Effects of Safety State Augmentation on Safe Exploration

Anonymous Author(s)
Affiliation
Address
email

Abstract

Safe exploration is a challenging and important problem in model-free reinforcement learning (RL). Often the safety cost is sparse and unknown, which unavoidably leads to constraint violations — a phenomenon ideally to be avoided in safety-critical applications. We tackle this problem by augmenting the state-space with a safety state, which is nonnegative if and only if the constraint is satisfied. The value of this state also serves as a distance toward constraint violation, while its initial value indicates the available safety budget. This idea allows us to derive policies for scheduling the safety budget during training. We call our approach Simmer (Safe policy IMproveMEnt for RL) to reflect the careful nature of these schedules. We apply this idea to two safe RL problems: RL with constraints imposed on an average cost, and RL with constraints imposed on a cost with probability one. Our experiments suggest that simmering a safe algorithm can improve safety during training for both settings. We further show that Simmer can stabilize training and improve the performance of safe RL with average constraints.

1 Introduction

Reinforcement learning (RL) is a framework for sequential decision-making that makes minimal prior assumptions about the environment where the agent has to act or make the decisions [42]. The policy for taking the actions is learned through interactions with the environment over time. RL has seen recent successes in playing video games with a computer [29], board games with a human [38] and is on a path toward real-life applications such as video compression [43], and plasma control [17]. There are still, however, some unsolved challenges for a successful deployment of RL such as efficient learning of constrained or safe Markov Decision Processes (MDPs) [3]. The constraints are typically modeled by a discounted sum of nonnegative costs that have to be smaller than some pre-defined value called *the safety budget*.

Exploration is a crucial component of RL and is still an active area of research [22, 37, 31]. In the context of safe RL, while exploring we do not want to incur the safety cost and constraint violations making exploration a harder task. We can distinguish two main research directions for minimizing constraint violation in safe RL: a model-based approach (which includes partially and fully known environments) and a model-free approach. In the model-free case, which we focus on, the agent would almost certainly visit unsafe regions to learn the safe policy, and therefore it is next to impossible to avoid constraint violations completely. The general goal, in this case, is to minimize the number of violations during training. For example, [8] use a conservative safety critic and rejection sampling to choose a "safer" action, [45] design a curriculum learning approach, where the teacher resets the student violating safety constraints, and [18] learns to reset the policy if the safety constraint is violated [18].

In this work, we aim to improve model-free safe reinforcement learning by augmenting the state-space with one state encapsulating the safety information. This safety state is initialized with the safety

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

budget and the value of the safety state can serve as a measure of distance to the unsafe region. This form of the safety state was proposed in [15] and introduced in RL [40], where a safe RL algorithm with probability one constraints was derived. In this paper, we take a closer look at the purposes of the safety state in a broader context. First, we claim that safety state augmentation *is often crucial for the averaged constrained problem* as well and provide examples of such occurrences. In some particular problems, the use of safety state augmentation may potentially be avoided, however, this can be said about any state in the environment. Second, our main suggestion is that scheduling the safety budget at different training epochs can improve the algorithm's performance. Our claim extends equally over averaged and almost surely constrained reinforcement learning problems. In particular, we achieve the state of the art performance on some environments in the safety gym benchmark. Further, we claim that scheduling the initial safety budget can lead to reducing safety constraint violations during training for safe RL with probability one constraints. We design two algorithms automatically tuning the safety budget, one is based on a classical control engineering PI controller, while the other uses Q learning to decide on the safety budget.

Related work. Many of the current safe RL methods are extensions of the most successful RL algorithms: Trust region policy optimization (TRPO) [35], proximal policy optimization (PPO) [36], soft actor-critic (SAC) [20] etc. Safe versions of TRPO, PPO, and SAC with a Lagrangian approach were first presented in [34], which is still considered to be one of the major baselines. A direct extension of TRPO by adding constraints to the trust region update was proposed in [1]. Model-based approaches were also considered cf. [33, 13, 24, 27] and most of them took a Bayesian approach to model one-step transitions. To our best knowledge, the most successful approaches for safe reinforcement learning to date are PID-Lagrangian [41] and LAMBDA [4]. The former views the Lagrangian multiplier update as another control problem and employs a PID controller to it (cf. [5]). Specifically, the authors link the multiplier update to integral control and add proportional and derivative controllers to achieve a superior behavior. On the other hand, [4] is a model-based approach that uses Bayesian world models to enhance safety. A recent work [26] formulated safe RL as inference resulting in a sample efficient off-policy approach. Some approaches took a slightly different approach to formulate safety, e.g., [28] proposed a two-player framework with the cooperating task agent and safety agents, [14] proposed a safety layer that would be applied after the action is computed using a classical policy, [21] define a probabilistic shield for safety. We also note [49, 23] that consider safe deployment of RL policies in real-life settings.

