL [e.g. (Kearns and Singh, 2002; Brafman and Tennenholtz, 2002; Strehl et al., 2006)], the RL agent learns a model of the environment, including the transitions between states and the reward function, which can then be used to compute the long-term cumulative reward. Since learning in the real world may come with high-cost failures, a key benefit of model-based RL is the use of offline optimisation, which can sample trajectories from the environment model without requiring too many real world samples and therefore, failures. Consequently, model-based RL provides improved sample complexity and may also be beneficial for safety. Before model-based RL can be used in long-term safety-critical applications, at least four requirements must be satisfied, namely constraint-satisfaction, safe exploration, robustness to model errors, and applicability in non-episodic environments.

Constraint-satisfaction means that the agent must not violate any of the safety constraints defined by the user. The main framework for dealing with such constraints in this general manner is constrained Markov decision processes (CMDPs) (Altman, 1999), which aims to optimise the long-term reward subject to constraints over a long-term constraint-cost function; in particular, the RL agent is given a maximum budget of cumulative constraint-cost which cannot be exceeded. Traditionally, CMDPs are solved by linear programming, and more recently by policy optimisation methods for high-dimensional control problems (Achiam et al., 2017).

Safe exploration means that the agent must not venture too far into unknown or dangerous states and ensure a timely return to safe known states. Within CMDPs, there are a variety of approaches to safe return based on a few additional assumptions. For control problems where the aim is to stay stable in a particular goal state, Lyapunov functions can be used with additional assumption of continuity on the policy and environment dynamics

(Chow et al., 2018; Berkenkamp et al., 2017). An approach for general CMDPs is to define a supervising agent which can return the agent safely (Turchetta et al., 2020). However, this defers the problem of safe exploration to the supervising agent, which needs to be trained on realistic data. A model-free approach for safe exploration involves imposing an entropy constraint on policies to ensure exploration (Yang et al., 2021). As an alternative to the CMDP framework, one can also consider safety in terms of ensuring performance improves on a known policy (Thomas et al., 2015; Garcelon et al., 2020); in this case safe exploration can be ensured using off-policy evaluation (Thomas et al., 2015), which samples only from a known baseline policy, or by ensuring the online exploration is monotonically improving the performance (Garcelon et al., 2020).

Robustness to model errors means that the agent must be prepared for the worst-case when its model is inaccurate. This approach has been considered primarily by the framework of robust Markov decision processes, which accounts for uncertainty on the transition dynamics model (Iyengar, 2005; Nilim & Ghaoui, 2005; Wiesemann et al., 2013) and which has been integrated in CMDPs by Russel et al. (2020). Others additionally incorporate uncertainty over the reward model with applications to non-stationary MDPs (Jaksch et al., 2010; Lecarpentier & Rachelson, 2019), although these approaches are in the unconstrained setting and the optimistic bias in Jaksch et al. (2010) is likely to be unsafe. A setting with both reward and constraint-cost uncertainty has also been considered but this assumes known transition dynamics (Zheng & Ratliff, 2020).

Applicability in non-episodic (also called “continuing”) environments means that the agent will be embedded in a long-term environment with no resets. The potential of continued sequential dependencies makes non-episodic environments challenging; it is therefore no surprise that non-episodic environments became a topic for recent benchmark formulations [(see (Naik et al., 2021; Platanios et al., 2020)]. In safety-critical settings, the continued lifetime means that any failure is unforgiving—unlike in the episodic video games. Therefore, the above-mentioned challenge of returnability becomes critically important in such environments. Somewhat similar to Turchetta et al. (2020), Eysenbach et al. (2018) defines a separate policy for resetting, with the main aim to abort and return to an initial state, mimicking the episodic setting without manual resets being required. Similar principles may be applied to the safety-constrained setting as well, where one could potentially target return to a larger set of safe and known states.

With respect to the above requirements, Explicit Explore or Exploit (E^3) (Kearns & Singh, 2002), a model-based RL algorithm for near-optimal RL with polynomial sample and time complexity guarantees, provides a unique starting point for safe model-based RL. E^3 approximates the MDP with which the agent is interacting and the algorithm accounts for model errors and the resulting value estimate errors by distinguishing between known states, which have been estimated correctly, and unknown states, which have not been estimated correctly. E^3 also provides a natural way to deal with exploration and exploitation in a non-episodic environment, namely by alternating limited-step trajectories, each of which is long enough to assess value function statistics correctly, and then explicitly choosing either an exploration policy or an exploitation policy. The exploitation policy is chosen when an optimal policy is available from the given starting state of the trajectory, and an exploration policy is chosen otherwise, in an attempt to find an unknown state. Recent work (Henaff, 2019) has also shown the practical feasibility of the approach for continuous state spaces, comparing favourably to state-of-the-art deep RL.

What is currently missing from E^3 is a suitable way to deal with constraint-satisfaction across the known and unknown states, as well as a suitable method for providing a safe return, or “escape”, from the unknown states back to the known states. Considering

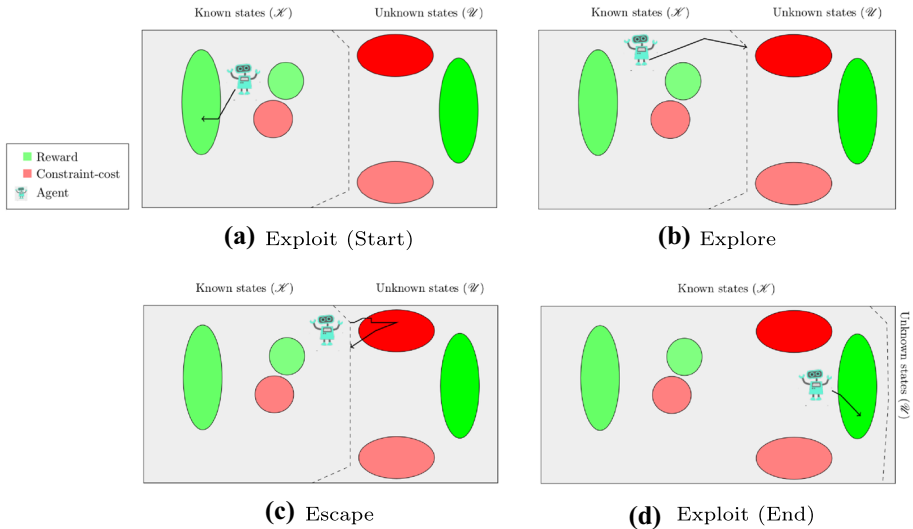


Fig. 1 Diagram illustrating E^4 . E^4 repeatedly takes limited-step trajectories in either known or unknown states. **a** The “exploitation policy” has the aim of solving the CMDP within the known states. If from a known state it has near-optimal constrained limited-step performance, then this policy is applied. **b** From another state, the exploitation policy may not be near-optimal; in this case an “exploration policy” aims to reach the unknown states, in an attempt to make them known. **c** Once in the unknown states, a balanced wandering behaviour ensures that actions are equally often taken for each unknown state, thereby yielding reliable statistics for different actions. Once there is an observed future risk of constraint-violation, an “escape policy” goes back to the known states as quickly as possible. **d** At the final stages of the E^4 algorithm, nearly all states have been frequently visited with balanced action-visitations, making them known. This then allows finding a near-optimal exploitation policy for the full CMDP

an additional “escape” policy in addition to an exploration and exploitation policy, this paper formulates an algorithm, called Explicit Explore, Exploit, or Escape (E^4 ; see Fig. 1). E^4 extends the E^3 algorithm to satisfy safety constraints throughout the lifetime of the RL agent through the following algorithmic contributions.

- Using the CMDP rather than an MDP framework, E^4 additionally models another stream of reinforcement signals called constraint-costs, which are constrained to a long-term budget and which represent the cost of state-action pairs.
- A correction is formulated for the constraint-cost budget in offline optimisation to account for the potential model errors in known states.
- Using a specialised “escape policy”, optimised to return to the known states as quickly as possible, E^4 halts the balanced wandering behaviour of E^3 as soon as there is a risk of exceeding the budgetary constraint in unknown states.
- Analytical formulas are derived to determine an allowed pseudo-budget in known and unknown states to ensure the budget over the full CMDP is not exceeded.
- To ensure the escape policy returns without exceeding the pseudo-budget in unknown states, even under worst-case assumptions, we propose four possible approaches, making modifications to existing robust and constrained optimisation algorithms:
 - using the robust CMDP policy gradient (Russel et al., 2020) as is in unknown states and optionally in known states;

- using the robust linear programming technique by Zheng and Ratliff (2020), which is applicable to episodic and undiscounted CMDPs with known transition dynamics, in the known states, by accounting for the discount factor and by considering a limited-step trajectory within the full lifetime of the E^4 agent;
- incorporating an uncertainty over transition dynamics into the constraints of constrained DP (Altman, 1998); and
- incorporating constraints into robust DP (Nilim & Ghaoui, 2005; Iyengar, 2005) by reformulating the CMDP as a Lagrangian MDP (Taleghan & Dietterich, 2018).

Applying these algorithmic principles, we demonstrate that E^4 finds a near-optimal policy within the set of policies that satisfies the budgetary constraints. Doing so yields similar sample complexity as E^3 but adds time complexity due to robust optimisation. Compared to the works mentioned above, the key distinctive features of E^4 are 1) safety-constraint satisfaction throughout the lifetime; 2) the explicit explore, exploit, or escape structure; 3) modelling the transition dynamics, reward function, constraint-cost function, and their uncertainty for robust offline optimisation; and 4) the aim to find a near-optimal constrained policy, rather than to improve on a known policy.

2 Preliminaries and definitions

The framework for E^4 is based on the CMDP task-modelling framework. In CMDPs, at each time step, an agent receives a state, performs an action, receives a reward, and a constraint-cost; the goal of the agent is to maximise the long-term cumulative reward whilst not exceeding a pre-defined budget of constraint-cost. Formally, a CMDP is a tuple $\langle \mathcal{S}, \mathcal{A}, P^*, r, c, d, \gamma \rangle$, where \mathcal{S} is the finite state space of size S ; \mathcal{A} is the finite action space of size A ; $P^* : \mathcal{S} \times \mathcal{A} \rightarrow \Delta^{\mathcal{S}}$ is the true transition dynamics of the environment with $\Delta^{\mathcal{S}} = \{p \in \mathbb{R}^{\mathcal{S}} : p^T \mathbf{1} = 1\}$ being the probability simplex in \mathcal{S} ; $r(s, a)$ is the average reward obtained for performing action $a \in \mathcal{A}$ in state $s \in \mathcal{S}$; $c(s, a)$ is the average constraint-cost for performing action $a \in \mathcal{A}$ state $s \in \mathcal{S}$; d is the budget on the expected asymptotic constraint-cost (see below); and $\gamma \in [0, 1)$ is the discount factor used for computing the long-term value and constraint-cost function based on the expected cumulative reward and constraint-cost, respectively. When at time t the RL agent is presented with a state s_t , its objective within the CMDP is to find a policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$ that maximises the expected asymptotic value,

$$V_{\pi}(s_t) = \mathbb{E}_{\pi, P^*} \left[\sum_{k=0}^{\infty} \gamma^k r(s_{t+k}, a_{t+k}) \right], \quad (1)$$

while the expected asymptotic constraint-cost, defined as

$$C_{\pi}(s_t) = \mathbb{E}_{\pi, P^*} \left[\sum_{k=0}^{\infty} \gamma^k c(s_{t+k}, a_{t+k}) \right], \quad (2)$$

must satisfy $C_{\pi}(s_t) \leq d$. Moreover, assumptions on the CMDP are three-fold. First, the reward distribution for a state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$ has mean $r(s, a) \in [0, r_{\max}]$ and variance $\text{Var}(s) \leq \text{Var}_{\max}^r$. Second, the constraint-cost distribution for a state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$ has a mean $c(s, a) \in [0, c_{\max}]$ and variance $\text{Var}^c(s) \leq \text{Var}_{\max}^c$. Third, a general upper-bound for the expected T-step value is denoted by $G_{\max}^r(T) = \sum_{t=0}^{\max_{\infty}} \gamma^t r_{\max} = \frac{r_{\max}}{1-\gamma}$ while analogously the general upper-bound for the expected T-step constraint-cost is

denoted by $G_{\max}^c(T) = \sum_{t=0}^{\infty} \gamma^t c_{\max} = \frac{c_{\max}}{1-\gamma}$. In practice, the above-mentioned maximal variance, reward, value, etc. do not need to be known exactly; any upper bound suffices, although tighter bounds improve the results. Knowledge of r_{\max} and c_{\max} are a basic requirement since the other maxima can be upper-bounded from them. To give a few examples of when r_{\max} and c_{\max} can be known exactly or upper-bounded, one can consider applications where rewards or constraint-costs represent a limited resource, such as energy, food, or money, or a physical force, such as friction or torque. A further property of the CMDP being studied is that it has a limited “diameter”. Completely analogous to the diameter (Jaksch et al., 2010) for MDPs, we define the diameter of a CMDP M as the maximal expected number of actions from one arbitrary state $s \in \mathcal{S}$ to any other $s' \in \mathcal{S}$ under the best possible stationary deterministic policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$ for the choice of (s, s') , or

$$D(M) := \max_{s \neq s'} \min_{\pi} \mathbb{E}[W(s'|s, M, \pi)], \quad (3)$$

where $W(s'|s, M, \pi)$ is the number of actions from s to s' given the model M and policy π .

While the above functions span over an infinite horizon, it is often of interest to compute their T -step approximation to make decisions based on limited data. The corresponding T -step approximations instead sum over T steps and are denoted by $V_{\pi}(s_t, T)$ and $C_{\pi}(s_t, T)$, respectively. From the class of policies Π , the class of constrained policies for any state $s \in \mathcal{S}$, a budget d , a number of time steps T , is denoted as $\Pi_c(s, d, T) = \{\pi \in \Pi : C_{\pi}(s, T) \leq d\}$. Denoting $\Pi_c(s, d)$ for $T \rightarrow \infty$, $\Pi_c(s, d) \neq \emptyset$ is non-empty by assumption, and therefore any other choice of T yields $\Pi_c(s, d, T) \neq \emptyset$. An optimal constrained policy for a T -step trajectory in M is then defined as $\pi^* = \arg \max_{\pi \in \Pi(s, d, T)} V_{\pi}(s, T)$ where for the asymptotic case of course $T \rightarrow \infty$. In addition to the expected asymptotic constraint-cost, which computes the expectation over different possible trajectories, the analysis also uses a path-based constraint-cost and value, which computes the average cost for a particular state-action trajectory and is denoted for a T -step path p as $C(p) = \sum_{t=0}^{T-1} \gamma^t c(s_t, a_t)$ and $V(p) = \sum_{t=0}^{T-1} \gamma^t r(s_t, a_t)$, respectively.

With M being the CMDP to be solved, this paper takes a model-based approach in which an approximate model $\hat{M} = \langle \mathcal{S}, \mathcal{A}, \hat{P}, \hat{r}, \hat{c}, d', \gamma \rangle$ is continually improved by sampling from M , with hat-notations emphasising the model and its components are estimated by the sample mean, and where a potentially lower budget d' , $c_{\max} \leq d' \leq d$, is used when the model is induced over a subset of the full CMDP (see Sect. 3.2), where the relation with c_{\max} is an assumption of the algorithmic approach to ensure at least one action can be performed before potentially exceeding d' . For states that have not yet been visited frequently, we formulate an uncertainty set (or ambiguity set), \mathcal{P} , as the set of transition dynamics models consistent with the samples from (or prior knowledge on) the true but unknown transition dynamics, P^* . For example, one can define a confidence interval around $\hat{P}_{s,a}$, the estimated (also called “nominal”) transition dynamics model, with a budget $\psi_{s,a}$ for state-action pair (s, a) :

$$\mathcal{P}_{s,a} = \{P \in \Delta^{\mathcal{S}} : \|P - \hat{P}_{s,a}\|_1 \leq \psi_{s,a}\}. \quad (4)$$

Within the states that have not been frequently visited yet, the uncertainty set is then used to optimise the constraints robustly, that is, with $\max_{p \in \mathcal{P}} \hat{C}_{\pi}(s_t) \leq d$, where \hat{C}_{π} indicates the estimate of the expected asymptotic constraint-cost, with more precise definitions and notations to follow.

3 Main theorem

With the preliminary definitions in mind, this section states and proves the main theorem, which postulates the existence of an algorithm, namely E^4 , which explores a CMDP safely within the constraint-cost budget and which finds a near-optimal constrained policy within polynomial time.

Theorem 1 There exists an algorithm (E^4) that outputs with probability $1 - \delta$ a near-optimal constrained policy π for s with $V_\pi(s) \geq V_{\pi^*}(s) - \epsilon$ and $C_\pi(s) \leq d$ with sample complexity and time complexity that is polynomial in $1/\epsilon, 1/\delta, S$, the horizon time $1/(1 - \gamma)$, r_{\max} , and c_{\max} . Moreover, the non-stationary policy $\pi'' = \{\pi_i\}_{i=1}^\infty$ induced by the exploration process of E^4 yields expected asymptotic constraint-cost $C_{\pi''}(s_t) \leq d$ with probability at least $1 - (UA\delta_\nu + \delta)$.

The E^4 algorithm is based on the model-based CMDP framework presented in the previous section. The Constrained Simulation Lemma (Sect. 3.1) shows that a sufficiently accurate CMDP model of the reward function, constraint-cost function, and transition dynamics allow, with high probability, an ϵ -correct value function and constraint-cost function approximation. Moreover, it is sufficient to visit each state-action pair a number of times that is polynomial in $S, T, 1/(1 - \gamma), r_{\max}, 1/\epsilon$, and $1/\delta$ to obtain such an accurate CMDP model. This leads to a natural definition of “known states”, which have been visited often with high frequency for each action and thereby have accurate models, and “unknown states”, which have not been visited often enough. Each of the two cases can then be treated as a different Induced CMDP (see Sect. 3.2). Therefore when starting a limited T -step trajectory from a state $s \in \mathcal{S}$, its strategy for known states differs from that in the unknown states.

Starting from a known state, the agent may already have a near-optimal policy available. If so, then this policy is best exploited for the following T -step trajectory. If not, then the model must be further improved by exploring unknown states. This is formalised by the l -safe Explore-or-Exploit Lemma (Sect. 3.3). This lemma shows that for policies satisfying a given a constraint-budget within the known states, there either exists a “exploitation policy” which is a near-optimal constrained policy or there exists an “exploration policy” which with high probability finds an unknown state. The l -safe Explore-or-Exploit Lemma then shows how to define a correction to the original budget to account for the approximation error of the model, resulting in a high-probability guarantee for satisfaction of the original constraint across the known states.

When exploring unknown states, the agent must attempt to make these states known by performing many state-action visitations. A key challenge is that this risks exceeding the constraint-cost budget because the agent has little knowledge of the unknown states. With worst-case assumptions on the transition dynamics and constraint-costs, the Safe Escape Lemma (Sect. 3.4) provides a high-probability guarantee for an optimised “escape policy” to find a path from the unknown states to the known states whilst satisfying the constraint-cost budget defined in the unknown states. Since the escape policy will not select each action equally frequently from a given state, the escape policy provides no guarantee on making states known. To safely gain as much knowledge on the unknown states as possible, E^4 performs a balanced wandering behaviour, which selects the least-taken action for a given state, as long as timely escape is ensured by the escape policy, as formalised by the Safe Balanced Wandering Lemma (Sect. 3.5).

To integrate the results for the known and unknown states, the Escape Budget Lemma (Sect. 3.6) shows how to set the budget for the unknown states, d' , to satisfy the budget on the full CMDP (d) given a budget within the known states d'' ; this also helps to define the conditions for constraint-satisfiability. The remainder of the section (Sect. 3.8) then puts all the lemmas together and demonstrates the validity of Theorem 1.

3.1 Constrained simulation lemma

Since the model is estimated based on the sample mean, repeatedly visiting the same state allows estimating the transition dynamics, the reward function, and the constraint-cost function. Therefore, the true CMDP M can be approximated by simulation model \hat{M} with a sufficient number of samples. If the simulation model is sufficient, the expected asymptotic value and constraint-cost of the CMDP can be approximated with high accuracy, a finding formalised by the Constrained Simulation Lemma (Sect. 2). The set of “known states” is then defined as those states which have sufficient visitations and are therefore approximated correctly.

The Constrained Simulation Lemma relies on the relation between the expected asymptotic and T-step value and constraint-cost. Specifically, the following Lemma shows that for sufficiently large T , the expected asymptotic value and constraint-costs are ϵ -close to their T-step counterparts.

Lemma 1 Constrained T-step Estimation Lemma: *Let M be the CMDP, π be a policy in M , and $T \geq \frac{1}{1-\gamma} \ln \left(\frac{\max(r_{\max}, c_{\max})}{\epsilon(1-\gamma)} \right)$. Then for any state $s \in S$ we have*

- (a) $V_{\pi}(s, T) \leq V_{\pi}(s) \leq V_{\pi}(s, T) + \epsilon$; and
- (b) $C_{\pi}(s, T) \leq C_{\pi}(s) \leq C_{\pi}(s, T) + \epsilon$.

For these reasons, we call $\frac{1}{1-\gamma} \ln \left(\frac{\max(r_{\max}, c_{\max})}{\epsilon(1-\gamma)} \right)$ the ϵ -horizon time of the CMDP M .

Proof of (a): Define MDP $M^- = \langle S, \mathcal{A}, r, \gamma, P \rangle$ with the same states, action space, reward function, discount factor, and transition dynamics as M . Let π be any policy in M^- . The result follows directly from the original T-step estimation Lemma (see Lemma 2 in Kearns and Singh (2002)) as this holds for any policy in any MDP. In short, the first inequality $V_{\pi}(s, T) \leq V_{\pi}(s)$ follows from rewards being non-negative and the expected asymptotic value additionally considering time steps $t = T + 1, \dots, \infty$; the second inequality follows after requiring $\epsilon \leq \gamma^T \frac{r_{\max}}{1-\gamma}$ (the maximal remaining reward not accounted for after time step T) and then solving for T .

Proof of (b): The proof is completely analogous as the transition dynamics and policy are the same while the constraint-cost function is analogously bounded in $[0, c_{\max}]$. \square

The basic reasoning behind the Constrained Simulation Lemma is that, if an estimator \hat{M} of the true CMDP M is sufficiently accurate, then \hat{M} can be used as a simulation model to simulate realistic trajectories of M ; with offline optimisation algorithms, this can then yield nearly correct estimates of the expected asymptotic value and the expected asymptotic constraint-cost. To develop this reasoning, a suitable definition for the accuracy of such a simulation model is formalised below.

Definition 1 CMDP \hat{M} is an α -approximation of CMDP M if and only if:

- 1) for all state-action-pairs $(s, a) \in \mathcal{S} \times \mathcal{A}$ and : $r(s, a) - \alpha \leq \hat{r}(s, a) \leq r(s, a) + \alpha$;
- 2) for all transitions $(s, a, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$: $P_{s,a}^*(s') - \epsilon \leq \hat{P}_{s,a}(s') \leq P_{s,a}^*(s') + \epsilon$.
- 3) for all state-action-pairs $(s, a) \in \mathcal{S} \times \mathcal{A}$: $c(s, a) - \alpha \leq \hat{c}(s, a) \leq c(s, a) + \alpha$.

With this definition in mind, the Constrained Simulation Lemma states that as long as T is chosen according to the ϵ -horizon time of M and the simulation \hat{M} is α -correct, the simulated value and constraint-costs are ϵ -correct with respect to the true expected asymptotic value and constraint-cost.

Lemma 2 *Constrained Simulation Lemma.* Let $T \geq \frac{1}{1-\gamma} \ln(\frac{\max(r_{\max}, c_{\max})}{\epsilon(1-\gamma)})$, let CMDP \hat{M} be an α -approximation of CMDP M with $\alpha = O((\epsilon/(STG))^2)$ where $G = \max(G_{\max}^r(T), G_{\max}^c(T))$, and let π be a policy in M . Then for any state s ,

- (a) $V_{\pi}(s) - \epsilon \leq \hat{V}_{\pi}(s) \leq V_{\pi}(s) + \epsilon$; and
- (b) $C_{\pi}(s) - \epsilon \leq \hat{C}_{\pi}(s) \leq C_{\pi}(s) + \epsilon$.

The proof extends the original Simulation Lemma (see Lemma 4 in Kearns and Singh (2002)) to include the constraint-costs. As it is extensive and otherwise analogous, it is provided in the appendix.

Having established the relation between the expected asymptotic constraint-cost and value on the one hand and the α -approximation on the other hand, the following lemma provides the number of samples for an accurate estimate of the expected asymptotic constraint-cost and value with error of at most ϵ .

Lemma 3 *Known State Lemma.* Given state s has been visited

$$m = O((STG/\epsilon)^4 \text{Var}_{\max} \ln(1/\delta)) \quad (5)$$

times, where $G = \max(G_{\max}^r(T), G_{\max}^c(T))$ and $\text{Var}_{\max} = \max(\text{Var}_{\max}^r, \text{Var}_{\max}^c)$, and from s each action has been executed at least $\lfloor m/A \rfloor$ times, then with probability of at least $1 - \delta$ we have

- a) $|P_{s,a}^*(s') - \hat{P}_{s,a}(s')| = O(\epsilon/(STG))^2$ for any $(s, a, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$;
- b) $|r(s, a) - \hat{r}(s, a)| = O((\epsilon/(STG))^2)$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$; and
- c) $|c(s, a) - \hat{c}(s, a)| = O((\epsilon/(STG))^2)$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$.

Proof The proof is based on Hoeffding's inequality and Chernoff bounds (see Appendix A for details). \square

Following the above lemma, the notion of “known states” can now be defined directly from Eq. 5. Intuitively, contrasting known versus unknown states allows to distinguish regions for which expected asymptotic value and constraint-cost is estimated correctly, and regions for which they are not. This contrast between known and unknown states will be critical to provide safety guarantees for the RL algorithm later on.

Definition 2 Known states \mathcal{K} and unknown states \mathcal{U} : a state $s \in \mathcal{S}$ is called known if $\forall a \in \mathcal{A} : n(s, a) \geq m_{\text{known}}$ where $m_{\text{known}} \in \mathbb{N}$ is $O((STG/\epsilon)^4 \text{Var}_{\max} \ln(1/\delta))$ following

Eq. 5. Alternatively, $s \in \mathcal{S}$ is called **unknown** if $\exists a \in \mathcal{A} : n(s, a) < m_{\text{known}}$. Known states are collectively referred to as $\mathcal{K} = \{s \in \mathcal{S} | \forall a \in \mathcal{A} : n(s, a) \geq m_{\text{known}}\}$ while unknown states are referred to as $\mathcal{U} = \{s \in \mathcal{S} | \exists a \in \mathcal{A} : n(s, a) < m_{\text{known}}\}$.

3.2 Induced CMDPs

Since only known states are modelled to a sufficient accuracy, it is useful to consider CMDP simulations over the known states only. The following defines the notion of an “induced CMDP”, which limits the original CMDP to a subset of the original state space (e.g., the known states). The induced CMDPs allow two useful results. First, since the induced CMDP is also a CMDP like in the Constrained Simulation Lemma, an analogous result holds for induced CMDPs over the set of known states (see Lemma 4); that is, an estimated induced CMDP is ϵ -correct with probability at least $1 - \delta$ when compared to the true induced CMDP. Second, the expected T -step value and constraint-cost within induced CMDPs provide lower bounds for the expected T -step value and constraint-cost, respectively, in the original CMDP (see Lemma 5).

Definition 3 Given a CMDP M , define for a subset of states $\mathcal{S} \subset \mathcal{S}$ an **induced CMDP** $M_{\mathcal{S}}$, which is equal in all but the following respects:

- its state space is $\mathcal{S} \cup s_0$ where s_0 is a terminal state of the CMDP.
- transitions to states in $s \notin \mathcal{S}$ are redirected to a terminal state s_0 , which terminates the episode and yields terminal reward $r(s_0) = 0$ and terminal constraint-cost $c(s_0) = 0$.
- for each $(s, a) \in \mathcal{S}$, rewards and constraint-costs always yield their mean $r(s, a)$ and $c(s, a)$ deterministically, with zero variance.

When a CMDP is induced over the known states \mathcal{K} (see Definition 2), the resulting CMDP is denoted as $M_{\mathcal{K}}$ and is called the **known-state CMDP** of M .

In the following, whenever a value function, constraint-cost, or other quantity is computed in an induced CMDP $M_{\mathcal{S}}$, this will be denoted by the conditioning operator $|M_{\mathcal{S}}$; for example, $V_{\pi}(s|M_{\mathcal{K}})$ denotes the expected asymptotic value of π in the known state CMDP. For brevity and consistency with the earlier sections, when the conditioning is on the full CMDP of interest, M , the conditioning is omitted; for example, $V_{\pi}(s)$ denotes the expected asymptotic value of π in the full CMDP of interest.

Now the so-called Induced Simulation Lemma applies the Constrained Simulation Lemma to known-state CMDPs.

Lemma 4 *Induced Simulation Lemma* Let M be the CMDP of interest and let $M_{\mathcal{K}}$ be its known state CMDP, and let $T \geq \frac{1}{1-\gamma} \ln \left(\frac{\max(r_{\max}, c_{\max})}{\epsilon(1-\gamma)} \right)$. Then, for any policy π and any state $s \in \mathcal{S}$, with probability $1 - \delta$, we have

$$V_{\pi}(s|M_{\mathcal{K}}) - \epsilon \leq \hat{V}_{\pi}(s|\hat{M}_{\mathcal{K}}) \leq V_{\pi}(s|M_{\mathcal{K}}) + \epsilon$$

and

$$C_{\pi}(s|M_{\mathcal{K}}) - \epsilon \leq \hat{C}_{\pi}(s|\hat{M}_{\mathcal{K}}) \leq C_{\pi}(s|M_{\mathcal{K}}) + \epsilon.$$

Proof The proof follows directly from the Constrained Simulation Lemma (see Lemma 2) and the Known State Lemma (see Lemma 3). \square

As one is interested in what happens in the “real world”, the below lemma relates the induced CMDP to the original CMDP M .

Lemma 5 *Induced underestimation lemma: Let $\mathcal{S} \subset \mathcal{S}$. For any $s \in \mathcal{S}$, and any policy π in M , we have*

$$V_{\pi}(s, T|M_{\mathcal{S}}) \leq V_{\pi}(s, T)$$

and

$$C_{\pi}(s, T|M_{\mathcal{S}}) \leq C_{\pi}(s, T).$$

Proof Let π be a policy in M and let $s \in \mathcal{S}$. Then we have, analogous to the original lemma, the following results because π stops at time $T_{\text{stop}} \leq T$ and rewards and constraint-costs are positive. To illustrate, for the value function, we have:

$$V_{\pi}(s, T|M_{\mathcal{S}}) = \sum_p \mathbb{P}_{\pi, P^*}[p] \sum_{t=0}^{T_{\text{stop}}-1} \gamma^t r_t \leq V_{\pi}(s, T)$$

due to rewards being positive and the paths being terminated earlier at a time $T_{\text{stop}} \leq T$ at terminal state $s_{T_{\text{stop}}} \in \mathcal{S} \setminus \mathcal{S}$. The result for the constraint-cost is completely analogous. \square

3.3 l -safe Explore-or-Exploit lemma

Following the definition of induced CMDPs, this section provides results for exploration and exploitation within the known states. For any induced CMDP which has satisfiable constraint-cost budget, the Constrained Explore-or-Exploit lemma (Lemma 6) proves the existence of either an exploitation policy, which solves the CMDP within the known states near-optimally, and an exploration policy, which gives a high-probability guarantee of visiting an “unknown state”. This allows the l -safe Explore-or-Exploit Lemma (Lemma 7), which provides safety, i.e. guarantees on constraint-satisfaction, within the known states by defining a lower budget l that accounts for worst-case modelling errors (hence the term l -safe). The budget l subtracts the model errors from Lemmas 1 and 4 to the Constrained Explore-or-Exploit Lemma.

Based on defining an induced CMDP over a subset of states \mathcal{S} (e.g. the known states, in which case $\mathcal{S} = \mathcal{K}$), the Constrained Explore-or-Exploit Lemma proves either the existence of an optimal constrained policy on the full CMDP or a constrained policy that finds a state not in \mathcal{S} in T -steps with high probability.

Lemma 6 *Constrained Explore-or-Exploit Lemma* Let M be any CMDP, let \mathcal{S} be any subset of states $\mathcal{S} \subset \mathcal{S}$, and let $M_{\mathcal{S}}$ be the induced CMDP over \mathcal{S} with a given budget d . For any $s \in \mathcal{S}$, for any T , and any $\epsilon \geq 0$, we have **either a)** there exists a policy $\pi \in M_{\mathcal{S}}$ for which $V_{\pi}(s, T|M_{\mathcal{S}}) \geq V_{\pi^*}(s, T) - \epsilon$, where $\pi^* = \arg \max_{\pi \in \Pi_{\epsilon}(s, d, T)} V_{\pi}(s, T)$ is the optimal constrained T -step policy, and which satisfies $C_{\pi}(s, T|M_{\mathcal{S}}) \leq d$; **or b)** there exists a policy

π in $M_{\mathcal{S}}$ which reaches the terminal state s_0 in $S \setminus \mathcal{S}$ in at most T steps with probability $p > \epsilon / G_{\max}^r(T)$, and $C_{\pi}(s, T | M_{\mathcal{S}}) \leq d$.

Proof The proof and lemma is analogous to the original Explore-or-Exploit Lemma, except that the optimal constrained policy rather than the optimal policy is considered (see Appendix B). \square

Now by integrating the Constrained Simulation Lemma and the Constrained Explore-or-Exploit Lemma, a characterisation of the safety can be provided in terms of the probability of constraint-satisfaction during exploitation or exploitation in the known-state CMDP. This safety guarantee, based on the *estimated* known-state CMDP, is provided by the l -safe Explore-or-Exploit Lemma below.

Lemma 7 *l -safe Explore-or-Exploit Lemma:* Let $\epsilon \in [0, d/2)$, $l \leq d - 2\epsilon$, $\delta \in [0, 1]$, \mathcal{K} , and $T \geq \ln \left(\frac{\max(r_{\max}, c_{\max})}{\epsilon(1-\gamma)} \right) \frac{1}{1-\gamma}$ (i.e. at least the ϵ -horizon time). Furthermore, let $M_{\mathcal{K}}$ be the known-state CMDP and $\hat{M}_{\mathcal{K}}$ its estimate. Then with probability $1 - \delta$, a policy π that satisfies $\hat{C}_{\pi}(s, T | \hat{M}_{\mathcal{K}}) \leq l$ for $s \in \mathcal{K}$ will satisfy $C_{\pi}(s | M_{\mathcal{K}}) \leq d$.

Proof Due to the Induced Simulation Lemma and the Constrained T-step Estimation Lemma, we have for any state $s \in S$ with probability $1 - \delta$:

$$C_{\pi}(s | M_{\mathcal{K}}) \leq \hat{C}_{\pi}(s | \hat{M}_{\mathcal{K}}) + \epsilon \leq \hat{C}_{\pi}(s, T | \hat{M}_{\mathcal{K}}) + 2\epsilon \leq l + 2\epsilon \leq d.$$

\square

3.4 Safe escape lemma

Contrasting to the known states, any model based on the unknown states is inaccurate, making it difficult to guarantee safe return to the known states before exceeding the constraint-cost budget. To overcome the challenge, we construct a “Worst-case Escape CMDP”, a CMDP induced over the unknown states that rewards escaping back to the known states before exceeding a constraint-cost budget and makes worst-case assumptions on the constraint-cost. Additionally incorporating worst-case transition dynamics through the use of an uncertainty set, the resulting robust CMDP provides a probabilistic guarantee for safety within the unknown states. The proof relies on robust constraint-satisfiability, a condition that is further discussed in Sect. 4.3.

Definition 4 Worst-case Escape CMDP Given a CMDP M , the Worst-case Escape CMDP $M_{\mathcal{U}}$ induced on \mathcal{U} is equivalent to the Induced CMDP on \mathcal{U} except that $c(s) = c_{\max}$ for all $s \in \mathcal{U}$.

Intuitively, the policy that optimises a Worst-case Escape CMDP escapes to the known states before violating the constraint-cost budget under worst-case assumptions.

The following result shows that a policy that is optimised robustly on an estimated Worst-case Escape CMDP will be able to escape the unknown states with high probability while also satisfying the constraint-cost budget within both the (non-estimated) Worst-case Escape CMDP and the full CMDP.