Other formulations of safe reinforcement learning were considered in the literature. For example, [11, 13, 48] proposed to use conditional-value-at risk (CVaR) constraints, while [40, 9] proposed to enforce constraints with probability one. Further, as we discussed above eliminating the number of constraint violations typically requires strong assumptions, e.g., finite state space [44], [47] [39], knowledge of a partial model [25] or initial safe policy [6]. These results are in the spirit of safe RL with control-theoretic notions [12, 7, 30, 10, 2, 16, 19], which make significant prior assumptions to guarantee safety. Finally, we note that the closest algorithm in the literature to our method is the curriculum learning approach to Safe RL from [45]. Due to space limitations we discuss these differences in Appendix.

2 Simmer: Safe policy improvement for reinforcement learning

2.1 MDP with safety state augmentation

We consider a constrained reinforcement learning setting defined for a Markov Decision Process (MDP) $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{P}, r, \gamma_r \rangle$ with transition probability $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \to [0, 1]$ acting on state $\mathcal{S}$ and $\mathcal{A}$ spaces, with the reward $r : \boldsymbol{S} \times \boldsymbol{A} \times \boldsymbol{S} \to \mathbb{R}_+$ ($\mathbb{R}_+$ stands for the nonnegative orthant $[0, +\infty)$) and the discount factor $\gamma_r \in (0, 1]$. The MDP is endowed with the following optimization problem

$$\max_{\pi(\cdot|\boldsymbol{s})} \mathbb{E} \sum_{t=0}^{T-1} \gamma_r^t r(\boldsymbol{s}_t, \boldsymbol{a}_t, \boldsymbol{s}_{t+1}), \text{ subject to: } g\left(\sum_{t=0}^{T-1} \gamma_l^t \boldsymbol{l}(\boldsymbol{s}_t, \boldsymbol{a}_t, \boldsymbol{s}_{t+1})\right) \leq \boldsymbol{d}, \tag{1}$$

with the time horizon $T > 0$, the safety discount factor $\gamma_l \in (0, 1]$, the safety cost $\boldsymbol{l} : \boldsymbol{S} \times \boldsymbol{A} \times \boldsymbol{S} \to \mathbb{R}_+^m$, and the inequality should be understood as the element-wise comparison. The statistic $g(\cdot) : \mathbb{R}_+^m \to \mathbb{R}_+^m$ is a design choice and there exits many different possibilities (e.g., CVaR, chance constraints etc [11]). In this paper, we consider two most relevant options, from our point of view: a) a constraint with probability one, i.e., $g_{\text{po}}(\boldsymbol{z}) = \mathbb{P}(\boldsymbol{z} \geq 0) - 1$, b) a constraint on average $g_{\text{av}}(\boldsymbol{z}) = \mathbb{E}\boldsymbol{z}$.

Similarly to [40], we augment the safety state information into the state-space by introducing the state $\boldsymbol{z}_t = \gamma^{-t}(\boldsymbol{d} - \sum_{k=0}^{t-1} \gamma_l^k \boldsymbol{l}(\boldsymbol{s}_k, \boldsymbol{a}_k, \boldsymbol{s}_{k+1}))$, which has the following update $\boldsymbol{z}_{t+1} = (\boldsymbol{z}_t - \boldsymbol{l}(\boldsymbol{s}_t, \boldsymbol{a}_t, \boldsymbol{s}_{t+1}))/\gamma_l$, with $\boldsymbol{z}_0 = \boldsymbol{d}$. Noting that the update is Markovian this state can be easily augmented into the MDP. The variable z_t has the interpretation of the remaining safety budget, and by definition enforcing the constraint on $\boldsymbol{z}_T$ is equivalent to enforcing the constraint on the accumulated cost. Now we can rewrite the safe RL problem as follows:

$$\max_{\pi(\cdot|\boldsymbol{s},\boldsymbol{z})} \mathbb{E} \sum_{t=0}^{T-1} \gamma_r^t r(\boldsymbol{s}_t, \boldsymbol{a}_t, \boldsymbol{s}_{t+1}), \text{ subject to: } g(\boldsymbol{z}_T) \geq 0. \tag{2}$$

While [40] considered only the case $g_{\text{po}}(\boldsymbol{z})$, we argue that the case of $g_{\text{av}}(\boldsymbol{z})$ deserves additional attention in the context of safety state augmentation. For completeness, we review the main points of the approach [40] in Appendix. Note that the policy now also depends on the safety state $\boldsymbol{z}$ and this feature deserves a more thorough discussion.