Lemma 8 *Worst-case Escape Lemma* Let M be the full CMDP, and $M_{\mathcal{U}}$ be the Worst-case Escape CMDP on \mathcal{U} , and $\hat{M}_{\mathcal{U}}$ its estimate. Let d' be the constraint-cost budget within the Worst-case Escape CMDP $M_{\mathcal{U}}$. Let $T' = \inf\{T \in \mathbb{N} : \sum_{t=0}^{T-1} \gamma^t c_{\max} \geq d'\}$, and $T \geq T'$. Moreover, let $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta^{\mathcal{S}}$ be an uncertainty set, such that for each state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$, the true transition dynamics model $P_{s,a}^*$ is contained in $\mathcal{P}_{s,a}$ with probability at least $1 - \delta_{\psi}$.¹ If π is a policy in $\hat{M}_{\mathcal{U}}$ that satisfies the constraint of the Worst-case Escape CMDP robustly, that is,

$$\max_{P \in \mathcal{P}} \hat{C}_{P,\pi}(s, T | \hat{M}_{\mathcal{U}}) \leq d'$$

for the initial state $s \in \mathcal{U}$, then with probability at least $1 - UA\delta_{\psi}$, π in $M_{\mathcal{U}}$ meets the constraint-cost budget $C_{P^*,\pi}(s | M_{\mathcal{U}}) \leq d'$ and takes at most T' steps in \mathcal{U} .

Proof Let π be a policy in $\hat{M}_{\mathcal{U}}$ that satisfies the constraint on a state $s \in \mathcal{U}$ robustly, i.e. $\max_{P \in \mathcal{P}} \hat{C}_{P,\pi}(s, T | \hat{M}_{\mathcal{U}}) \leq d'$. Due to union bound, we have with probability at least $1 - UA\delta_{\psi}$ that $P^* \in \mathcal{P}$ (i.e. the true transition dynamics model is contained in the uncertainty set) where $U = |\mathcal{U}|$. Therefore, with constraint-cost estimates equal to their true values c_{\max} for all steps taken in the worst-case CMDP, we have with probability at least $1 - UA\delta_{\psi}$,

$$C_{P^*,\pi}(s, T | M_{\mathcal{U}}) \leq \max_{P \in \mathcal{P}} \hat{C}_{P,\pi}(s, T | \hat{M}_{\mathcal{U}}) \leq d'.$$

Since the budget is not exceeded and T' is the infimum of the set $\{T \in \mathbb{N} : \sum_{t=0}^{T-1} \gamma^t c_{\max} \geq d'\}$, the number of steps taken in \mathcal{U} (with $c_t = c_{\max}$) is at most T' . \square

3.5 Balanced wandering

To be able to make an unknown state known, each action must be performed often enough, following the Known State Lemma (Lemma 3). In E^3 , this is achieved through balanced wandering.

Definition 5 For a given state $s \in \mathcal{S}$, **balanced wandering** takes the action that has been performed the least,

$$a^* = \arg \min_{a \in \mathcal{A}} n(s, a). \quad (6)$$

In E^4 , taking many arbitrary actions with balanced wandering may exceed the constraint-cost budget. How can the agent perform balanced wandering while achieving a safe escape? The answer is to perform balanced wandering as long as the predicted constraint-cost following the escape policy allows doing so, as formalised by the Safe Balanced Wandering Lemma.

¹ Note that the proof only applies the uncertainty set on the unknown states, \mathcal{U} , since after visiting a state in \mathcal{K} , the Worst-case Escape CMDP is terminated. Therefore, for the sake of the lemma, one might equally define the uncertainty set on the unknown states only, i.e. in the form $\mathcal{P} : \mathcal{U} \times \mathcal{A} \rightarrow \Delta^{\mathcal{S}}$. However, using the full state space is a convenience data structure in case one desires to also construct an uncertainty set for the known states (see Algorithm 1).

Lemma 9 *Safe Balanced Wandering Lemma* Let $M_{\mathcal{U}}$ be the Worst-case Escape CMDP on \mathcal{U} , and $\hat{M}_{\mathcal{U}}$ its estimate. Let d' be the constraint-cost budget within the Worst-case Escape CMDP $M_{\mathcal{U}}$. Let $T' = \inf\{T \in \mathbb{N} : \sum_{t=0}^{T-1} \gamma^t c_{\max} \geq d'\}$, and $T \geq T'$. Let $\mathcal{P} : S \times \mathcal{A} \rightarrow \Delta^S$ be an uncertainty set, such that for each state-action pair $(s, a) \in S \times \mathcal{A}$, the true transition dynamics model $P_{s,a}^*$ is contained in $\mathcal{P}_{s,a}$ with probability at least $1 - \delta_{\psi}$. Let π be a policy in $\hat{M}_{\mathcal{U}}$ that satisfies the constraint of the Worst-case CMDP robustly, that is,

$$\max_{P \in \mathcal{P}} \hat{C}_{P,\pi}(s, T | \hat{M}_{\mathcal{U}}) \leq d'$$

for the initial state $s \in \mathcal{U}$. Let μ be a strategy that performs balanced wandering but switches to policy π at time t when

$$C(p_t) + \gamma^t \max_{P \in \mathcal{P}} \max_{a_t \in \mathcal{A}} \sum_{s_{t+1} \in S} P_{s_t, a_t}(s_{t+1}) \hat{C}_{P,\pi}(s_{t+1}, T | \hat{M}_{\mathcal{U}}) \geq d' - c_{\max},$$

where $p_t = \{s, s_1, \dots, s_{t-1}\}$ is the t -step path taken so far from s . Then with probability at least $1 - UA\delta_{\psi}$, μ is safe, in the sense that

$$C_{P^*, \mu}(s, T | M_{\mathcal{U}}) \leq d'.$$

Proof Let $t \in \{0, \dots, T' - 1\}$ be the first step for which

$$C(p_t) + \gamma^t \max_{P \in \mathcal{P}} \max_{a_t \in \mathcal{A}} \sum_{s_{t+1} \in S} P_{s_t, a_t}(s_{t+1}) \hat{C}_{P,\pi}(s_{t+1}, T | \hat{M}_{\mathcal{U}}) \geq d' - c_{\max}.$$

Since one step expends at most c_{\max} , this implies

$$\begin{aligned} C(p_t) + \gamma^t \max_{P \in \mathcal{P}} \hat{C}_{P,\pi}(s, T | \hat{M}_{\mathcal{U}}) &\leq C(p_t) \\ &+ \gamma^t \max_{P \in \mathcal{P}} \max_{a_t \in \mathcal{A}} \sum_{s_{t+1} \in S} P_{s_t, a_t}(s_{t+1}) \hat{C}_{P,\pi}(s_{t+1}, T | \hat{M}_{\mathcal{U}}) \leq d' \end{aligned}$$

and therefore, by Lemma 8, with probability $1 - UA\delta_{\psi}$ that

$$C(p_t) + \gamma^t C_{P^*, \mu}(s_t, T | M_{\mathcal{U}}) \leq d'.$$

Since this holds for any path p_t from s , it also holds on expectation over t -step paths from s , such that with probability $1 - UA\delta_{\psi}$

$$C_{P^*, \mu}(s, T | M_{\mathcal{U}}) \leq d'.$$

□

3.6 Escape budget lemma

Having introduced how to use induced CMDPs for safety within known and unknown states, this section now turns to the full CMDP M : that is, when following T -step policies on induced CMDPs $M_{\mathcal{K}}$ and $M_{\mathcal{U}}$, does this still guarantee safety across the entire CMDP? To resolve this question, the following lemma, called the Escape Budget Lemma, formulates an “**escape budget**” $d' \leq d$, the highest constraint-cost in $M_{\mathcal{U}}$ that still provides a probabilistic safety guarantee over the entire CMDP M .

First follows a definition of safe return states, which indicate states from which a T -step trajectory is known to stay in \mathcal{K} following some policy π and which satisfies a much lower budget but has no requirements on the value (so typically yields a poor value).

Definition 6 A d_s -safe return state is a known state $s \in \mathcal{K}$ for which there exists a stationary policy $\pi_{d_s} : \mathcal{S} \rightarrow \mathcal{A}$ that takes at least T steps from s in \mathcal{K} with probability $1 - \delta$ and that satisfies $C_{\pi_{d_s}}(s, T | \gamma = 1) \leq d_s - \epsilon$. This policy π is called a d_s -safe return policy and the set of such states is denoted by \mathcal{K}_{d_s} .

While the CMDP of interest, M , has $\gamma \in [0, 1)$, the above definition is based on $\gamma = 1$ to ensure accounting for the worst-case impact of safe return on future constraint-satisfaction. Since T is the ϵ -horizon time, this implies any non-stationary policy π_n which first applies π_{d_s} for T steps and then any policy in M will yield $C_{\pi_n}(s) \leq d_s$ for any starting state $s \in \mathcal{K}_{d_s}$. The Escape Budget Lemma takes place in a setting where there is an cycle alternating between three policies:

- an exploration policy in $M_{\mathcal{K}}$, which aims to find an unknown state;
- a safe balanced wandering policy in $M_{\mathcal{U}}$, which performs balanced wandering followed by an escape back to the known states, in an attempt to make states known; and
- a safe return policy in $M_{\mathcal{K}}$, which acts T steps within \mathcal{K} with low expected T -step constraint-cost, to ensure long-term constraint-satisfaction despite the constraint-costs of exploration attempts.

This alternation defines a non-stationary policy $\pi = \{\pi_i\}_{i=0}^{\infty}$, where i indicates the particular policy being used, where $i = 0, 3, 6, \dots$ index exploration policies, $i = 1, 4, 7, \dots$ index safe balanced wandering policies, and $i = 2, 5, 8, \dots$ index safe return policies. Within this setting, the lemma determines the available escape budget based on the initial path cost in the known states, the worst-case expected cost in the unknown states (d'), and the worst-case expected asymptotic constraint-cost of the safe return policy (d'').

Lemma 10 Escape Budget Lemma Let M be the full CMDP, $M_{\mathcal{K}}$ be the known-state CMDP, $M_{\mathcal{U}}$ be the Worst-case Escape CMDP, and let $\mathcal{K}_{d_s} = \mathcal{K}$ for some level $d_s \leq d$. Let p be a path from $s \in \mathcal{K}$ to a starting state $s' \in \mathcal{U}$. Let $d'' \leq d - \epsilon$ be the budget in $M_{\mathcal{K}}$ and d'_i be the “escape” budget in $M_{\mathcal{U}}$ at the i ’th visitation to $M_{\mathcal{U}}$. Let $\pi = \{\pi_i\}_{i=0}^{\infty}$ be a non-stationary policy, defined on $M_{\mathcal{K}}$ with expected asymptotic cost $C_{\pi_i}(s | M_{\mathcal{K}}) \leq d''$ for $i = 0, 3, 6, \dots$, on $M_{\mathcal{U}}$ with expected asymptotic cost $C_{\pi_i}(s | M_{\mathcal{K}}) \leq d'_{i//3}$ ² for $i = 1, 4, 7, \dots$, and on $M_{\mathcal{K}}$ as a d_s -safe return policy for $i = 2, 5, 8, \dots$. Further, let $\Pi(s, d'') \neq \emptyset$. Let $d'_{\min} = d - d'' - 2d_s + \epsilon$ and $T_{\min} = \inf\{T \in \mathbb{N} : \sum_{t=0}^{T-1} \gamma^t c_{\max} \geq d'_{\min}\}$. Finally, define $D^+ := D(M) + 1$, where $D(M)$ is the diameter of M . Define the following requirements:

- Diameter requirement:** the diameter satisfies $D^+ \leq T_{\min} \leq T'_{i//3}$;
- Known-state requirements:** for any recent path p from $M_{\mathcal{K}}$ to unknown state $s_0 \in \mathcal{U}$, require $d'_i \leq \frac{1}{\gamma^{T_k}}(d - \gamma^{T_k+1}d_s - C(p_k))$;

² The notation $//$ indicates integer division

- (c) **Unknown-state requirement:** $d'_i \leq d - \gamma d_s$; and
 (d) **Safe return requirement:** $d_s \leq \frac{1}{2}(d + \epsilon - d'' - D^+ c_{\max})$.

Then these requirements are not conflicting and imply $C_\pi(s_t) \leq d$ for all $s_t \in \mathcal{S}$ at all times t . Moreover, the budget allows at least one time step of balanced wandering before escape.

Proof The proof will first show that the requirements (a), (b), (c), and (d) are not conflicting with each other. Then, the proof selects an arbitrary known state s_t at any time t and shows how the requirements yield $C_\pi(s_t) \leq d$. Then, the proof analogously selects an unknown state s_t at any time t and shows how the requirements yield $C_\pi(s_t) \leq d$. Finally, the proof shows that at least one time step of balanced wandering is allowed due to the definition of D^+ .

0) Requirements are not conflicting Requirements (b) and (c) are not conflicting with each other: since they are both upper bounds, one can set $d'_{i//3} = \min\left(\frac{1}{\gamma^{T_k}}\left(d - \gamma^{T_k+T'_{i//3}}d_s - C(p_k)\right), d - \gamma^{T'_{i//3}}d''\right)$. Now it is required to show that (a) does not conflict with this setting of $d'_{i//3}$. First, note that the worst case, where the discount factor is $\gamma = 1$ and where $d_s - \epsilon + d''$ is expended in the known states and d_s is expended from a safe return state, yields the lower bound for $d'_{i//3}$,

$$d - d'' + \epsilon - 2d_s = d_{\min} \leq d'_{i//3},$$

and therefore $T_{\min} \leq d'_{i//3}$. Filling in d_s according to requirement (d) yields

$$\begin{aligned} d'_{i//3} &\geq d - d'' + \epsilon - 2d_s \\ &\geq d - d'' + \epsilon - 2\left(\frac{1}{2}(d - d'' + \epsilon - D^+ c_{\max})\right) = D^+ c_{\max}, \end{aligned}$$

and therefore

$$T'_{i//3} \geq d'_{i//3}/c_{\max} \geq D^+,$$

which is consistent with requirement (a).

1) Known states Let $s \in \mathcal{K}$. If the agent is at the starting state of its d_s -safe return then the proof is finished by definition of $d_s \leq d$. Otherwise, the agent is either somewhere along the trajectory of the d_s -safe return policy or it is along the trajectory of the exploration policy. In both cases, the agent forms a path p_k from s to a starting unknown state $s_u \in \mathcal{U}$, then follows a safe balanced wandering policy that escapes to a d_s -safe return state s_k with probability $1 - UA\delta_\psi$, and then from s_k takes T steps using a d_s -safe return policy. Denoting the length of the path p_k as T_k , the cost of the path in the known states as $C(p_k)$, the cost of the path in the unknown states as $C(p_u)$, and the expected asymptotic constraint-cost from the safe return states as $C_\pi(s_k)$, the statement to prove is

$$C_\pi(s) = \mathbb{E}_{P^*, \pi} [C(p_k) + \gamma^{T_k} C(p_u) + \gamma^{T_k+T_u} C_\pi(s_k)] \leq d.$$

By definition of the diameter, requirement (a) ensures that there exists an escape policy π' which satisfies $C_{\pi'}(s|M_{\mathcal{U}}) \leq d'_{i//3}$ from any $s \in \mathcal{U}$. Therefore, setting $\pi_i = \mu$, where μ is a safe balanced wandering policy for the budget $d'_{i//3}$, we have $\mathbb{E}[C(p_u)] \leq d'_{i//3}$ and $T_u \geq 1$. Moreover, the safe return policy ensures that $C_\pi(s) \leq d_s$ such that

$$\begin{aligned}
C(p_k) + \gamma^{T_k} \mathbb{E}[C(p_u) + \gamma^{T_u} C_\pi(s_k)] \\
\leq C(p_k) + \gamma^{T_k} d'_{i//3} + \gamma^{T_k+1} d_s \\
\leq d,
\end{aligned}$$

where the last line follows from requirement (b). Since the inequality holds for any path p_k , it also holds on expectation over paths.

2) Unknown states Let $s \in \mathcal{U}$. The agent forms a path p_u from s to a starting unknown state $s_k \in \mathcal{K}$. Denoting the length of the path as T_u , based on the cost of the path $C(p_u)$ and the expected asymptotic cost over the terminal state in \mathcal{K} , the statement to prove is

$$C_\pi(s) = \mathbb{E}_{p^*, \pi}[C(p_u) + \gamma^{T_u} C_\pi(s_k)] \leq d.$$

Note that $T_u \geq 1$, and via requirement (a) and (b), $C_\pi(s|M_{\mathcal{U}}) = C_{\pi_i}(s|M_{\mathcal{U}}) \leq d'_{i//3}$ for some stationary policy π_i , and $C_\pi(s_k) \leq d_s$. Further filling in requirement (c) yields the statement to prove,

$$\begin{aligned}
C_\pi(s) &\leq d'_{i//3} + \gamma d_s \\
&\leq d - \gamma d_s + \gamma d_s = d.
\end{aligned}$$

3) At least one step of balanced wandering Let $j \in \{0, 1, \dots\}$. As shown in part 0) of the proof, $T'_j \geq D^+ = D(M) + 1$. By definition of the diameter, an escape policy can return to the known states within at most $D(M)$ time steps. Therefore, at least one step of balanced wandering can be performed without exceeding the budget d'_j . \square

For the $C_\pi(s_t) \leq d$ for all t in the current trajectory, it is required to perform computation b) over all known states recently visited and take the minimum, and then again the minimum between the result and the quantity in c). For such a computation, one can take the last T steps within \mathcal{K} and subtract ϵ for the computation of d' . Note that for unknown states, one only needs to perform c) on starting states because it assumes the worst-case c_{\max} for all coming T' steps whereas intermediate unknown states will only yield c_{\max} for a lower number of steps.

The Escape Budget Lemma hereby provides a general strategy for determining a desired escape budget. Below is an example of a high constraint-cost trajectory in \mathcal{K} and the resulting budget d' . The below trajectory is shorter than usual purely for demonstration purposes; any trajectory (except the initial $i = 0$ trajectory) would involve at least T steps from the known states due to the d_s -safe policy taking T steps prior to the exploration policy.

Example 1 Let $c_{\max} = 1$, $\gamma = 0.98$, and $\epsilon = 1.0$. The setting yields an ϵ -horizon time $T = \lceil \frac{1}{1-\gamma} \ln \left(\frac{c_{\max}}{\epsilon(1-\gamma)} \right) \rceil = 196$. Further, $d = 25$, and safe return policies exist for all $s \in \mathcal{K}$ with $d_s = 5.0$, both of which are well below the general upper bound $G_{\max}^c(T) + \epsilon = 49.05$. The diameter of the full CMDP is known to be $D(M) = 3$ and therefore set $D^+ = 4$. Compute the known-state budget as $d'' = d - 2d_s - D^+ + \epsilon = 12.0$. Computing $T_{\min} = \inf\{T \in \mathbb{N} : \sum_{t=0}^{T-1} \gamma^t c_{\max} \geq d - d'' - 2d_s + \epsilon\} = 3$, $D^+ = T_{\min}$ implies the constraint on the unknown states is certainly satisfiable for any trajectory. The current trajectory of constraint-cost yields $[0, c_{\max}, 0, 0, c_{\max}, c_{\max}/2, c_{\max}, c_{\max}, c_{\max}]$, resulting in a list of path-based constraint-costs to check, $C_p = [3.29, 3.66, 2.95, 3.28, 3.65, 2.94, 2.71, 1.9, 1.0]$ for $t = 0, \dots, 8$. For condition (b), looping over $C \in C_p$ and computing

$d' \leftarrow \frac{1}{\gamma^{T_k}}(d - \gamma^{T_k+1}d_s - C(p))$, d' is minimal for $t = 4$, namely $d' = 17.53$. Condition (c) does not require any adjustment to d' . This setting of d' leaves $T' = 21$ steps from which at least $T' - D(M) = 18$ steps of balanced wandering can be allocated. Following escape, the safe return policy will take T steps in the known states, where it yields $C_\pi(s, T|\gamma = 1) \leq d_s - \epsilon = 4.0$, before an exploration policy starts an exploration attempt.

3.7 Simulated budget satisfaction lemma

Previously, the l -safe Explore-or-Exploit Lemma (Lemma 7) showed that, if a policy yields $d - 2\epsilon$ -safety for the expected T -step constraint-cost in a simulation of the Known-state CMDP, then it also yields d -safety for the expected asymptotic constraint-cost in the real Known-state CMDP. The converse does not follow automatically; if a policy π yields d -safety for the expected asymptotic constraint-cost in a real Known-state CMDP, then $d - 2\epsilon$ -safety is not guaranteed for the T -step expected constraint-cost in a simulation of the Known-state CMDP – making it possible that the constraint cannot be satisfied. The following lemma provides the conditions on the real-world CMDP under which the $d - 2\epsilon$ -safety constraint can be satisfied in simulation.

Lemma 11 *Simulated Budget Satisfaction Lemma* Let M be a full CMDP, $M_{\mathcal{K}}$ be a known-state CMDP induced over a set of known states \mathcal{K} in M , $s \in \mathcal{K}$, π be a policy that satisfies $\pi \in \Pi_c(s, l - \epsilon)$ over M , where $l = d - 2\epsilon$. Further, let $\hat{M}_{\mathcal{K}}$ be an estimation of $M_{\mathcal{K}}$. Then π will satisfy $\hat{C}_\pi(s, T|\hat{M}_{\mathcal{K}}) \leq l$.

Proof Note that $C_\pi(s|M_{\mathcal{K}}) \leq C_\pi(s)$ due to constraint-costs being positive. Therefore, by the Constrained T-step Estimation Lemma (Lemma 1) and the Constrained Simulation Lemma (Lemma 2),

$$\hat{C}_\pi(s, T|\hat{M}_{\mathcal{K}}) \leq \hat{C}_\pi(s|\hat{M}_{\mathcal{K}}) \leq C_\pi(s|M_{\mathcal{K}}) + \epsilon \leq C_\pi(s) + \epsilon \leq l - \epsilon + \epsilon = l.$$

Thereby, $\hat{C}_\pi(s, T|\hat{M}_{\mathcal{K}}) \leq l$. □

3.8 Proof of Theorem 1

With all the lemmas in place, the following presents the step-by-step proof of Theorem 1.

(a) Returning a near-optimal constrained policy within polynomial time First, we prove the statement that for $\epsilon \geq 0$, E^4 outputs with probability at least $1 - \delta$ a near-optimal constrained policy π for s with $V_\pi(s) \geq V_{\pi^*}(s) - \epsilon$ and $C_\pi(s) \leq d$ with sample complexity and time complexity that is polynomial in $1/\epsilon$, $1/\delta$, S , the horizon time $1/(1 - \gamma)$, r_{\max} , and C_{\max} .

Select $s \in \mathcal{S}$ and $\epsilon \geq 0$ arbitrarily. Let m_{known} be $O((STG/\epsilon)^4 \text{Var}_{\max} \ln(1/\delta'))$, where $\delta' = \delta/4$. Define T as the $\epsilon/4$ -horizon time. If $s \in \mathcal{K}$, then by assumption $\Pi_c(s, d, T) = \{\pi \in \Pi : C_\pi(s, T) \leq d\} \neq \emptyset$. By Lemma 6, either there exists a near-optimal constraint-satisfying policy $\pi \in \Pi_c(s, d, T)$ from s with $V_\pi(s, T|M_{\mathcal{K}}) \geq V_{\pi^*}(s, T) - \epsilon/2$ or there exists a policy $\pi' \in \Pi_c(s, d, T)$ that finds an unknown state with probability at least $\epsilon/(2G'_{\max}(T))$, regardless of the choice of T and ϵ . If $\hat{V}_\pi(s, T|\hat{M}_{\mathcal{K}}) \geq V_{\pi^*}(s) - \epsilon/2$, then due to Lemmas 5, 1, and 2, it follows with probability $1 - \delta'$ that

$$V_{\pi^*}(s) - \epsilon/2 \leq \hat{V}_{\pi}(s, T|\hat{M}_{\mathcal{H}}) \leq \hat{V}_{\pi}(s, T) \leq \hat{V}_{\pi}(s) \leq V_{\pi}(s) + \epsilon/4$$

and that

$$V_{\pi}(s, T|M_{\mathcal{H}}) \geq \hat{V}_{\pi}(s, T|\hat{M}_{\mathcal{H}}) - \epsilon/4 \geq V_{\pi^*}(s) - 3\epsilon/4 \geq V_{\pi^*}(s, T) - 3\epsilon/4,$$

which implies offline optimisation will find an ϵ -optimal exploitation policy π , and the proof is finished in this case. Otherwise, if $\hat{V}_{\pi'}(s, T|\hat{M}_{\mathcal{H}}) < V_{\pi^*}(s) - \epsilon/2$ this implies there exists an exploration policy that can be found by E^4 since

$$V_{\pi'}(s, T|M_{\mathcal{H}}) \leq \hat{V}_{\pi'}(s, T|\hat{M}_{\mathcal{H}}) + \epsilon/4 < V_{\pi^*}(s) - 3\epsilon/4 \leq V_{\pi^*}(s, T) - \epsilon/2.$$

E^4 starts an exploration attempt using π' , which takes a T -step trajectory to find an unknown state, based on the exploration known-state CMDP. Such attempts may fail repeatedly but have a success probability of at least $\epsilon/(2G_{\max}^r(T))$. Upon success, the algorithm performs balanced wandering in $M_{\mathcal{H}}$. This cycle of attempted explorations is repeated and due to the repeated visitation, states will become known after a number of visitations $m_{\text{known}} = O((STG/\epsilon)^4 \text{Var}_{\max} \log(1/\delta'))$. In the worst case, the algorithm must make all the states known before finding a near-optimal constrained exploitation policy for s , requiring up to Sm_{known} steps of balanced wandering. In the worst case, due to $T' \geq D(M) + 1$ following the requirements of Lemma 10, there is only one step of balanced wandering per successful exploration attempt. With this in mind, Chernoff bound analysis (see Appendix C) shows that with probability $1 - \delta'$, the total number of exploration attempts before making all states known is bounded by $O\left(\frac{G_{\max}^r(T)}{\epsilon} Sm_{\text{known}} \ln(S/\delta')\right)$. Since each T -step trajectory takes at most T actions, the number of actions taken by E^4 before a near-optimal policy can be returned will be bounded by

$$O\left(T \frac{G_{\max}^r(T)}{\epsilon} Sm_{\text{known}} \log(S/\delta')\right). \quad (7)$$

Therefore, with $T \geq \ln\left(\frac{\max(r_{\max}, c_{\max})}{\epsilon(1-\gamma)}\right) \frac{1}{1-\gamma}$ and $m_{\text{known}} = O((STG/\epsilon)^4 \text{Var}_{\max} \log(1/\delta'))$, the sample complexity to output a near-optimal constrained policy for state $s \in \mathcal{S}$ is polynomial in $1/\epsilon$, $1/\delta$, S , the horizon time $1/(1-\gamma)$, r_{\max} and c_{\max} .³ Since offline optimisation is repeated at every attempted exploration, the time complexity of E^4 is

$$O\left(\text{Opt } T \frac{G_{\max}^r(T)}{\epsilon} Sm_{\text{known}} \log(S/\delta')\right), \quad (8)$$

where Opt refers to the time complexity of the offline optimisation. Given the above, all that remains to be shown is that 1) Opt is polynomial-time; and 2) offline optimisation converges to the global optimum. These statements are demonstrated for different offline optimisation algorithms in Sect. 4.1. Summing the three observed failure probabilities, the above results hold with probability at least $1 - \delta$. \square

(b) Safe exploration Here, we prove the statement that at any time t , the non-stationary policy $\pi^n = \{\pi_i\}_{i=1}^{\infty}$ induced by the exploration process of E^4 yields $C_{\pi^n}(s_t) \leq d$, as specified in the CMDP objective, with probability at least $1 - (UA\delta_{\psi} + \delta)$.

³ The dependencies on r_{\max} and c_{\max} follow from the factors $G_{\max}^r(T)$, G^4 , and $\text{Var}_{\max} \leq \max(c_{\max}, r_{\max})^2/4$.

Select an arbitrary time point t in the RL agent's lifetime and observe the state s_t , let $\epsilon > 0$, and let d'' be the known-state budget, and let $l = d'' - 2\epsilon$. Moreover, it is assumed that $\Pi_c(s_t, l - \epsilon, T) = \{\pi \in \Pi : C_\pi(s_t, T) \leq l - \epsilon\} \neq \emptyset$, $\Pi_c(s_t, d_s - \epsilon, T) = \{\pi \in \Pi : C_\pi(s_t, T) \leq d_s - \epsilon\} \neq \emptyset$ and that the optimal value within $\Pi_c(s_t, l - \epsilon, T)$ is ϵ -close to that of the optimal value within $\Pi_c(s_t, d, T)$.⁴

If $s_t \in \mathcal{K}$, then either the agent is performing d_s -safe return or the agent is doing exploration/exploitation. If the agent is at the starting state of its d_s -safe return then the proof is finished by definition of d_s (see Definition 6). If not, the exploration or exploitation policy will be activated after fewer than T steps of d_s -safe return. The setting $\Pi_c(s_t, l - \epsilon, T) = \emptyset$ implies that the offline optimisation constraints are satisfiable on level l in $\hat{M}_{\mathcal{K}}$ with probability $1 - \delta'$. By Lemma 7, any policy π with $\hat{C}_\pi(s_t, T|\hat{M}_{\mathcal{K}}) \leq l$ yields $C_\pi(s_t|M_{\mathcal{K}}) \leq d''$ as required. By Lemma 6 either the exploitation policy $\pi' \in \Pi_c(s_t, d'', T)$ exists or the exploration policy $\pi'' \in \Pi_c(s_t, d'', T)$ exists (or both). By assumption, the optimal value within $\Pi_c(s_t, d, T)$ is ϵ -close to that of the optimal value within $\Pi_c(s_t, l - \epsilon, T)$ and consequently that in $\Pi_c(s_t, d'', T)$ as well. Combining the above, either a) E^4 will be able to find a policy $\pi' \in \Pi_c(s_t, d'', T)$ that is near-optimal from s_t , or b) E^4 finds a policy π'' which has $C_{\pi''}(s_t, T|M_{\mathcal{K}}) \leq d''$ and which with probability $p > \epsilon/G'_{\max}(T)$ reaches an unknown state in T steps. If the agent stays within \mathcal{K} , then the proof for the known states is complete since then $C_\pi(s_t) = C_\pi(s_t|M_{\mathcal{K}}) \leq d'' \leq d$. Otherwise, by Lemma 10, the setting of d' implies that the non-stationary policy π^n , which alternates between l -safe exploration/exploitation, d' -safe balanced wandering, and d_s -safe return, yields $C_{\pi^n}(s_t) \leq d$ – provided that the agent does not exceed the budget d' in the unknown states (see following paragraph).

If $s_t \in \mathcal{U}$, then the E^4 agent will use the safe balanced wandering policy μ until returning to the known states. Let \mathcal{P} be an uncertainty set such that 1) for all $(s, a) \in \mathcal{S} \times \mathcal{A}$, $\mathcal{P}_{s,a}$ contains $P_{s,a}^*$ with probability at least $1 - \delta_\psi$, and 2) its worst-transition diameter is at most $T' - 1$ (see Eq. 20). From the definition of $T' = \inf\{T \in \mathbb{N} : \sum_{t=0}^{T-1} \gamma^t c_{\max} \geq d'\}$, there exists a policy π'' (the escape policy) that satisfies

$$\max_{P \in \mathcal{P}} \hat{C}_{P, \pi''}(s, T|\hat{M}_{\mathcal{U}}) \leq d'.$$

By Lemma 8, the above implies $C_{P^*, \pi''}(s, T|M_{\mathcal{U}}) \leq \max_{P \in \mathcal{P}} \hat{C}_{P, \pi''}(s, T|\hat{M}_{\mathcal{U}}) \leq d'$ with probability at least $1 - UA\delta_\psi$. Now consider the trajectory within the unknown states. μ first performs steps of balanced wandering until taking one more step would not allow the escape policy π'' to return safely, i.e. as soon as $C(p_t) + \gamma^t \max_{P \in \mathcal{P}} \max_{a_t \in \mathcal{A}} \sum_{s_{t+1} \in \mathcal{S}} P_{s_t, a_t}(s_{t+1}) \hat{C}_{P, \pi''}(s_{t+1}, T|\hat{M}_{\mathcal{U}}) \geq d' - c_{\max}$. By the proof of Lemma 9, the total cost within the unknown states is

$$\begin{aligned} c(s_t) + \gamma^t C_{P^*, \pi}(s, T|M_{\mathcal{U}}) &\leq c(s_t) + \gamma^t C_\mu(s_t, T|M_{\mathcal{U}}) \\ &\leq C(p_t) + \gamma^t C_\mu(s_t, T|M_{\mathcal{U}}) \leq d' \end{aligned}$$

with probability at least $1 - UA\delta_\psi$.

This constraint-satisfaction implies that $C_{\pi^n}(s_t) \leq d$ when exploring or exploiting from $s_t \in \mathcal{K}$. Similarly, due to requirement (c) in Lemma 10, it follows that for $s_t \in \mathcal{U}$,

⁴ Note that the last assumption is not needed for safety throughout exploration; instead its purpose is to be able to return an ϵ -optimal exploitation policy despite the difference between d'' and d . If the assumption does not hold on the initial known-state budget d'' , one potential way to realise the assumption practically is to reset d'' to be close to d after most states have been made known, thereby returning the exploitation policy only when the known-state budget is sufficiently close to d .

$C_{\pi^n}(s_t) \leq d$. Further, the safe return policy may fail to stay within \mathcal{K} with probability δ' . Therefore, combining all the previous yields $C_{\pi^n}(s_t) \leq d$ with probability at least $1 - (UA\delta_\psi + \delta)$. \square

4 Explicit Explore, Exploit, or Escape algorithm

With the above theory in mind, E^4 is now further developed as a practical algorithmic framework, with a discussion of different prototype implementations. Within the main loop of the algorithm, three different policies are optimised: an exploration policy, an exploitation policy, and an escape policy. We discuss a range of offline methods to optimise these policies, including policy gradient, linear programming and dynamic programming approaches (Sect. 4.1). The general flow of how these policies fit together is then discussed in Sect. 4.2, with the complete algorithm being summarised in Algorithm 1. The algorithm and theory relies on two conditions for robust constraint-satisfiability, namely that the diameter of the CMDP must be limited and the availability of tight uncertainty sets that still capture the true transition dynamics with high probability. These conditions are discussed along with a variety of example uncertainty sets (Sect. 4.3).

4.1 Offline optimisation

The algorithm starts by performing offline optimisation of three different policies, the exploration policy, the exploitation policy, and the escape policy, each of which are defined over a different induced CMDP. The remainder of this subsection provides four possible approaches for offline optimisation in E^4 , comparing their time complexity, applicability, and scalability. Using one of these algorithms allows fulfilling the two remaining conditions on offline optimisation in Sect. 3, namely 1) the algorithm is polynomial-time; and 2) the algorithm converges to the global optimum. Therefore, using these algorithms for E^4 yields near-optimal constrained policies in polynomial time.

4.1.1 Robust-constrained policy gradient

A first proposed approach to optimising these policies is based on an existing solution, namely the RCM DP policy gradient (Russel et al., 2020), where distributional robustness is integrated into CMDPs. This is proposed for the Worst-case Escape CMDP but can optionally also be used for the Known-state CMDPs (see Algorithm 1). As Russel et al. (2020) note, one can incorporate the argmax over either the value function or the constraint-cost and they choose the value function; in E^4 , safety is the main concern, and therefore robustness is incorporated by an argmax over the constraint-cost. The robust CMDP objective is solved by computing the saddle point of the Lagrangian for a given budget d ,

$$\min_{\lambda \geq 0} \max_{\pi_\theta} L(\lambda, \pi_\theta) = \hat{V}_{\pi_\theta, P}(s, T | \hat{M}_{\mathcal{J}}) - \lambda (d - \hat{C}_{\pi_\theta, P}(s, T | \hat{M}_{\mathcal{J}})), \quad (9)$$

where $\hat{M}_{\mathcal{J}}$ is an induced CMDP over $\mathcal{J} \subset \mathcal{S}$, \mathcal{P} is the uncertainty set, and $P = \arg \max_{P' \in \mathcal{P}} \hat{C}_{\pi_\theta, P'}(s, T | \hat{M}_{\mathcal{J}})$. Substituting $f(\theta) = \hat{V}_{\pi_\theta, P}(s | \hat{M}_{\mathcal{J}})$ and $g(\theta) = d - \hat{C}_{\pi_\theta, P}(s, T | \hat{M}_{\mathcal{J}})$, the aim is to find the policy π_θ such that the gradient is a null-vector; that is,

$$\nabla_{\theta} L(\lambda, \pi_{\theta}) = \frac{df(\theta)}{d\theta} - \lambda \frac{dg(\theta)}{d\theta} = \mathbf{0} \quad (10)$$

and

$$\nabla_{\lambda} L(\lambda, \pi_{\theta}) = g(\theta) = \mathbf{0}. \quad (11)$$

To optimise the above objective, sampling of limited-step trajectories is repeated for a large number of independent iterations starting from a randomly selected state in the subset of the state space over which the CMDP is induced (see Algorithm 2). Based on the large number of trajectories collected, one then performs gradient descent in λ and gradient ascent in θ . A benefit of this approach is scalability to a high number of parameters.