2.2 Do we need the safety state?

As a simple demonstration consider the cartoon in Figure 1. A robot needs to reach the goal while crossing the hazards region, which is marked by the red circle, and the safety cost is acquired for every time unit spent in the region. Both green and blue paths are safe, i.e., satisfy the constraint. However, at the crossing of these paths robot needs to know which path it took to this state. Switching from the green path to the blue one will lead to a constraint violation. Standard safe RL algorithms will have trouble with such scenarios while adding the safety state solves the issue. We confirm this observation in our experiments.

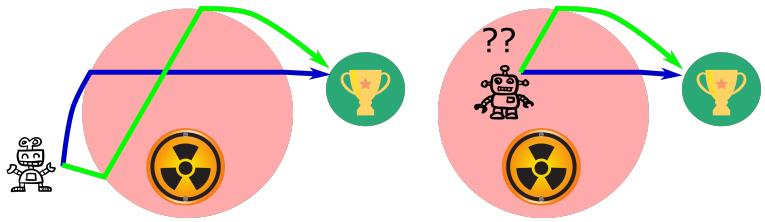

Figure 1: A robot needs to reach the goal while crossing the hazards region (marked by the red circle) and the safety cost is acquired for every time unit spent in the region. Both green and blue paths are safe, i.e., satisfy the constraint. However, at the crossing of these paths robot needs to know which path it took to this state. Switching from the green path to the blue one will lead to constraint violations.

Let us now discuss how this logic can be mathematically formalized. This problem is well studied in the optimal control literature in the context of problems with known dynamics and with terminal or end-point constraints. In the deterministic case, this problem reduces to the standard maximum principle with end-point constraints [46]. Further, the authors [40] showed that the policy dependence on the remaining safety budget is crucial for safe reinforcement learning with probability one constraints. The problem was also studied in the context of stochastic diffusions [32], where the authors derived the representation of the optimal value function. In both cases, it was shown that the optimal policy depends on the whole state, which in our case includes the safety state $\boldsymbol{z}$. While in some cases one may not need to augment the safety state, in many situations it is critical.

2.3 Simmering Safe Reinforcement Learning

In this paper, we propose to use the initial safety budget $\boldsymbol{d}$ as another tuning dial for the algorithms. Specifically, we will exploit the link between the safety budget $\boldsymbol{d}$ with the initial state of the safety state. We argue that adjustment of $\boldsymbol{d}$ during training from some initial value $\boldsymbol{d}^{\text{start}}$ to the target value $\boldsymbol{d}^{\text{target}}$ can lead to improved exploration in terms of safety and performance.

Let us now present the mathematical formulation. At every epoch k we pick a test safety budget $\boldsymbol{d}_k$, collect the data set $\mathcal{D}_k$, compute the returns $\mathbb{E}_{\mathcal{D}_k} \sum_{t=0}^{T-1} \gamma_r^t r(\boldsymbol{s}_t, \boldsymbol{a}_t)$ and the costs $\hat{g}_{\mathcal{D}_k}(\boldsymbol{z}_T)$, where the function $\hat{g}_{\mathcal{D}_k}(\cdot)$ is the empirical mean or the maximum for the averaged and the probability one

constrained problems, respectively. Using this information we aim at solving the following problem at every epoch (but we only apply a pre-determined number of gradient steps):

$$\max_{\pi(\cdot|\boldsymbol{s},\boldsymbol{z})} \mathbb{E}_{\mathcal{D}_k} \sum_{t=0}^{T-1} \gamma_r^t r(\boldsymbol{s}_t, \boldsymbol{a}_t), \text{ subject to: } \hat{g}_{\mathcal{D}_k}(\boldsymbol{z}_T) \geq 0, \boldsymbol{z}_0 = \boldsymbol{d}_k \tag{3}$$

where $s_0 \in \mathcal{S}_0$ and $\boldsymbol{z}_0 = \boldsymbol{d}_k$ are the initial states of the augmented MDP. Note that for off-policy algorithms the data set $\mathcal{D}_k$ can potentially grow with epochs, while for on-policy algorithms the data set $\mathcal{D}_k$ will be emptied on every epoch.

The key steps of our approach are as follows:

- Let there be a finite number of test safety budgets, i.e., we assume that $\boldsymbol{d}_k$ can take the following values $\{\boldsymbol{d}^{0,\text{ref}}, \ldots, \boldsymbol{d}^{\text{K}-1,\text{ref}}\}$ with $\boldsymbol{d}^{\text{start}} = \boldsymbol{d}^{0,\text{ref}}$ and $\boldsymbol{d}^{\text{target}} = \boldsymbol{d}^{\text{K}-1,\text{ref}}$;
- We assume that the value $\boldsymbol{d}^{\text{start}} = \boldsymbol{d}^{0,\text{ref}}$ is such that the task can be solved;
- At for every k we perform only one epoch of optimization solving Problem 3.

Now we can find $\boldsymbol{d}_k$ to achieve our goals and as we will show in the next section the problem of safety during training can be formalized as a two-level decision-making problem. However, we will show that even a naïve approach of scheduling $\boldsymbol{d}_k$ can lead to improvement in performance and safety.

2.4 Application: Simmering for safety during training

Exploration can often lead to constraint violation during training due to the inherent stochasticity of exploration. While there is a significant effort in research for safe exploration it typically requires significant prior assumptions on MDPs. We propose to embrace the philosophy of classical reinforcement learning and proceed with minimal assumptions. We propose to choose the test safety budget $\boldsymbol{d}_k$ using another decision-making problem:

$$\begin{aligned} \max_{\boldsymbol{u}_k \in [-\delta\boldsymbol{d}, \delta\boldsymbol{d}]} \quad & -\sum_k \text{ReLU}\left(-\hat{g}_{\mathcal{D}_k}(\boldsymbol{z}_T^k(\boldsymbol{d}^{\text{target}}))\right), \\ & \boldsymbol{d}_{k+1} = \text{clip}\left(\boldsymbol{d}_k + \boldsymbol{u}_k, \boldsymbol{d}^{\text{start}}, \boldsymbol{d}^{\text{target}}\right), \boldsymbol{d}_0 = \boldsymbol{d}^{\text{start}}, \end{aligned} \tag{4}$$

where $\boldsymbol{z}_T^k(\boldsymbol{d}^{\text{target}})$ is the violation of the target constraint obtained as a result of solving one epoch of Problem 3, the action set is $[-\delta\boldsymbol{d}, \delta\boldsymbol{d}]$ and $\delta\boldsymbol{d}$ is pre-determined, $\text{clip}(x, y, z)$ clips the value x to a lower bound y and an upper bound z, i.e., the function returns y if $x \leq y$, z if $z \leq x$ and x otherwise.

The intuition behind our formulation is quite straightforward. Assume we start with a very strict safety budget $\boldsymbol{d}^{\text{start}}$ training with which does not lead to significant violations with respect to the target safety budget $\boldsymbol{d}^{\text{target}}$. Then by gradually increasing the safety budget $\boldsymbol{d}$ from $\boldsymbol{d}^{\text{start}}$ to $\boldsymbol{d}^{\text{target}}$, we can reduce the number of constraint violations during training. We see this formulation as the first step toward eliminating safety violations during training.

We formalized our problem as reinforcement learning over a partially observable process with the non-stationary observations of the learning process. This two-level RL problem allows us to use a broad spectrum of tools available in sequential decision-making differing our approach from the existing in the literature. However, solving Problem 4 appears to be a daunting task, especially in the online setting as we intend. Hence, we will employ heuristic solutions, which may prove more effective than the quest to find an optimal solution. Intuitively, a gradual increase in $\boldsymbol{d}_k$ by assigning scheduling increasing reference values (denoted by $\boldsymbol{d}_k^{\text{ref}}$) would less likely lead to constraint violation due to stricter exploration constraints. However, fixing the schedule *a priori* limits the ability of the algorithm to react to constraint violations. To alleviate this issue we propose two approaches: PI Simmer (PI-controlled safety budget) and Q Simmer (Online Q-learning with non-stationary rewards). In both algorithms the intuition is similar. If the current accumulated costs are well below the safety budget $\boldsymbol{d}_k^{\text{ref}}$, then we are very safe and the safety budget can be further increased. If the accumulated costs are around the safety budget, then we could stay at the same level or increase the safety budget. If the current accumulated costs are well above the safety budget $\boldsymbol{d}_k^{\text{ref}}$, then the safety budget should be decreased. This intuition is based on our previous empirical observations, i.e., lower safety budgets would often cause lower bursts in accumulated safety costs. We hypothesize that the policy can quickly learn "extremely unsafe" actions thus providing low safety cost bursts for low safety budgets.