A common wisdom is that gradient descent methods do not perform well in multi-modal and non-convex landscapes, where they may get stuck in local optima. Similarly, the theory behind the RCMDP policy gradient only provides a convergence result for a local optimum (Russel et al., 2021, 2020). However, with sufficient over-parametrisation – and more precisely, with a number of neurons n_h polynomial in the number of samples n_{ε} , number of hidden layers l , and the (inverse of) the maximal distance between data points v^{-1} – a neural network can learn the global optimum to arbitrary accuracy within polynomial number of samples (Allen-Zhu et al., 2019). In addition to general results for the L_2 -norm, Allen-Zhu et al. (2019) provides similar results for arbitrary Lipschitz-continuous loss functions; one such result states that if the loss function is σ -gradient dominant (see Zhou and Liang (2017) for its definition), then gradient descent finds with probability $1 - e^{-\Omega(\ln^2(n_{\varepsilon}))}$ an ϵ -optimal parameter vector within a number of iterations I polynomial in n_h , l , v^{-1} , σ^{-1} , and ϵ^{-1} . Per iteration, all operations have polynomial time complexity: look-ups, list appending, random number generation, gradient computation, division, and the forward pass of a neural network. Therefore, the resulting algorithm has overall polynomial time complexity and converges with high probability to the global optimum.

4.1.2 Robust linear programming

A second and novel approach is a robust variant of linear programming (LP) for RCMDPs, extending the traditional LP framework of Altman (1998). Note that Zheng and Ratliff (2020) have previously used a “robust” version of LP in the context of UCRL, which estimates an upper confidence bound on the cost. This section proposes a similar variant of robust linear programming, with the key difference that 1) our version is discounted, non-episodic for an extension of E^3 whereas their version is based on an undiscounted, episodic framework for an extension of UCRL; 2) separate exploration, exploitation, and escape policies are optimised; and 3) the error term in the budget (right hand side) as opposed to the constraint-cost, although this is merely a superficial difference. A downside of this approach is that similar to Zheng and Ratliff (2020), this formulation does not account for uncertainty in the transition dynamics. The agent may therefore get stuck in unknown states longer than anticipated.

The approach uses the occupation measure (Altman, 1999), which can be interpreted as the total proportion of discounted time spent in a particular state-action pair. The occupation measure allows to formulate the asymptotic constraint-cost as a simple weighted sum of its immediate cost. For a CMDP $\hat{M}_{\mathcal{S}}$, a policy π , and a state-action pair (s, a) , the T -step estimate of the occupation measure is defined as

$$f(s, a) = (1 - \gamma) \sum_{t=0}^T \gamma^t \mathbb{P}_{\pi}(s, a | \hat{M}_{\mathcal{S}}), \quad (12)$$

leading to the definition of a policy as

$$\pi(a|s) = \frac{f(s, a)}{\sum_{a' \in \mathcal{A}} f(s, a')}. \quad (13)$$

Now the state-action value and constraint-cost functions can be expressed in terms of f_{π} , namely

$$\hat{V}_{\pi}(s, a, T | \hat{M}_{\mathcal{S}}) = f(s, a) \hat{r}(s, a | \hat{M}_{\mathcal{S}})$$

and

$$\hat{C}_{\pi}(s, a, T | \hat{M}_{\mathcal{S}}) = f(s, a | \hat{M}_{\mathcal{S}}) \hat{c}(s, a | \hat{M}_{\mathcal{S}}).$$

For the exploitation and exploration policies in E^4 , this yields the linear programming problem

$$\max_f f^T \hat{r} \quad \text{s.t.} \quad f^T \hat{c} \leq l, \quad (14)$$

where $f^T \in \mathbb{R}^{SA}$ and $\hat{r} \in \mathbb{R}^{SA}$. For a Worst-case Escape CMDP, the robust LP is not recommended due to requiring known transition dynamics but if the constraint-cost estimate is set to c_{\max} , one may optimise the linear programming problem

$$\max_f f^T \hat{r} \quad \text{s.t.} \quad f^T c_{\max} \leq d'. \quad (15)$$

Various general polynomial-time algorithms exist for linear programming, including the ellipsoid method (Khachiyan, 1979) and the projective method (Karmarkar, 1984), which have $O(n^6 L)$ and $O(n^{3.5} L)$ on $O(L)$ digit numbers. Following Karmarkar (1984), interior point methods have been further developed, providing ϵ -optimal guarantees with much improved time complexity (see Potra and Wright (2000) for a variety of algorithms).

4.1.3 Dynamic programming approaches

Dynamic programming (DP) as used in the original E^3 provides convergence to the optimal value and a time complexity of $O(S^2 T)$. In E^4 , DP cannot be directly applied due to the robust constrained setting. There is currently no suitable algorithm for constrained-robust DP. We propose two possible approaches but leave a full analysis for further research. The first approach starts with constrained DP (Altman, 1998) and then incorporates the uncertainty set while the second approach starts with robust DP (Nilim & Ghaoui, 2005; Iyengar, 2005) and incorporates constraints by reformulating the CMDP as a Lagrangian MDP (Taleghan & Dietterich, 2018).

Dual linear program The technique by Altman (1998) goes as follows when applied to a CMDP $\hat{M}_{\mathcal{S}}$, which has transition dynamics \hat{P} , constraint-cost function \hat{c} , and reward function \hat{r} . Note that the solution to the unconstrained DP problem

$$\phi(s) = \max_{a \in \mathcal{A}} \left[(1 - \gamma) \hat{r}(s, a) + \gamma \sum_{s' \in \mathcal{S}} \hat{P}_{s,a}(s') \phi(s') \right] \quad \forall s \in \mathcal{S}.$$

can be rewritten as an LP of the form

$$\begin{aligned} & \max \phi(s) \\ \text{s.t.} \quad & \phi(s) \geq (1 - \gamma) \hat{r}(s, a) + \gamma \sum_{s' \in \mathcal{S}} \hat{P}_{s,a}(s') \phi(s') \quad \forall s \in \mathcal{S} \quad \forall a \in \mathcal{A}. \end{aligned} \quad (16)$$

For constrained DP, the value can be defined based on the min-max of the Lagrangian

$$\hat{V}_{\pi_\theta}(s, T | \hat{M}_{\mathcal{S}}) = \min_{\lambda \geq 0} \max_{\pi_\theta} L(\lambda, \pi_\theta),$$

where min and max can be interchanged as strong duality holds without requiring Slater's condition. Representing $L(\lambda, \pi_\theta)$ in terms of the immediate rewards and constraint-cost results in the LP

$$\begin{aligned} & \hat{V}_{\pi_\theta}(s, T | \hat{M}_{\mathcal{S}}) = \max_{\lambda, \phi} \phi(s) - \lambda d \\ \text{s.t.} \quad & \phi(s) \geq (1 - \gamma)(\hat{r}(s, a) + \lambda \hat{c}(s, a)) + \gamma \sum_{s' \in \mathcal{S}} \hat{P}_{s,a}(s') \phi(s') \quad \forall s \in \mathcal{S} \quad \forall a \in \mathcal{A}. \end{aligned} \quad (17)$$

In E^4 , the transition dynamics are uncertain. Adding the constraints for all transition dynamics in the uncertainty set and using the estimators yields the LP

$$\begin{aligned} & \hat{V}_{\pi_\theta, \mathcal{P}}(s, T | \hat{M}_{\mathcal{S}}) = \max_{\lambda, \phi} \phi(s) - \lambda d \\ \text{s.t.} \quad & \phi(s) \geq (1 - \gamma)(\hat{r}(s, a) + \lambda \hat{c}(s, a)) + \gamma \sum_{s' \in \mathcal{S}} P_{s,a}(s') \phi(s') \quad \forall s \in \mathcal{S} \quad \forall a \in \mathcal{A} \\ & \forall P \in \mathcal{P}. \end{aligned} \quad (18)$$

Robust dynamic programming with Lagrangian MDP Another proposed approach for dynamic programming in E^4 is to apply robust DP (Nilim & Ghaoui, 2005; Iyengar, 2005) to a Lagrangian MDP (Taleghan & Dietterich, 2018). Robust DP models the value as

$$\hat{V}(s, T | \hat{M}_{\mathcal{S}}) = \max_{a \in \mathcal{A}} \hat{r}(s, a) + \gamma \max_{P \in \mathcal{P}_{s,a}} P^T \hat{V}(\cdot, T | \hat{M}_{\mathcal{S}}), \quad (19)$$

for worst-case transition $P \in \Delta^{\mathcal{S}}$ and the values for each following state $\hat{V}(\cdot, T | \hat{M}_{\mathcal{S}}) \in \mathbb{R}^{\mathcal{S}}$. To solve the “inner problem”, $\max_{P \in \mathcal{P}_{s,a}} P^T \hat{V}(\cdot, T | \hat{M}_{\mathcal{S}})$, one uses a bisection algorithm, yielding time complexity $O(S \log(G/\epsilon_s))$, where G is an upper bound to the value function and ϵ_s is the desired accuracy of the approximation. The overall problem (Eq. 19) gives ϵ -optimal guarantees in polynomial time complexity $O(TS^2 \log(1/\epsilon)^2)$, adding only $O(\log(1/\epsilon)^2)$ time cost compared to traditional DP.

To make robust DP work for E^4 , one can then construct a Lagrangian MDP which redefines the reward function as a linear combination of the original reward function and the constraint-cost function, such that the value is similar to the Lagrangian in Eq. 9.

A downside of using robust DP is the requirement of (s, a) -rectangular uncertainty sets (see Sect. 4.3.2). More general uncertainty sets can be converted into

(s, a) -rectangular format by projection onto a larger rectangular subspace (Nilim & Ghaoui, 2005) but this will result in a higher upper bound on the worst-case constraint-cost.

4.2 Exploration, exploitation and escape policies

The three different kinds of policies are optimised over different CMDPs introduced earlier in Sect. 3. This section discusses how these policies are used, what they represent, and when they are activated.

Due to the Constrained Exploit-or-Exploit Lemma (Lemma 6), for any satisfiable budget l and any $\epsilon \geq 0$, from a starting state s , either a) there exists a policy $\pi \in M_{\mathcal{S}}$ for which $V_{\pi}(s, T | M_{\mathcal{K}}) \geq V_{\pi^*}(s, T) - \epsilon$, where $\pi^* = \arg \max_{\pi} V_{\pi}(s)$ s.t. $C_{\pi}(s) \leq l$ for all $s \in \mathcal{S}$; or b) there exists a policy π in $M_{\mathcal{K}}$ which reaches the terminal state s_0 in $\mathcal{S} \setminus \mathcal{K}$ in at most T steps with probability $p > \epsilon / G_{\max}^r(T)$. The estimated exploitation and exploration policies represent case a) and b), respectively, performing l -safe exploitation and l -safe exploration. The exploitation policy is ideally activated whenever it is known that $\hat{V}_{\pi}(s | \hat{M}_{\mathcal{K}}) \geq V_{\pi^*}(s) - \epsilon$. In practice, this knowledge is often not available. However, one may use a strategy similar to what has been proposed for the exploitation and exploration policy in E^3 (Kearns & Singh, 2002); that is, one activates the exploration policy first and continues it as long as $p > \epsilon / G_{\max}^r(T)$ remains likely, and then, as soon as $p < \epsilon / G_{\max}^r(T)$ with high probability, one activates the exploitation policy which is then guaranteed to exist by Lemma 6.

With the above in mind, the different policies and their corresponding induced CMDPs can now be discussed. First, the l -safe exploitation policy has the aim of solving a CMDP induced over the known states $M_{\mathcal{K}}$, differing from M only in the sense that: 1) it terminates with $r(s_0) = 0$ and $c(s_0) = 0$ as soon as it reaches the terminal state s_0 , which is when an unknown state in \mathcal{U} is entered; and 2) the allowed constraint-cost budget is more limited, namely it is set to $l = d'' - 2\epsilon$ following Lemma 7. Second, the l -safe exploration policy has the aim of solving a CMDP induced over the known states $M'_{\mathcal{K}}$. $M'_{\mathcal{K}}$ differs from $M_{\mathcal{K}}$ in the sense that it terminates with $r(s_0) = r_{\max}$ as soon as it reaches the terminal state s_0 , which is when an unknown state in \mathcal{U} is entered, and that it receives $r(s) = 0$ for all $s \in \mathcal{K}$. Third, the d' -safe escape policy has the aim of solving a CMDP induced over the unknown states $M_{\mathcal{U}}$. It differs from M in the sense that: a) it terminates with $r(s_0) = 0$ and $c(s_0) = 0$ as soon as it reaches the terminal state s_0 , which is when a known state in \mathcal{K} is entered; and b) the allowed constraint-cost budget d' is set according to Lemma 10. Finally, to ensure constraint-satisfaction, Lemma 10 requires an additional T -step trajectory of safe return within the known states, a trajectory which yields a cumulative constraint-cost of at most $d_s - \epsilon$. Such a safe return policy can be similarly formulated as a CMDP (e.g. by rewarding 1 for known states and 0 for unknown states) but is often readily available from domain knowledge; for example, domains such as those in Example 2 include an action which makes the agent stay in a particular state or set of states for extended periods of time.

Before the main loop, $\epsilon > 0$ is chosen and the ϵ -horizon time $T \leftarrow \frac{1}{1-\gamma} \ln\left(\frac{\max(r_{\max}, c_{\max})}{\epsilon(1-\gamma)}\right)$ and safety budget $l \leftarrow d'' - 2\epsilon$ are chosen accordingly. At the start of each iteration of the main loop, exploration and exploitation policies are optimised. Via Lemma 7, l -safe exploration or exploitation policies will with high probability be constraint-satisfying despite the T -step estimated constraint-cost in $\hat{M}_{\mathcal{K}}$ differing from the asymptotic true constraint-cost in $M_{\mathcal{K}}$. Once an unknown state is visited, d' and T' are set according to Lemma 10. As long as the different policies satisfy their respective constraint-cost budget within their induced

CMDP, l for the exploration and exploitation policies and d' for the escape policy, then the overall non-stationary policy applying these policies sequentially will satisfy the constraint-cost in the full CMDP M .

Algorithm 1 E⁴ algorithm: Explore or Exploit-Explore-Exploit-or-Escape for CMDPs.

Require: starting set of known states \mathcal{K} , known-state budget d'' , safe return budget d_s , learning rate schedules η_1 and η_2 , and $\epsilon > 0$.

```

1: Start from random  $s \in \mathcal{K}$ .
2: Define  $l \leftarrow d'' - 2\epsilon$ .  $\triangleright l$ -safe exploration or exploitation
3:  $T = \frac{1}{1-\gamma} \ln\left(\frac{\max(r_{\max}, c_{\max})}{\epsilon(1-\gamma)}\right)$ .  $\triangleright \epsilon$ -horizon time
4: while true do
5:    $\triangleright$  perform off-line optimisations (see Algorithm 2)
6:    $\triangleright$  Exploitation CMDP  $\hat{M}_{\mathcal{K}}$ 
7:    $\hat{\pi} \leftarrow \text{OFFLINE-OPTIMISATION}(\hat{\pi}, \hat{M}_{\mathcal{K}}, \mathcal{P}, l, T)$ .
8:    $\triangleright$  Exploration CMDP  $\hat{M}'_{\mathcal{K}}$ 
9:    $\hat{\pi}' \leftarrow \text{OFFLINE-OPTIMISATION}(\hat{\pi}', \hat{M}'_{\mathcal{K}}, \mathcal{P}, l, T)$ .
10:  if  $\hat{V}_{\hat{\pi}}(s, T | \hat{M}_{\mathcal{K}}) \geq V_{\pi^*}(s) - \epsilon$  then  $\triangleright$  Idealised ( $\pi^*$  is unknown)
11:     $\triangleright$  in relaxed version, explore first, then exploit
12:     $\pi \leftarrow \hat{\pi}$ .  $\triangleright$  exploit
13:  else
14:     $\pi \leftarrow \hat{\pi}'$ .  $\triangleright$  explore
15:  end if
16:   $\triangleright$  explore or exploit
17:  for  $t = 0, \dots, T-1$  do
18:     $a \leftarrow \pi(s)$ .  $\triangleright$  take action
19:    Sample reward  $r$  and constraint-cost  $c$  based on  $(s, a)$ .
20:     $s \sim P_{s,a}^*$ .  $\triangleright$  sample next state
21:    if  $s \in \mathcal{U}$  then
22:      Set budget  $d'$ .  $\triangleright$  Lemma 10
23:       $T' \leftarrow \inf\{T \in \mathbb{N} : \sum_{t=0}^T \gamma^t c_{\max} \geq d'\}$ .
24:       $\hat{\pi}'' \leftarrow \text{OFFLINE-OPTIMISATION}(\hat{\pi}'', \hat{M}_{\mathcal{U}}, \mathcal{P}, d', 2T')$ .
25:      break.
26:    end if
27:  end for
28:   $\triangleright$  balanced wandering with escape (see Lemma 8, 9)
29:   $\pi \leftarrow \hat{\pi}''$ .  $\triangleright$  Worst-case Escape policy
30:  wandering  $\leftarrow$  true  $\triangleright$  start with balanced wandering
31:  Initialise path cost  $C_p \leftarrow 0$ .
32:  for  $t = 0, \dots, T'-1$  do
33:    if wandering then
34:       $a \leftarrow \arg \min_{a' \in \mathcal{A}} n(s, a')$ .  $\triangleright$  balanced wandering
35:    else
36:       $a \leftarrow \pi(s)$ .  $\triangleright$  take action from escape policy
37:    end if
38:     $n(s, a) += 1$ .  $\triangleright$  count the visitation
39:    if  $n(s, a) = m_{\text{known}}$  then
40:       $\mathcal{K} \leftarrow \mathcal{K} \cup \{s\}$ .  $\triangleright$  add  $s$  to known states
41:    end if
42:    if  $s \in \mathcal{K}$  then
43:      break.
44:    end if
45:    Sample reward  $r$  and constraint-cost  $c$  based on  $(s, a)$ .
46:     $s \sim P_{s,a}^*$ .  $\triangleright$  sample next state
47:     $C_p \leftarrow C_p + \gamma^t c$ .  $\triangleright$  constraint-cost of past and future
48:     $\hat{C} \leftarrow \max_{P \in \mathcal{P}} \max_{a \in \mathcal{A}} \sum_{s \in \mathcal{S}} P_{s,a}(s') \hat{C}_{P,\pi}(s', T | \hat{M}_{\mathcal{U}})$ .
49:    if  $C_p + \gamma^{t+1} \hat{C} \geq d' - c_{\max}$  then wandering  $\leftarrow$  false.
50:    end if
51:  end for
52:  Perform safe return policy  $\pi_{d_s}$  for  $T$  steps within  $\mathcal{K}$ .
53: end while