Algorithm 1: PI SIMMER (basic version)

Inputs: $\{\boldsymbol{d}_k^{\text{ref}}\}_{k=0}^{N-1}$ - safety budget schedule, hyper-parameters - $\boldsymbol{K}_p$, $\boldsymbol{K}_i$, $\delta\boldsymbol{d}$.
Set $\boldsymbol{d}_0 = \boldsymbol{d}_0^{\text{ref}}$;
for $k = 0, \dots, N-1$ **do**
  Perform one epoch of learning for the safe policy with $\boldsymbol{d}_k$ and get the statistic $\hat{g}\left(\boldsymbol{z}_T(\boldsymbol{d}_k)\right)$;
  Compute the error term $\boldsymbol{e}_k = \boldsymbol{d}_k^{\text{ref}} - \hat{g}\left(\boldsymbol{z}_T(\boldsymbol{d}_k)\right)$;
  Compute action $\boldsymbol{u}_k^{\text{raw}} = \underbrace{\boldsymbol{K}_p\boldsymbol{e}_k}_{\text{P-Part}} + \underbrace{\boldsymbol{K}_i \sum_{i=k-T_i}^{k} \boldsymbol{e}_i}_{\text{I-Part}}$;
  Set $\boldsymbol{d}_{k+1} = \text{clip}(\boldsymbol{u}_k^{\text{raw}}, -\delta\boldsymbol{d}, \delta\boldsymbol{d}) + \boldsymbol{d}_k$;
end

PI Simmer. The idea for the PI controller is quite intuitive it takes the error term $\boldsymbol{e} = \boldsymbol{d}^{\text{ref}} - c$ and uses it for action computation. P stands for proportional control and links the error terms with actions by multiplying the error term by the gain $\boldsymbol{K}$. The proportional part delivers a brute force control by having a large control magnitude for large errors, but it is not effective if the instantaneous error values become small. Proportional control cannot achieve zero error $\boldsymbol{e}$ tracking. This happens since zero error results in a zero proportional action and hence the control over the error is lost. To deliver zero error, integral control is typically used, which sums up previous error terms and uses this sum to determine actions instead of the error. The integral term can be seen as "action acceleration" (or "momentum") term which is large if the past errors are large. This can potentially cause unwanted behaviors (such as, oscillations) if the corresponding gain $\boldsymbol{K}_i$ is too large. Both gains $\boldsymbol{K}$, $\boldsymbol{K}_i$ are usually chosen by tuning. PI controller can solve many hard control problems, but there are some implementation and engineering tricks and improvements. We introduce a simplified version in Algorithm 1. We provide a full version of the algorithm and rules-of-thumb for the choice of hyper-parameters in our ablation studies in Appendix.

Q Simmer. Consider an MDP with the states $\{\boldsymbol{d}_0^{\text{ref}}, \dots, \boldsymbol{d}_{K-1}^{\text{ref}}\}$, which for simplicity of notation we denote $\{0, \dots, K-1\}$, with the actions $\{\boldsymbol{a}_{-1}, \boldsymbol{a}_0, \boldsymbol{a}_{+1}\}$, where the action $\boldsymbol{a}_{+1}$ moves the state $\boldsymbol{s} = i$ to the state $\boldsymbol{s} = i+1$, $\boldsymbol{a}_{-1}$ moves the state $\boldsymbol{s} = i$ to the state $\boldsymbol{s} = i-1$ and the action $\boldsymbol{a}_0$ does not transfer the state. Note that the action $\boldsymbol{a}_{+1}$ is defined for all $i < K-1$ and the action $\boldsymbol{a}_{-1}$ is defined for all $i > 0$. Our design for the rewards of this MDP is guided by the following intuition.

$$\begin{array}{lll} \text{We are not safe} & \text{We are borderline safe} & \text{We are very safe} \\ \text{if } \boldsymbol{s} - \boldsymbol{o} \le -\delta: & \text{if } |\boldsymbol{s} - \boldsymbol{o}| \le \delta: & \text{if } \boldsymbol{s} - \boldsymbol{o} \ge \delta: \\ \boldsymbol{r} = \begin{cases} 2 & \boldsymbol{a} = -1, \\ -1 & \boldsymbol{a} = 0, \\ -1 & \boldsymbol{a} = 1, \end{cases} & \boldsymbol{r} = \begin{cases} -1 & \boldsymbol{a} = -1, \\ 1 & \boldsymbol{a} = 0, \\ 1 & \boldsymbol{a} = 1, \end{cases} & \boldsymbol{r} = \begin{cases} -1 & \boldsymbol{a} = -1, \\ 1 & \boldsymbol{a} = 0, \\ 2 & \boldsymbol{a} = 1. \end{cases} \end{array} \tag{5}$$

where $\boldsymbol{o} = \hat{g}\left(\boldsymbol{z}_T(\boldsymbol{d}_k)\right)$ is a safety violation statistic (over an epoch) of the safe RL algorithm, δ is a significance violation threshold. We use a Q-learning update to learn the Q function:

$$\boldsymbol{Q}(\boldsymbol{s}_t, \boldsymbol{a}_t) = (1 - l_r)\boldsymbol{Q}(\boldsymbol{s}_t, \boldsymbol{a}_t) + l_r(\boldsymbol{r}_t + \max_{\boldsymbol{b}} \boldsymbol{Q}(\boldsymbol{s}_{t+1}, \boldsymbol{b})) \tag{6}$$

where l_r is the learning rate and get the action with ε-greedy exploration strategy:

$$\boldsymbol{a}_t = \begin{cases} \text{argmax}_{\boldsymbol{b}}\, \boldsymbol{Q}(\boldsymbol{s}_t, \boldsymbol{b}) & \text{with probability } \varepsilon, \\ \text{random} & \text{with probability } 1 - \varepsilon. \end{cases} \tag{7}$$

We summarize the approach as Algorithm 2.

3 Experiments

3.1 Baselines, environments and code base

Environments: We use the safe pendulum environment defined in [13], we also use the custom-made safety gym environment with deterministic constraints, which we call static point goal [48]. In this

Algorithm 2: Q-SIMMER

Input: $\{\boldsymbol{d}_k^{\text{ref}}\}_{k=0}^{K-1}$ - an *a priori* chosen state space; hyper-parameters - N, τ, l_r, δ, ε;
Initialize $\boldsymbol{d}_0 = \boldsymbol{d}_0^{\text{ref}}$
for $k = 0, \ldots, N-1$ **do**
  Perform one epoch of learning for the safe policy with $\boldsymbol{d}_k$ and get the statistic $\hat{g}\left(\boldsymbol{z}_T(\boldsymbol{d}_k)\right)$;
  Update the Q function as in Equation 6 while computing the reward $\boldsymbol{r}$ as in Equation 5;
  Compute an action $\boldsymbol{a}_k$ as in Equation 7 and compute $\boldsymbol{d}_{k+1}$ accordingly;
end

environment with 46 states and two actions, a large hazard circle is placed before the goal and forces the agent to go around it to reach the goal similarly to our cartoon in Figure 1. We provide additional details in Appendix. The rest of our tests are performed on the safety gym benchmarks [34].
Code base: Our code is based on two repositories: safety starter agents [34], and PID Lagrangian [41]. We have implemented the safety state augmentation following the description in [40]. We use default parameters for both code bases unless stated otherwise.
Computational resources: We performed all computations on a PC equipped with 512GB of RAM, two Intel Xeon E5 CPUs, and four 16GB NVIDIA Tesla V100 GPUs.

Baselines (•) and our algorithms (■):
• **CPO, Lagrangian PPO (L-PPO) and TRPO (L-TRPO).** Standard baselines from [34].
• **PID-Lagrangian.** An algorithm stabilizing L-PPO learning [41].
• **LAMBDA.** A model-based method showing great performance on Safety Gym [4].
• **PO-PPO** PPO-based algorithm with probability one constraints from [40].
■ **PI-Simmer** - Scheduling safety budget using PI controller for PO-PPO;
■ **Q-Simmer** - Scheduling safety budget using Q learning for PO-PPO;
■ **L-PPO (PID-L) w SA** - L-PPO (PID-L) solving Problem 2 with safety state augmentation.
■ **Simmer L-PPO (PID-L)** - L-PPO (PID-L) w SA and safety budget scheduling.