```

Algorithm 2 Offline-Optimisation of RCMDPs

```

1: procedure OFFLINE-OPTIMISATION(Policy  $\pi_\theta$ , Induced
   CMDP  $\hat{M}_{\mathcal{S}}$ , uncertainty set  $\mathcal{P}$ , constraint-cost budget  $d$ ,
   trajectory length  $T$ , learning rate schedules  $\eta_1$  and  $\eta_2$ )
2:   Trajectories  $\xi = []$ .
3:   for  $i = 1, 2, \dots, I$  do ▷ independent iterations
4:     Randomly select starting state  $s \in \mathcal{S}$ .
5:     ▷ Simulate worst-case trajectory
6:     for  $t = 0, 1, \dots, T - 1$  do
7:       if  $s$  is terminal then
8:         break
9:       end if
10:      Action  $a \leftarrow \pi_\theta(s)$ .
11:      ▷ deterministic cost and reward in induced CMDP
12:       $c \leftarrow \hat{c}(s, a | \hat{M}_{\mathcal{S}})$ .  $r \leftarrow \hat{r}(s, a | \hat{M}_{\mathcal{S}})$ .
13:       $P \leftarrow \arg \max_{P' \in \mathcal{P}} \sum_{s' \in \mathcal{S}} P'_{s,a}(s') \hat{C}_{\pi_\theta, P'}(s', T | \hat{M}_{\mathcal{S}})$ .
14:      State  $s' \sim P_{s,a}$ . ▷ Worst-case transition
15:      Gradient  $\nabla \leftarrow \frac{\nabla_\theta \pi_\theta(a|s)}{\pi_\theta(a|s)}$ .
16:      Trajectory  $\xi \leftarrow \xi + [s, a, c, r, s']$ .
17:       $s \leftarrow s'$ .
18:   end for
19:    $T_{\text{stop}} \leftarrow t$ .
20:   ▷ Optimise the policy
21:   Initialise  $V \leftarrow 0$ .
22:   Initialise  $C \leftarrow 0$ .
23:   for  $t = T_{\text{stop}} - 1, T_{\text{stop}} - 2, \dots, 0$  do
24:      $V \leftarrow r_t + \gamma V$  ▷ Objective:  $\max_\theta V(\theta)$  s.t.  $C(\theta) \leq d$ 
25:      $C \leftarrow c_t + \gamma V$  ▷  $L(\theta) = V(\theta) - \lambda(C(\theta) - d)$ 
26:      $\theta \leftarrow \theta + \eta_1(k) * (V - \lambda C) \nabla_t$  ▷  $dL/d\theta$ 
27:      $\lambda \leftarrow \lambda + \eta_2(k) * (d - C)$  ▷  $dL/d\lambda$ 
28:   end for
29: end for
30: return  $\pi_\theta$ 
31: end procedure

```

4.3 Robust constraint-satisfiability

For E^4 to work as intended, the Worst-case Escape CMDP must have robust constraint-satisfiability. That is, when selecting the worst-case transitions from an uncertainty set \mathcal{P} , there must be a policy π that satisfies $\max_{P \in \mathcal{P}} \hat{C}_{P,\pi}(s, T | \hat{M}_{\mathcal{U}}) \leq d'$. This depends critically on two factors. First, constraint-satisfiability depends on the CMDP of interest. The parameter of the CMDP that we analyse here is the diameter; if this parameter is limited, then the constraints are satisfiable; this holds true in simulation as well as in the real world CMDP. Second, when performing offline optimisation, robust constraint-satisfiability further depends on the uncertainty set. If the uncertainty set is too broad or does not include the true transition dynamics, then constraint-satisfiability cannot be guaranteed unless for trivial constraints (e.g. where every path goes to the known states before the budget d' is exceeded). Because the transition dynamics

model is not known in the unknown states, we discuss how the uncertainty set can be formed based on prior knowledge and statistical inference.

4.3.1 The diameter of the CMDP

For constraint-satisfiability in the Worst-case Escape CMDP $M_{\mathcal{U}}$, the diameter must satisfy $D(M_{\mathcal{U}}) \leq T' = \inf\{T \in \mathbb{N} : \sum_{t=1}^T \gamma^{t-1} c_{\max} \geq d'\}$. In this case, the policy can escape back to \mathcal{K} within at most T' steps, before potentially exceeding the budget (on expectation). To ensure that at least one step of balanced wandering can be taken per attempted exploration, in line with the theory in Sect. 3 and Algorithm 1, the diameter must satisfy $D(M_{\mathcal{U}}) \leq T' - 1$. The robust optimisation further implies that the worst-case diameter in the uncertainty set must also be at most $D(M_{\mathcal{U}}) \leq T' - 1$. This worst-case assumption therefore can be written in terms of the uncertainty set (which may be restricted to the entries with $\mathcal{P}_{s,a}$ where $s \in \mathcal{U}$ for making the assumption weaker):

$$D(\mathcal{P}) := \max_{s \neq s'} \min_{\pi} \max_{P \in \mathcal{P}} \mathbb{E}[W(s'|s, \mathcal{P}, \pi)] \leq T' - 1. \quad (20)$$

Note that this “worst-transition diameter” differs from the worst-case diameter of Garcelon et al. (2020) in that the worst-case is over the uncertain transition dynamics rather than the policy and that the best policy is taken rather than the worst-case policy.

The assumptions on the diameter provide a worst-case guarantee. In practice, many domains (e.g. Example 2) have properties related to reversibility, such that the number of steps taken in unknown steps relates to the number of steps to escape back to the known states. In such cases, the diameter assumption can be significantly relaxed. In addition, more informative metrics than the diameter could potentially improve the budgetary requirements specified in Sect. 10.

4.3.2 Uncertainty sets

For unknown states, no samples are given at the start of the algorithm, implying that uncertainty sets constructed from their state-action visitations are too large to provide safety guarantees and to make efficient use of exploration attempts. This section discusses how to construct narrow uncertainty sets to obtain safe and efficient behaviour within the unknown states. We present existing uncertainty sets, including the (s, a) -rectangular uncertainty sets [(see e.g., (Wiesemann et al., 2013; Russel et al., 2020)], and the factor matrix representation (Goyal & Grand-Clement, 2018), as well as the use of expert knowledge to provide tight uncertainty sets. A gridworld example is then given as an illustration (see Example 2).

(s,a)-rectangular sets The most well-studied class of uncertainty set is the (s,a) -rectangular set, which defines a plausible interval for each $(s, a) \in \mathcal{S} \times \mathcal{A}$. An advantage of (s,a) -rectangular set is that they provide various polynomial-time results for robust optimisation (Wiesemann et al., 2013; Nilim & Ghaoui, 2005). A simple example is the set based on the L1-norm (Russel et al., 2020), which defines

$$\mathcal{P}_{s,a} = \{P \in \Delta^{\mathcal{S}} : \|P - \hat{P}_{s,a}\|_1 \leq \psi_{s,a}\}, \quad (21)$$

for all $(s, a) \in \mathcal{S} \times \mathcal{A}$. Traditionally, (s,a) -rectangular uncertainty sets are formed based on Hoeffding’s inequality, defining the budget $\psi_{s,a} = \sqrt{\frac{2}{n(s,a)} \ln\left(\frac{SA2^S}{1-\delta_\psi}\right)}$ based on the number of visitations of the state-action pairs for failure rate δ_ψ . This will unfortunately not work

when the number of visitations is small (or even zero) in unknown states, since then the uncertainty set is prohibitively broad. However, a useful alternative for the unknown states is the Bayesian Credible Region (Russel & Petrik, 2019), which defines the uncertainty set by finding the tightest budget with low posterior belief of failure,

$$\psi = \min_{\psi'} \{ \psi' \in \mathbb{R}^+ : \mathcal{P} \left(\|P_{s,a}^* - \hat{P}_{s,a}\|_1 > \psi'_{s,a} \right) < \delta_{\psi'} \}. \quad (22)$$

The posterior \mathcal{P} allows injecting prior knowledge via posterior sampling, either analytically (e.g. via Dirichlet prior) or alternatively, via simulation methods (e.g. Monte Carlo Markov Chain sampling methods).

r-rectangular sets based on factor matrix More efficient than (s,a)-rectangular sets are r-rectangular sets based on a factor matrix representation, which allows to efficiently treat different state-action pairs as being correlated. An uncertainty set $\mathcal{P} \subset \mathbb{R}^{S \times A \times S}$ is generated using a factor matrix $W = \{\mathbf{w}_1, \dots, \mathbf{w}_r\}$, a convex, compact subset of $\mathbb{R}^{r \times S}$, and a fixed set of coefficients $\mathbf{u}_{s,a}^i$ for all $i \in \{1, \dots, r\}$ and all $(s, a) \in S \times \mathcal{A}$. Specifically, one has

$$\mathcal{P}_{s,a} = \sum_{i=1}^r u_{sa}^i \mathbf{w}_i \quad (23)$$

for all $(s, a) \in S \times \mathcal{A}$.

The resulting representation is flexible; for example, if $r = SA$ and u_{sa}^i for all $i = 1, \dots, r$ one has the (s,a)-rectangular case. The factor matrix can be estimated via non-negative matrix factorisation (Xu & Yin, 2013) although currently this requires a relatively accurate estimate $\hat{P}_{s,a}$ (Goyal & Grand-Clement, 2018). In addition, factorisation has been used to provide E^3 with scalability to large or continuous state spaces (Henaff, 2019); as a result, factorisation methods seem promising for improving the scalability of E^4 .

Sets based on expert knowledge: action models and local inference When a domain expert has a high-precision model for the transition dynamics of unknown states, much tighter uncertainty sets can be formed, which is especially useful in the early stages of the E^4 algorithm, when the visitations are too few to provide reliable statistics. As an illustration, we consider the case where an agent is located in a discrete state space organised along Cartesian coordinates and its available actions are moving within a local neighbourhood, as is typical in many robotics applications. The expert then formulates a set of probabilistic models, each of which is valid locally in a subset of the state space (e.g. due to the position in the landscape).

The expert first formulates n action models $g_i(s, a)$, $i = 1, \dots, n$, which output the “typical” next state following action a in state s . Then, for each i a probability distribution $P_{g_i, \tau, \mathcal{N}}$ is formed by assigning high probability $1 - \tau$ to $s' = g_i(s, a)$ and probability τ/N to a local neighbourhood $\mathcal{N}(s'|s, a)$ (e.g. all states $s_e \in S$ with $\|s_e - s'\|_2 \leq 2$, representing plausible action outcomes), where N is the size of the neighbourhood. One may additionally iterate the definition of $P_{g_i, \tau, \mathcal{N}}(s, a, s')$ over different neighbourhoods and different error rates if these are uncertain. For simplicity, we illustrate this only the error rate in Example 2, i.e. we fix the neighbourhood and define $\mathcal{P}_{g_i} = \{P_{g_i, \tau, \mathcal{N}} : \tau \leq 0.1\}$.

The resulting uncertainty set, $\{\mathcal{P}_{g_1}, \dots, \mathcal{P}_{g_n}\}$, is typically small when significant domain knowledge exists. However, the size of this uncertainty set can be even further reduced by eliminating models with low probability. For this purpose, one can consider the probability of the recent path p under model g_i , $\mathbb{P}_{g_i}(p)$. Alternatively, the expert defines a transfer probability $\rho(g_j|s, a, g_i)$, reflecting the probability of model g_j

after taking action a in state s for which g_i was valid. Note that the dependencies on s , a , and/or g_i can be dropped in the case of statistical independence.

Example 2 Let S be a discrete set of Cartesian xy -coordinates in a 10-by-10 gridworld ($S = 100$). Let \mathcal{A} be the set of moves within a Von Neumann neighbourhood, with actions {north,west,south,east} moving one step in the corresponding direction and the action stay remaining in the same cell ($A = 5$). The optimal constrained policy is to cycle around the bounds of the gridworld, without hitting the wall. If any wall is hit, the constraint-cost is 1 and the agent remains in the same state; otherwise the constraint-cost is 0. Hitting the wall repeatedly in quick succession, for an asymptotic constraint-cost significantly larger than $d = 10$ (assuming $\gamma = 0.99$), is known to risk damaging the agent. Actions fail stochastically with a probability $\tau = 0.05$. Initially, we have $\mathcal{K} = \{(1, 1), (2, 1), \dots, (9, 1)\}$ and $\mathcal{U} = S \setminus \mathcal{K}$ as known and unknown states, respectively. The agent is aware of being initialised in the south-west wall.

The expert knows there are walls but only knows opposite ends are at least 5 steps away from each other. For the interior of the gridworld, $g_1(s, a) := \text{co}(s) + \text{co}(a)$ where co defines the coordinates for the state and for the action (e.g. one step north is given by $(0, 1)$). For the north, west, south, and east bounds of the gridworld, respectively, the action models $i = 2, \dots, 5$ make the corresponding action a null move (e.g., $g_2(s, \text{north}) := \text{co}(s)$). Each model is error-free, i.e. $\tau = 0$. For hitting the wall, the activation of the partial uncertainty set depends on the action and the previously valid ϕ_i . For example, if $a = \text{east}$, then $\rho(g_5|a) = 1$ (east bound) except when g_3 (west bound) was previously active.

The uncertainty set constructed from the uncertainty subsets, $\{\mathcal{P}_{g_1}, \dots, \mathcal{P}_{g_5}\}$, where $\mathcal{P}_{g_i} = \{P_{g_i, \tau, \mathcal{N}} : \tau \leq 0.1\}$. These subsets are sufficient to model hitting the wall from the different bounds of the gridworld with the given error rate. For example, if $s = (10, 1) \in \mathcal{U}$ and taking the next action east , the worst-case transition within the uncertainty set is hitting the wall with 100% probability ($\tau = 0$), a scenario considered by action model g_5 . For other actions, the worst-case transition within the uncertainty set is hitting the wall with 10% probability ($\tau = 0.10$). These worst-cases over-estimate the constraint-cost when compared to $\tau = 0.05$, which would yield only 5% probability.

The current budget for unknown states is $d' = 5$, allowing for 5 steps in unknown states with $c_{\max} = 1$. Although the diameter is significantly higher, From the first unknown state $(10, 1)$, the agent will behave as follows. Given the recent trajectory, the agent has the south and east bounds as active models $\{g_4, g_5\}$ due to taking more than 5 actions east from the known south-west walls. After offline optimisation iterations on the Worst-case Escape CMDP with uncertainty set $\{P_{g_4}, P_{g_5}\}$, the agent performs balanced wandering for a few steps and then escapes to $(10, 1)$. The agent then performs T actions of stay to ensure safe return.

After much exploration of the gridworld, eventually the agent makes all states known save for a few in the north bound. Now the uncertainty set is narrowed down to $\mathcal{P}_{g_2} = \{P_{g_2, \tau, \mathcal{N}} : \tau \leq 0.1\}$ and moreover, there are different routes to escape. Therefore, once the agent is optimised for this uncertainty set, the agent can make much quicker progress with balanced wandering. Therefore, these remaining states will quickly become known. Once the bounds are known, a near-optimal exploitation policy can now be found for the entire CMDP.

4.4 Practical considerations

Setting m_{known} in practice In general, the required number of samples for a state to be known is $m_{\text{known}} = O\left((STG/\epsilon)^4 \text{Var}_{\max} \ln \frac{1}{\delta}\right)$. In practice, m_{known} needs to be set to a fixed value, which requires uncovering constants hidden by the big O notation as well as replacing asymptotic or other unpractical assumptions with practical assumptions. First, in the Constrained Simulation Lemma, obtaining the constant K_1 in $\alpha = K_1(\epsilon/(STG))^2$ requires solving exactly for the two conditions specified in addition to the big O notation $\alpha = O(\epsilon)$, plugging in exact values of ϵ , T , $G_{\max}^r(T)$, and $G_{\max}^c(T)$. G will be equal to $G_{\max}^r(T)$ or $G_{\max}^c(T)$; since these have analogous requirements, let $G = G_{\max}^r(T)$. To obtain the first requirement, first note that for most practical settings, $\alpha \leq 1$. Then, set $K_1 \leq 1/144$ to obtain

$$(\alpha + 2\beta)TSG_{\max}^r(T) \leq 3\sqrt{\alpha}TSG_{\max}^r(T) \leq \epsilon/4.$$

Additionally assuming $S \geq 4$ yields the second requirement,

$$\Delta = \alpha/\beta = \sqrt{\alpha} \leq \epsilon/(4TG_{\max}^r(T)).$$

Then, note that combining the three conditions following from Hoeffding's inequality in the Known State Lemma while plugging in K_1 and summing over actions, which are modelled as a constant with relatively low value, and substituting $M = \max(\{r_{\max}, c_{\max}, 1\})$ appropriately instead of Var_{\max} (as $\epsilon \rightarrow 0$ does not hold in practice), yields a general formula for m_{known} :

$$m_{\text{known}} \geq \frac{1}{2}A \ln\left(\frac{2}{\delta}\right) M^2 K_1^{-2} (STG/\epsilon)^4. \quad (24)$$

Note that Eq. 24 implies E^4 is particularly suitable for domains with limited state space. Naturally, the failure rate δ should at all times be very limited but for reducing m_{known} one can dynamically change the other parameters, starting with relatively high ϵ and decreasing ϵ later on in the lifetime. This would also decrease T and G , allowing the algorithm to reach "satisfactory" performance levels across a large set of states more quickly.

Given certain assumptions, one may be able to significantly reduce m_{known} by using a different concentration inequality. For example, one may use Bernstein's inequality if the variance of rewards and costs are known, and relatively low compared to both r_{\max} or c_{\max} . For the reward function, define the random variable $X = |r(s, a) - r_t|/m$, where r_t is the observed reward at time t provided $s_t = s$ and $a_t = a$. Then, note that $\sum_{i=1}^m X_i = |r(s, a) - \hat{r}(s, a)|$. Then, if $\text{Var}_{\max} \leq \alpha$, we have

$$\begin{aligned} \mathbb{P}(|r(s, a) - \hat{r}(s, a)| \geq \alpha) &\leq 2e^{-\frac{m\alpha^2}{2(\text{Var}_{\max} + \alpha/3)}} = \delta \\ m &= 2\alpha^{-2}(\text{Var}_{\max} + \alpha/3) \ln(2/\delta) \end{aligned}$$

The requirement for the constraint-cost is completely analogous, defining $Y = |c(s, a) - c_t|/m$. For the transition dynamics one can define a Bernoulli variable Y which is 1 if (s, a) results in s' and 0 otherwise. Then, defining the random variable $Z = |P_{s,a}^*(s') - Y|/m$, we have the analogous result

$$m = 2\alpha^{-2}(\text{Var}(Z) + \alpha/3) \ln(2/\delta).$$

Therefore, defining $\mathcal{V} = \max(\text{Var}(Z), \text{Var}_{\max})$ and summing over actions, we have

$$m_{\text{known}} \geq 2A\alpha^{-2}(\mathcal{V} + \alpha/3) \ln(2\delta) K_1^{-2} (STG/\epsilon)^4, \quad (25)$$

which provides a decrease of a factor $\frac{1}{4}\alpha^2 M^2/(\mathcal{V} + \alpha/3)$. As an example, plugging in $M = \max(r_{\max}, c_{\max}) = 100$, $\mathcal{V} = 1$, $\alpha = 0.10$, a 24-fold improvement can be observed.