3.2 Improving Safety During Training for Pendulum Swing-Up

Improving safety during training is more suited for almost surely safe RL and we will take PO PPO as our baseline. In this setting, we can aim to reduce the number of individual trajectories that violate the constraints, and thus we can avoid estimating the statistic g. For PI Simmer we chose the following hyper-parameters $K = 0.01$, $K_i = 0.005$, $K_{\text{aw}} = 0.01$ and $\tau = 0.995$. The parameter K_{aw} resets the integral term avoiding the accumulation of error - the higher the value the more aggressive is reset. The parameter τ is the Polyak update for the error term, which pre-processes the error term for the PI controller. We discuss these parameters in detail in Appendix. We also note that except for an initial burst of violations both our approaches manage to keep the number of violations quite low. Overall we found that low values for K and K_i are beneficial to avoid overreaction to constraint violations. In this case, keeping the value of K_{aw} low is advisable as the action saturation does not occur too often. Finally, keeping τ close to one will force the controller to react to most of the constraint violations. For Q Simmer we chose $\delta = 1$, $\tau = 0.995$, $l_r = 0.05$, and $\varepsilon = 0.95$. It appears that a fast learning rate here can allow for learning, sufficiently fast forgetting of past rewards, but also to avoid catastrophic forgetting. As we consider a finite state MDP we can avoid using sophisticated techniques for online learning and use the simplest one — tuning learning rate. We chose δ and τ to avoid frequent state transitions.

We compare our algorithm to PO PPO in terms of the number of trajectories with constraint violations, and returns, and compare the progression of the schedule $\boldsymbol{d}_k^{\text{test}}$. Naturally, some individual trajectories still violate the constraint, but the number can be significantly reduced using Simmer RL as Figure 2 suggests. While PI Simmer outperforms Q Simmer in these runs, it is worth mentioning that PI Simmer uses more prior information than Q Simmer. Indeed, while composing a schedule is not hard, we still have to identify the switch points, which are learned by Q Simmer. We perform ablation studies on the parameters and discuss their choice in more detail in Appendix.

3.3 Guiding Exploration by Scheduling Safety Constraints

We now turn our attention to safe RL with constraints imposed on average costs. We test the performance of Simmer L-PPO and L-PPO w SA on the swing-up pendulum environment and present

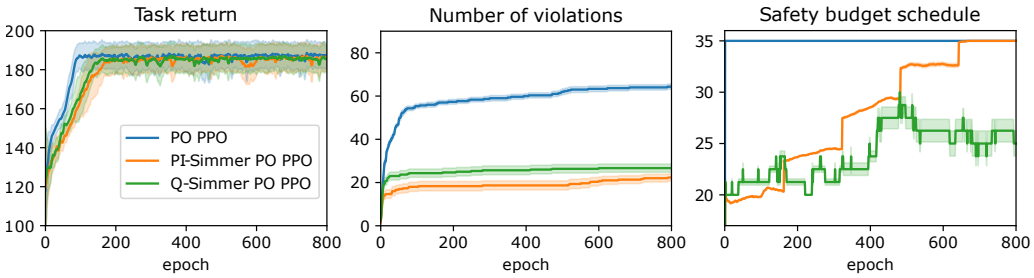

(a) Returns (b) Number of violations (c) Safety budgets

Figure 2: Reducing safety violations during training for Safe RL with constraints almost surely.

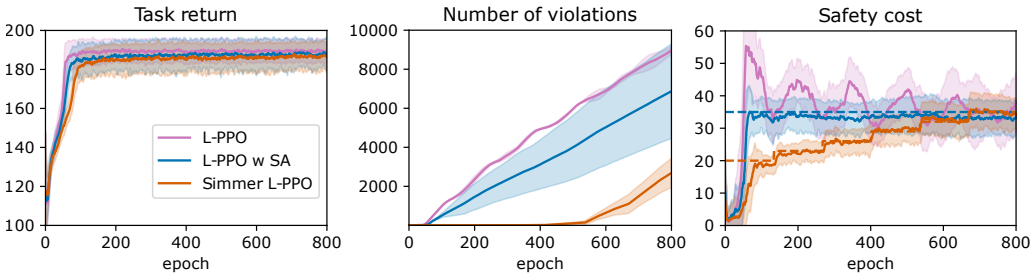

(a) Returns (b) Number of violations (c) Safety costs

Figure 3: Comparison of Simmer L-PPO, L-PPO with and without safety state augmentation. Mean returns and cost are computed over a hundred different trajectories obtained for three different seeds. Shaded areas represent standard deviations.