Finally, it is worth noting that in many settings, the reward and constraint-cost function are a deterministic, rather than stochastic, function of the state-action pair. In such cases, a single sample suffices to know the average reward and constraint-cost for a given state-action pair, and therefore the number of samples for a state to be known depends only on the result from the transition dynamics, that is,

$$m_{\text{known}} = \frac{1}{2}A \ln\left(\frac{2}{\delta}\right) K_1^{-2} (STG/\epsilon)^4. \quad (26)$$

From worst-case to practical assumptions In deriving the number of actions that need to be taken to reach a near-optimal policy for the entire CMDP, many worst-case assumptions were made. With more practical assumptions, one can get improved results for the sample complexity. First, the worst-case number of exploration attempts is $N = O(G_{\max}^r(T)/\epsilon \ln(S/\delta) S m_{\text{known}})$ can be improved by noting that one does not always need to make all states known before having the ϵ -optimal solution, yielding some $S' < S$. Second, both the number of exploration attempts and m_{known} can be improved by estimating $G_{\max}^r(T)$ and $G_{\max}^c(T)$ via empirical trajectories instead of using a generic upper bound. Third, the cost of c_{\max} at all time steps in the Worst-case Escape CMDP could be improved to $\hat{c}(s, a) + \alpha_u(s, a)$, where $\alpha_u(s, a)$ is the approximation error for state-action pairs that have been visited already (but are not yet known). For example, one could set $\alpha_u(s, a)$ based on the standard error $\alpha_u(s, a) = \sqrt{\frac{\text{Var}_{\max}^c}{n(s, a)}}$ such that more balanced wandering steps could be taken. Note that such changes to the Worst-case Escape MDP would also allow for the application of E^4 to CMDPs with larger diameter. Fourth, using a separate α (and/or α_u) for the reward, constraint-cost, and transition dynamics functions may account for the range and variance of these different random variables. Fifth, in practice, there will be an initial set of known states; this further reduces the number of states that need to be made known to a number $S'' < S' < S$. Sixth, many of the settings from Lemma 10 can be significantly optimised in practice: (a) the diameter requirement can be relaxed in domains such as Example 2; (b) the current path as well as predictions of future constraint-cost may indicate when the T -step safe return trajectory is not needed; (c) the bounds on the budget can be improved by using the specific setting of γ of the target CMDP instead of the worst-case setting of $\gamma = 1$; and (d) the known-state and unknown-state budgets can take further information into account, such as the current number of known and unknown states. Finally, the use of adaptive sampling, which has been used in prior work on E^3 (Domingo, 1999), may also be beneficial for practical E^4 implementations.

Can one use E^4 for constraints on individual paths? A particular trajectory p starting on a state $s \in \mathcal{S}$ may yield $C(p) > d$, even if $C_\pi(s) \leq d$, since C_π is formulated on expectation over paths. While E^4 has not been designed for constraints on individual paths, if the user desires probabilistic guarantees on the constraint-cost of paths with the same level d , one can form a confidence interval based on the standard deviation S_C , which if unknown can be upper-bounded by $S_C \leq \sum_{t=0}^{\infty} \gamma^t \sqrt{\text{Var}_{\max}^c} = \frac{1}{1-\gamma} \sqrt{\text{Var}_{\max}^c}$. Then the budget can be reformulated as $d_p \leftarrow d - 3S_C$ such that the budget of interest d is rarely exceeded. For example, if the asymptotic constraint-cost of paths is normally distributed with mean d_p and standard deviation S_C , then there is at most a 0.1% probability that any such path p has

constraint-cost $C(p) > d$. More generally, if one has no information on the distribution, one may use concentration equalities to provide upper bounds on failure probabilities. For example, with the Chebyshev-Cantelli inequality (Cantelli, 1928), the same example yields

$$\mathbb{P}(C(p) - d > 0) = \mathbb{P}(C(p) - d_p > 3S_C) < \frac{S_C^2}{S_C^2 + 9S_C^2} = \frac{1}{10}.$$

When it is not the case that $d - 3S_C$ is positive (and greater than $2c_{max}$), an alternative safe exploration algorithm would be required, one which is explicitly formulated to handle probability of individual failures. In E^4 , the reasoning is that having $C(p) > d$ is dangerous but usually not catastrophic, while in the above setting $C(p) > d$ is always catastrophic. To the authors' best knowledge, handling individual failures in this manner has not been done before as this requirement is too strict, but solving this exciting challenge would often be useful in safety-critical settings.

5 Conclusion

Safety is of critical concern in real world applications, and especially in RL, where an agent interacts with an initially unknown environment. Unlike game environments, failures in real world safety-critical applications will finish the lifetime of the RL agent with serious cost to the owner of the agent, and potentially, society at large. Therefore, instead of the model-free RL algorithms which are formulated with this trial-and-error setting in mind, a model-based RL approach may be more suitable for safety-critical applications. Model-based RL algorithms such as E^3 learn near-optimal policies with polynomial sample and time complexity, making them an attractive option for learning in the real world as the model can be used to perform offline computations without requiring too much real world trial-and-error. This paper integrates E^3 into a constrained Markov decision process framework to show that – as long as the constraints are satisfiable due to a limited diameter of the CMDP – there exists an algorithm (E^4) which with high probability finds a near-optimal *constrained* policy within polynomial time. The E^4 algorithm combines the use of an explicit exploration and an explicit exploitation policy with an additional escape policy that provides a safe return for states not reliably known by the model. The algorithm is attractive for safety-critical settings, not only due to the offline computations, but also because of the non-episodic setting in which there is only a single lifetime. The simulation model allows anticipating constraint-satisfaction failures and where there is uncertainty about the true environment, distributional robustness is used to ensure the worst-case scenario does not violate the constraints. Beyond theoretical results supporting the framework, a discussion highlights offline optimisation algorithms and shows how to formulate uncertainty sets for unknown states based on prior knowledge, empirical inference, or a combination thereof. Further practical considerations discussed include relaxing the worst-case assumptions that underlie the theory, aspects of the domain and the implementation that affect sample efficiency, and the practical interpretation of the various components of the algorithm.

A few exciting research directions emerge from this paper, including the development of scalable methods for constrained robust offline optimisation, defining uncertainty sets for unknown states, as well as the explicit use of exploration, exploitation, and escape policies within other model-based or even model-free RL algorithms. Given that the explicit

separation of exploration and exploitation that characterised E^3 has been taken up by recent works in different learning settings, including hard-exploration (Ecoffet et al., 2021) and meta-learning (Liu et al., 2020), E^4 may provide unique perspectives on how to solve such problems *safely*.

Appendix A: constrained simulation lemma

Below is the proof of the Constrained Simulation Lemma and the Known State Lemma.

Constrained Simulation Lemma Let $T \geq \frac{1}{1-\gamma} \ln \left(\frac{\max(r_{\max}, c_{\max})}{\frac{\epsilon}{2}(1-\gamma)} \right)$,⁵ let CMDP \hat{M} be an α -approximation of CMDP M with $\alpha = O((\epsilon/(STG))^2)$ where $G = \max(G_{\max}^r(T), G_{\max}^c(T))$, and let π be a policy in M . Then for any state s ,

- (a) $V_{\pi}(s) - \epsilon \leq \hat{V}_{\pi}(s) \leq V_{\pi}(s) + \epsilon$;
- (b) $C_{\pi}(s) - \epsilon \leq \hat{C}_{\pi}(s) \leq C_{\pi}(s) + \epsilon$.

Proof Let π be any policy in M . Define β -small transition as a transition for which $P_{s,a}^*(s') \leq \beta$. After T time steps, there is a probability of at most $T * S * \beta$ that a β -small transition is crossed (since at any one time step at most $S * \beta$). The proof distinguishes between case A) the path traversed includes at least one β -small transition; and B) the path does not include any β -small transition.

Case 1: at least one β -small transition

If we do have β -small transitions then:

- 1) T -step walks of π from s that cross at least one β -small transition contribute at most $TS\beta G_{\max}^r(T)$ (on expectation) to $V_{\pi}(s, T)$ and at most $TS\beta G_{\max}^c(T)$ (on expectation) to $C_{\pi}(s, T)$.
- 2) Since \hat{M} is an α -approximation of M , this means that β -small transitions satisfy $P_{s,a}^*(s') \leq \beta + \alpha$. Therefore, T -step walks of π from s that cross at least one β -small transition contribute at most $TS(\alpha + \beta)G_{\max}^r(T)$ (on expectation) to $\hat{V}_{\pi}(s, T)$ and at most $TS(\alpha + \beta)G_{\max}^c(T)$ (on expectation) to $\hat{C}_{\pi}(s, T)$. These also give the upper bounds to the contribution to the discrepancy of \hat{M} to M . So, the remainder of the proof consists of choosing an appropriate β and then solving for α to bound the discrepancy to the desired ϵ in the following equation:

$$(\alpha + 2\beta)TSG_{\max}^r(T) \leq \epsilon/4 \quad (27)$$

$$(\alpha + 2\beta)TSG_{\max}^c(T) \leq \epsilon/4. \quad (28)$$

In the general case, we will then solve the approximation for paths with no β -small transitions and then adding the above-mentioned $\epsilon/4$ to account for the maximal contribution of such β -small transitions.

⁵ The $\epsilon/2$ - and ϵ -horizon times are linearly dependent and therefore their difference vanishes in the big O notation for α (Kearns & Singh, 2002).

Case 2: no β -small transitions

For T -step trajectories with no β -small transitions, we can convert the additive approximation

$$P_{s,a}^*(s') - \alpha \leq \hat{P}(s, a, s') \leq P_{s,a}^*(s') + \alpha$$

into a multiplicative approximation

$$(1 - \Delta)P_{s,a}^*(s') \leq \hat{P}_{s,a}(s') \leq (1 + \Delta)P_{s,a}^*(s'),$$

where $\Delta = \alpha/\beta$ because $P > \beta$. To illustrate, if $P > \beta$ then if $X \leq P + \alpha$, this means that $X \leq \alpha + \beta < (1 + \Delta) * P$. Therefore, for any T -step path p without β -small transitions, we have

$$(1 - \Delta)^T \mathbb{P}_{P,\pi}(p) \leq \mathbb{P}_{\hat{P},\pi}(p) \leq (1 + \Delta)^T \mathbb{P}_{P,\pi}(p).$$

A similar argument follows for the value function. For all state-actions $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have $r(s, a) - \alpha \leq \hat{r}(s, a) \leq r(s, a) + \alpha$ and therefore

$$V(p) - T\alpha \leq \hat{V}_\pi(p) \leq V(p) + T\alpha$$

holds for any path p . Then, since we took any arbitrary path without β -small transitions, taking the expectation over such paths will yield the same inequality, while we can further additionally account for $\epsilon/4$ as the added contribution of T -step walks *with* β -small transitions. Then converting it to a multiplicative approximation over the expected T -step value yields

$$(1 - \Delta)^T (V_\pi(s, T) - T\alpha) - \epsilon/4 \leq \hat{V}_\pi(s, T) \leq (1 + \Delta)^T (V_\pi(s, T) + T\alpha) + \epsilon/4.$$

Then provide an upper bound for Δ based on a Taylor expansion:

$$\ln \left((1 + \Delta)^T \right) = T \ln(1 + \Delta) = T(\Delta - \Delta^2/2 + \Delta^3/3 - \dots) \geq T\Delta/2$$

Requiring that

$$(1 + \Delta)^T (V_\pi(s, T) + T\alpha) + \epsilon/4 \leq V_\pi(s, T) + \epsilon/2 \quad (29)$$

we need both conditions to hold:

- 1) $(1 + \Delta)^T V_\pi(s, T) \leq V_\pi(s, T) + \epsilon/8.$
- 2) $(1 + \Delta)^T T\alpha \leq \epsilon/8.$

Solving for Δ in condition 1) yields

$$\begin{aligned} (1 + \Delta)^T (V_\pi(s, T)) &\leq V_\pi(s, T) + \epsilon/8 \\ (1 + \Delta)^T &\leq (1 + \epsilon/(8G_{\max}^r(T))) \\ T\Delta/2 &\leq T \ln(1 + \Delta) \leq \ln \left(1 + \epsilon/(8G_{\max}^r(T)) \right) \approx \epsilon/(8G_{\max}^r(T)) \\ \Delta &\leq \epsilon/(4TG_{\max}^r(T)). \end{aligned}$$

Line 2 and 4 of the above imply together that $(1 + \Delta)^T$ is a constant such that $\alpha T = O(\epsilon)$.

The requirement for the constraint-cost is completely analogous. For all $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have $c(s) - \alpha \leq \hat{c}(s) \leq c(s) + \alpha$ and therefore

$$C(p) - T\alpha \leq \hat{C}_\pi(p) \leq C(p) + T\alpha \quad (30)$$

holds. Requiring that

$$(1 + \Delta)^T (C_\pi(s, T) + T\alpha) + \epsilon/4 \leq C_\pi(s, T) + \epsilon/2 \quad (31)$$

we need both conditions to hold:

- 1) $(1 + \Delta)^T C_\pi(s, T) \leq C_\pi(s, T) + \epsilon/8$.
- 2) $(1 + \Delta)^T T\alpha \leq \epsilon/8$.

Analogous to the value function, this yields

$$\Delta \leq \epsilon / (4TG_{\max}^c(T)), \quad (32)$$

and $\alpha T = O(\epsilon)$.

Solving for α

From the requirements for the value function (see Eq. 27), we have $(\alpha + 2\beta)TSG_{\max}^r(T) \leq \epsilon/4$. Choosing $\beta = \sqrt{\alpha}$, we require

$$(\alpha + 2\beta)TSG_{\max}^r(T) \leq 3\sqrt{\alpha}TSG_{\max}^r(T) \leq \epsilon/4,$$

and

$$\Delta = \alpha/\beta = \sqrt{\alpha} \leq \epsilon / (4TG_{\max}^r(T)),$$

and $\alpha T = O(\epsilon)$. The choice of $\alpha = O((\epsilon/(STG_{\max}^r(T)))^2)$ provides these results.

Similarly, we require for the constraint-cost function (see Eq. 27) that $(\alpha + 2\beta)TSG_{\max}^c(T) \leq \epsilon/4$. Analogously, requiring

$$(\alpha + 2\beta)TSG_{\max}^c(T) \leq 3\sqrt{\alpha}TSG_{\max}^c(T) \leq \epsilon/4,$$

and

$$\Delta = \alpha/\beta = \sqrt{\alpha} \leq \epsilon / (4TG_{\max}^c(T)),$$

and $\alpha T = O(\epsilon)$. The choice of $\alpha = O((\epsilon/(STG_{\max}^c(T)))^2)$ provides these results.

Therefore, to meet the requirements for both the value function and the constraint-cost function, selecting $\alpha = O((\epsilon/(STG))^2)$, where $G = \max(G_{\max}^r(T), G_{\max}^c(T))$, yields the above requirements.

Taking the expectation over the T -step paths in Eq. 29 and Eq. 30, we have

$$V_\pi(s, T) - \epsilon/2 \leq \hat{V}_\pi(s, T) \leq V_\pi(s, T) + \epsilon/2$$

and applying the Constrained T-step estimation Lemma (Lemma 1) to \hat{V} and \hat{C} , we have:

$$\hat{V}_\pi(s, T) \leq \hat{V}_\pi(s) \leq \hat{V}_\pi(s, T) + \epsilon/2.$$

Similarly, we have

$$V_{\pi}(s, T) \leq V_{\pi}(s) \leq V_{\pi}(s, T) + \epsilon/2.$$

Therefore, we have

$$\hat{V}_{\pi}(s) \leq \hat{V}_{\pi}(s, T) + \epsilon/2 \leq V_{\pi}(s, T) + \epsilon \leq V_{\pi}(s) + \epsilon$$

and

$$V_{\pi}(s) - \epsilon \leq V_{\pi}(s, T) - \epsilon/2 \leq \hat{V}_{\pi}(s, T) \leq \hat{V}_{\pi}(s).$$

Analogous results follow for the constraint-cost, concluding the proof. \square

Known State Lemma Given state s has been visited

$$m = O((STG/\epsilon)^4 \text{Var}_{\max} \ln(1/\delta))$$

times, where $G = \max(G_{\max}^r(T), G_{\max}^c(T))$ and $\text{Var}_{\max} = \max(\text{Var}_{\max}^r, \text{Var}_{\max}^c)$, and from s each action has been executed at least $\lfloor m/A \rfloor$ times, then with probability of at least $1 - \delta$ we have

- a) $|P_{s,a}^*(s') - \hat{P}_{s,a}(s')| = O((\epsilon/(STG))^2)$ for any $(s, a, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$;
- b) $|r(s, a) - \hat{r}(s, a)| = O((\epsilon/(STG))^2)$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$; and
- c) $|c(s, a) - \hat{c}(s, a)| = O((\epsilon/(STG))^2)$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$.

Proof

- a) **Transition dynamics:** Note that $P_{s,a}^*(s') \in [0, 1]$. Therefore, applying Hoeffding's inequality (Hoeffding, 1963), and setting δ equal to the resulting upper bound, we have

$$\begin{aligned} \mathbb{P}\left(|P_{s,a}^*(s') - \hat{P}_{s,a}(s')| \geq t\right) &\leq 2e^{-2mt^2} = \delta \\ -mt^2 &= \frac{1}{2} \ln\left(\frac{\delta}{2}\right) \\ m &= \frac{\frac{1}{2} \ln\left(\frac{2}{\delta}\right)}{t^2}. \end{aligned}$$

Setting $t = (\epsilon/(STG))^2$, we see that

$$m = \frac{\frac{1}{2} \ln\left(\frac{2}{\delta}\right)}{(\epsilon/(STG))^4} = \frac{1}{2} \ln\left(\frac{2}{\delta}\right) (STG/\epsilon)^4.$$

Thus $O((STG/\epsilon)^4 \ln \frac{1}{\delta})$ samples are needed to obtain $O((\epsilon/(STG))^2)$ -accurate estimates of each transition probability.

- b) **Rewards:** Note that $G^2 \geq r_{\max}^2$ and that $(STG/\epsilon)^4$ grows as $\epsilon \rightarrow 0$ so restricting $\epsilon^2 \in [0, 2\text{Var}_{\max}^r]$ includes the worst-case asymptotic behaviour of $(STG/\epsilon)^4 \text{Var}_{\max}^r \ln(1/\delta)$. For any $\epsilon \in [0, 2\text{Var}_{\max}^r]$ and $t = (\epsilon/(STG))^2 \in [0, 2\text{Var}_{\max}^r/r_{\max}^2]$ we have (Phillips, 2012):

$$\begin{aligned}\mathbb{P}(|r(s, a) - \hat{r}(s, a)| \geq t) &\leq 2e^{-\frac{mt^2}{4\text{Var}_{\max}^r}} = \delta \\ \frac{mt^2}{4\text{Var}_{\max}^r} &= \ln(2/\delta) \\ m &= 4 \ln(2/\delta) \text{Var}_{\max}^r t^{-2}.\end{aligned}$$

Therefore, $O\left((STG/\epsilon)^4 \text{Var}_{\max}^r \ln \frac{1}{\delta}\right)$ samples are needed to obtain $O((\epsilon/(STG))^2)$ - accurate estimates of the average reward for a given state.

c) **Constraint-costs:** Analogously, with $t = (\epsilon/(STG))^2 \in [0, 2\text{Var}_{\max}^c/c_{\max}^2]$, we have:

$$\begin{aligned}\mathbb{P}(|c(s, a) - \hat{c}(s, a)| \geq t) &\leq 2e^{-\frac{mt^2}{4\text{Var}_{\max}^c}} = \delta \\ \frac{mt^2}{4\text{Var}_{\max}^c} &= \ln(2/\delta) \\ m &= 4 \ln(2/\delta) \text{Var}_{\max}^c t^{-2}.\end{aligned}$$

Therefore, $O\left((STG/\epsilon)^4 \text{Var}_{\max}^c \ln \frac{1}{\delta}\right)$ samples are needed to obtain $O((\epsilon/(STG))^2)$ - accurate estimates of the average constraint-cost for a given state.

Combining worst-cases: Taking the worst case of the three conditions on m , and setting $\text{Var}_{\max} = \max(\text{Var}_{\max}^r, \text{Var}_{\max}^c)$, we have

$$m = O\left((STG/\epsilon)^4 \text{Var}_{\max} \ln \frac{1}{\delta}\right).$$

□

Appendix B: explore-or-exploit lemma

Constrained Explore-or-Exploit Lemma Let M be any CMDP, let \mathcal{S} be any subset of states $\mathcal{S} \subset \mathcal{S}$, and let $M_{\mathcal{S}}$ be the induced CMDP over \mathcal{S} with a given budget d . For any $s \in \mathcal{S}$, for any T , and any $\epsilon \geq 0$, we have **either a)** there exists a policy $\pi \in M_{\mathcal{S}}$ for which $V_{\pi}(s, T|M_{\mathcal{S}}) \geq V_{\pi^*}(s, T) - \epsilon$, where $\pi^* = \arg \max_{\pi \in \Pi_c(s, d, T)} V_{\pi}(s, T)$ is the optimal constrained T -step policy, and which satisfies $C_{\pi}(s, T|M_{\mathcal{S}}) \leq d$; **or b)** there exists a policy π in $M_{\mathcal{S}}$ which reaches the terminal state s_0 in $\mathcal{S} \setminus \mathcal{S}$ in at most T steps with probability $p > \epsilon/G_{\max}^r(T)$, and $C_{\pi}(s, T|M_{\mathcal{S}}) \leq d$.

Proof Case a) $V_{\pi}(s, T|M_{\mathcal{S}}) \geq V_{\pi^*}(s, T) - \epsilon$:

In this case, condition a) is satisfied.

Case b) $V_{\pi}(s, T|M_{\mathcal{S}}) < V_{\pi^*}(s, T) - \epsilon$:

For some set of policies Π^* in M , we have $V_{\pi}(s, T|M_{\mathcal{S}}) = V_{\pi^*}(s, T)$ for all $\pi \in \Pi^*$.⁶ Since $\pi^* \in \Pi_c(s, d, T)$, this implies via the Induced Underestimation Lemma that for a subset $\Pi_{\mathcal{S}}^* \subseteq \Pi^*$, we have additionally $C_{\pi}(s, T|M_{\mathcal{S}}) \leq C_{\pi^*}(s, T) \leq d$ for all $\pi \in \Pi_{\mathcal{S}}^*$. Now select

⁶ π may not be the optimal constrained policy; since the T -step value only approximates the asymptotic value $V(s) = \lim_{T' \rightarrow \infty} V(s, T')$, there may be many such policies.

$\pi \in \Pi_{\mathcal{S}}^*$. Decompose T-paths starting from s into T-paths that only contain states in \mathcal{S} (p_k -type paths) and T-paths that also contain states $s \notin \mathcal{S}$ (p_u -type paths):

$$\begin{aligned} V_{\pi}(s, T) &= \sum_p \mathbb{P}_{\pi, P}[p] V_{\pi}(p) \\ &= \sum_{p_k} \mathbb{P}_{\pi, P}[p_k] V_{\pi}(p_k) + \sum_{p_u} \mathbb{P}_{\pi, P}[p_u] V_{\pi}(p_u). \end{aligned}$$

For any p_k -type path, we have $\mathbb{P}_{\pi, P}[p_k] = \mathbb{P}_{\pi, P^{\mathcal{M}_{\mathcal{S}}}}[p_k]$ and $V(p_k) = V(p_k | \mathcal{M}_{\mathcal{S}}) \leq V(s, T | \mathcal{M}_{\mathcal{S}})$, where equalities follow from the paths all containing the known set only (since for the induced model the transition dynamics are the same and rewards outside \mathcal{S} are not being considered) and the inequality follows from positive rewards and $V(s, T | \mathcal{M}_{\mathcal{S}})$ considering additionally p_u -type paths if any. Therefore,

$$\begin{aligned} \sum_{p_k} \mathbb{P}_{\pi, P}[p_k] V_{\pi}(p_k) &= \sum_{p_k} \mathbb{P}_{\pi, P^{\mathcal{M}_{\mathcal{S}}}}[p_k] V_{\pi}(p_k | \mathcal{M}_{\mathcal{S}}) \\ &\leq V_{\pi}(s, T | \mathcal{M}_{\mathcal{S}}) \\ &< V_{\pi^*}(s, T) - \epsilon \end{aligned}$$

Therefore, since $V_{\pi^*}(s, T) = V_{\pi}(s, T) = \sum_{p_k} \mathbb{P}_{\pi, P}[p_k] V_{\pi}(p_k) + \sum_{p_u} \mathbb{P}_{\pi, P}[p_u] V_{\pi}(p_u)$, this implies that:

$$\sum_{p_u} \mathbb{P}_{\pi, P}[p_u] V_{\pi}(p_u) > \epsilon.$$

Since $V_{\pi}(p) \leq G_{\max}^r(T)$ for any T -step path, it follows that

$$\sum_{p_u} \mathbb{P}_{\pi, P}[p_u] > \epsilon / G_{\max}^r(T)$$

Therefore, condition b) is satisfied. \square

Appendix C: Chernoff bound for number of attempted explorations

For any state $s \in \mathcal{S}$, allocate at most $\frac{\delta}{S}$ probability to the state failing to become known (i.e. the state receives fewer than m_{known} visits). Let $p = \epsilon / G_{\max}^r(T)$, let N be the number of exploration attempts, and let $m_{\text{known}} = (1 - \Delta)Np$ and $\Delta = \frac{Np - m_{\text{known}}}{Np}$. Then, by Chernoff bounds, the probability of the number of successful exploration attempts being smaller than m_{known} is bounded by

$$\mathbb{P}[X \leq (1 - \Delta)Np] \leq \left(\frac{e^{-\Delta}}{(1 - \Delta)^{1-\Delta}} \right)^{Np} = \frac{\delta}{S},$$

which leads to a total number of exploration attempts

$$Np(\Delta + (1 - \Delta) \ln(1 - \Delta)) = \ln\left(\frac{S}{\delta}\right)$$

$$Np\Delta = \ln\left(\frac{S}{\delta}\right) \quad N \text{ is maximal for } \Delta \rightarrow 1 \text{ and } \lim_{x \rightarrow 0+} x \ln(x) = 0$$

$$Np - m_{\text{known}} = \ln\left(\frac{\delta}{S}\right)$$

$$N = p^{-1} \left(\ln\left(\frac{S}{\delta}\right) + m_{\text{known}} \right)$$

$$N = O\left(p^{-1} \ln\left(\frac{S}{\delta}\right) m_{\text{known}}\right).$$

Filling in $p = \epsilon / G_{\max}^r(T)$ and repeating over all $s \in \mathcal{S}$, the desired result is obtained

$$N = O\left(\frac{G_{\max}^r(T)}{\epsilon} \ln\left(\frac{S}{\delta}\right) S m_{\text{known}}\right).$$

Author Contributions David M. Bossens: Conceptualization, Formal analysis, Methodology, Writing - Original Draft, Project administration. Nicholas Bishop: Conceptualization, Writing - Review & Editing.

Funding David M. Bossens was supported by the UKRI Trustworthy Autonomous Systems Hub, EP/V00784X/1. Nicholas Bishop was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) Doctoral Training Partnership grant.

Availability of data and material Not applicable.

Code Availability Not applicable.

Declarations

Conflict of interest None.

Ethical approval Not applicable.

Consent to participate Not applicable.

Consent for publication Not applicable.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

Achiam, J., Held, D., Tamar, A., & Abbeel, P. (2017). Constrained policy optimization. *34th International Conference on Machine Learning, ICML 2017*, 1:30–47. [arXiv:1705.10528](https://arxiv.org/abs/1705.10528)

- Allen-Zhu, Z., Li, Y., & Song, Z. (2019). A convergence theory for deep learning via over-parameterization. *36th International Conference on Machine Learning, ICML 2019*, 2019:362–372. [arXiv:1811.03962](https://arxiv.org/abs/1811.03962)
- Altman, E. (1998). Constrained Markov decision processes with total cost criteria: Lagrangian approach and dual linear program. *Mathematical Methods of Operations Research*, 48(3), 387–417. <https://doi.org/10.1007/s001860050035>
- Altman, E. (1999). *Constrained Markov decision processes*. Cambridge: CRC Press. <https://doi.org/10.1109/CDC.2012.6426596>
- Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mané, D. (2016). Concrete problems in AI safety. pp. 1–29. [arXiv:1606.06565](https://arxiv.org/abs/1606.06565)
- Berkenkamp, F., Turchetta, M., Schoellig, A. P., & Krause, A. (2017). Safe model-based reinforcement learning with stability guarantees. *Advances in Neural Information Processing Systems*, 2017(Nips):909–919. [arXiv:1705.08551](https://arxiv.org/abs/1705.08551)
- Brafman, R. I., & Tennenholtz, M. (2002). R-MAX—a general polynomial time algorithm for near-optimal reinforcement learning. *Journal of Machine Learning Research*, 3, 213–231.
- Cantelli, F. P. (1928). Sui confini della probabilità. In *Atti del Congresso Internazionale del Matematici* 6, pp. 47–60, Bologna. <https://doi.org/10.1162/153244303765208377>
- Chow, Y., Nachum, O., Duenez-Guzman, E., & Ghavamzadeh, M. (2018). A lyapunov-based approach to safe reinforcement learning. *Advances in Neural Information Processing Systems*, 2018(NeurIPS):8092–8101. [arXiv:1805.07708](https://arxiv.org/abs/1805.07708)
- Domingo, C. (1999). Faster near-optimal reinforcement learning: Adding adaptiveness to the e3 algorithm. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 1720, 241–251. https://doi.org/10.1007/3-540-46769-6_20
- Ecoffet, A., Huizinga, J., Lehman, J., Stanley, K. O., & Clune, J. (2021). First return, then explore. *Nature*, 590(7847), 580–586. <https://doi.org/10.1038/s41586-020-03157-9> [arXiv:2004.12919](https://arxiv.org/abs/2004.12919)
- Everitt, T., Lea, G., & Hutter, M. (2018). AGI safety literature review. *IJCAI International Joint Conference on Artificial Intelligence*, 2018:5441–5449. <https://doi.org/10.24963/ijcai.2018/768>, [arXiv:1805.01109](https://arxiv.org/abs/1805.01109)
- Eysenbach, B., Gu, S., Ibarz, J., & Levine, S. (2018). Leave no trace: Learning to reset for safe and autonomous reinforcement learning. *6th International Conference on Learning Representations, ICLR 2018 - Conference Track Proceedings*. [arXiv:1711.06782](https://arxiv.org/abs/1711.06782)
- Garcelon, E., Ghavamzadeh, M., Lazaric, A., & Pirodda, M. (2020). Conservative Exploration in Reinforcement Learning. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS 2020)*.
- Goyal, V., & Grand-Clement, J. (2018). Robust Markov Decision Process: Beyond Rectangularity. pp. 1–27. [arXiv:1811.00215](https://arxiv.org/abs/1811.00215)
- Henaff, M. (2019). Explicit explore-exploit algorithms in continuous state spaces. *Advances in Neural Information Processing Systems*, 32(NeurIPS). [arXiv:1911.00617](https://arxiv.org/abs/1911.00617)
- Hoeffding, W. (1963). Probability Inequalities for Sums of Bounded Random Variables. *Journal of the American Statistical Association*, 58(301), 13–30.
- Iyengar, G. N. (2005). Robust dynamic programming. *Mathematics of Operations Research*, 30(2), 257–280. <https://doi.org/10.1287/moor.1040.0129>
- Jaksch, T., Ortner, R., & Auer, P. (2010). Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11, 1563–1600.
- Karamarkar, N. (1984). A new polynomial-time algorithm for linear programming. *Combinatorica*, 4(April), 373–395.
- Kearns, M., & Singh, S. (2002). Near-optimal reinforcement learning in polynomial time. *Machine Learning*, 49(2–3), 209–232. <https://doi.org/10.1023/A:1017984413808>
- Khachiyan, L. G. (1979). A polynomial algorithm in linear programming. *Dokl. Akad. Nauk SSSR*, 244(5):1093–1096. <https://doi.org/10.18287/0134-2452-2015-39-4-459-461>
- Lecarpentier, E., & Rachelson, E. (2019). Non-stationary markov decision processes a worst-case approach using model-based reinforcement learning. *Advances in Neural Information Processing Systems*, 32(NeurIPS). [arXiv:1904.10090v4](https://arxiv.org/abs/1904.10090v4)
- Liu, E. Z., Raghunathan, A., Liang, P., & Finn, C. (2020). Decoupling Exploration and Exploitation for Meta-Reinforcement Learning without Sacrifices. Retrieved from [arXiv:2008.02790](https://arxiv.org/abs/2008.02790)
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533. <https://doi.org/10.1038/nature14236> [arXiv:1604.03986](https://arxiv.org/abs/1604.03986)
- Naik, A., Abbas, Z., White, A., & Sutton, R. S. (2021). 2021. In *NERL Workshop at ICLR: Towards Reinforcement Learning in the Continuing Setting*.
- Nilim, A., & Ghaoui, L. E. (2005). Robust control of Markov decision processes with uncertain transition matrices. *Operations Research*, 53(5), 780–798. <https://doi.org/10.1287/opre.1050.0216>

- Phillips, J. M. (2012). *Chernoff-Hoeffding Inequality and Applications.*, 1, 1–6. [arXiv:1209.6396](#).
- Platanios, E. A., Saparov, A., & Mitchell, T. (2020). Jelly Bean World: A Testbed for Never-Ending Learning. In *ICLR, 2020*, 1–17.
- Potra, F. A., & Wright, S. J. (2000). Interior point methods. *Journal of Computational and Applied Mathematics*, 124, 281–302. https://doi.org/10.1007/978-3-319-58356-3_17
- Rummery, G. A., & Niranjan, M. (1994). Online Q-learning Using Connectionist Systems. Technical report, Cambridge University Engineering Department.
- Russel, R. H., Benosman, M., & Van Baar, J. (2020). Robust Constrained-MDPs: Soft-Constrained Robust Policy Optimization under Model Uncertainty. Retrieved from [arXiv:2010.04870](#)
- Russel, R. H., Benosman, M., Van Baar, J., & Corcodel, R. (2021). Lyapunov Robust Constrained-MDPs: Soft-Constrained Robustly Stable Policy Optimization under Model Uncertainty. Retrieved from [arXiv:2108.02701](#)
- Russel, R. H., & Petrik, M. (2019). Beyond confidence regions: Tight Bayesian ambiguity sets for robust MDPs. *Advances in Neural Information Processing Systems*, 32(NeurIPS). [arXiv:1902.07605](#)
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal Policy Optimization Algorithms. *arXiv preprint*, pages 1–12. Retrieved from [arXiv:1707.06347https://doi.org/10.1016/j.jdeveco.2016.04.001](#)
- Strehl, A. L., Li, L., & Littman, M. L. (2006). Incremental model-based learners with formal learning-time guarantees. *Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence, UAI 2006*, pp. 485–493. [arxiv.org/abs/1206.6870](#)
- Taleghani, M. A., & Dietterich, T. G. (2018). Efficient exploration for constrained MDPs. *AAAI Spring Symposium - Technical Report*, 2018:313–319.
- Thomas, P. S., Theodorou, G., & Ghavamzadeh, M. (2015). High confidence policy improvement. *32nd International Conference on Machine Learning, ICML 2015*, 3(2002):2370–2378.
- Turchetta, M., Kolobov, A., Shah, S., Krause, A., & Agarwal, A. (2020). Safe Reinforcement Learning via Curriculum Induction. (NeurIPS). Retrieved from [arXiv:2006.12136](#)
- Watkins, C. J. C. H., & Dayan, P. (1992). Q-learning. *Machine Learning*, 8(3–4), 279–292. <https://doi.org/10.1007/BF00992698>
- Wiesemann, W., Kuhn, D., & Rustem, B. (2013). Robust Markov decision processes. *Mathematics of Operations Research*, 38(1), 153–183. <https://doi.org/10.1287/moor.1120.0566>.
- Xu, Y., & Yin, W. (2013). A block coordinate descent method for regularized multiconvex optimization with applications to nonnegative tensor factorization and completion. *SIAM Journal on Imaging Sciences*, 6(3), 1758–1789. <https://doi.org/10.1137/120887795>
- Yang, Q., Simao, T. D., Tindemans, S. H., & Spaan, M. T. J. (2021). WCSAC: Worst-case soft actor critic for safety-constrained reinforcement learning. *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence*.
- Zheng, L., & Ratliff, L. J. (2020). Constrained Upper Confidence Reinforcement Learning with Known Dynamics. In *Proceedings of Machine Learning Research*, 120, 1–10.
- Zhou, Y., & Liang, Y. (2017). Characterization of gradient dominance and regularity conditions for neural networks. (1965). Retrieved from [arXiv:1710.06910](#)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.