training results in Figure 3. Here we use safety starter agents as a base learner for all algorithms. The major observations are that L-PPO w SA delivers *almost no constraint violations* with respect to the mean cost estimate, and therefore using Q Simmer and PI Simmer becomes redundant. In the meantime, L-PPO has even trouble to converge. Note that we have used *the same hyper-parameters for all algorithms*, which are default parameters in safety starter agents and the learning rate 0.03. While the behavior of the L-PPO algorithm can certainly be improved with tuning, we note that simply augmenting the safety state leads to improved performance as well as stability of the algorithm. Further, we observe that Simmer L-PPO leads to a fewer number of violations, however, the rate of violations for the safety budget of 35 is fairly similar to L-PPO w SA.

Interestingly, a similar picture occurs with more advanced baselines such as PID Lagrangian [41] and more complicated environments. Here, we used the static point goal environment designed in [48], where a large static hazard region is placed in front of the goal, which is similar to our motivational example in Figure 1. The results for PID-L and PID-L with state augmentation are depicted in Figures 4a and 4b, respectively, suggest that the presence of the safety state stabilizes training and leads to a more consistent constraint satisfaction for different safety budgets d. Note that hyper-parameters for all the runs are the same for both algorithms. We further apply the naïve simmer approach to both baselines with results depicted in Figure 4c. In both cases, the safety budget takes values 1, 5, 10, 15, and 20, and increased after equal time intervals. Note that in both cases now training curves are quite stable, although state augmentation delivers an extra boost. In all our experiments we used the same hyper-parameters for all versions of PID-L, i.e., $K = 0.1$, $K_i = 0.01$, $\gamma_l = 0.99$.

Overall, our experiments suggest that safety during training with constraints imposed on average costs becomes a much easier problem with safety state augmentation. Indeed, even intuitively every outlier trajectory with a large constraint violation should bias the average cost and should instruct the algorithm not to follow this path. We note that in both cases above we have sparse costs, i.e., the agent encounters unsafe regions and incurs costs while navigating toward the goal. While an algorithm without the safety state will receive information on constraint violation after an episode, the safety state would constantly inform the algorithm of the distance toward a violation. This is one of the reasons why safety state augmentation can learn the task better. Note that simmering

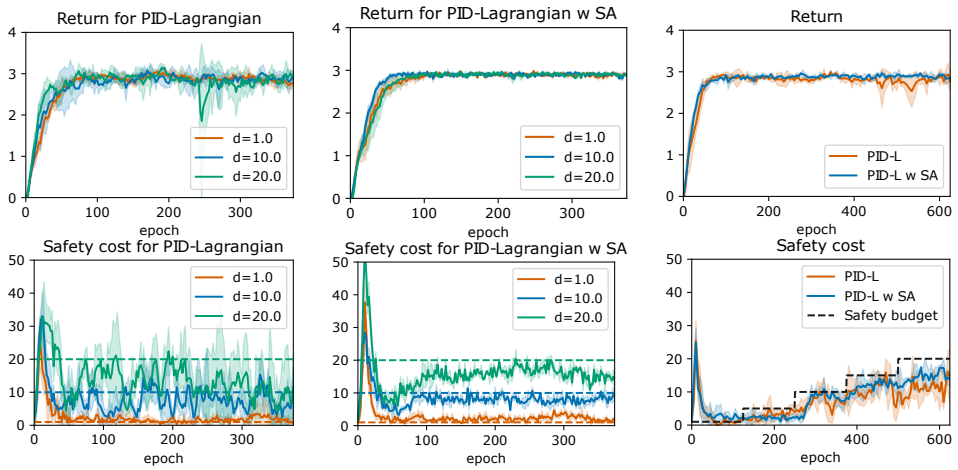

(a) PID-L. (b) PID-L w SA. (c) Simmer RL.

Figure 4: Comparing PID Lagrangian with and without safety state augmentation. d is the safety budget used in training. All the curves are means and shaded areas are standard deviations computed over three seeds. These results suggest that safety state augmentation can stabilize training and deliver safer solutions.

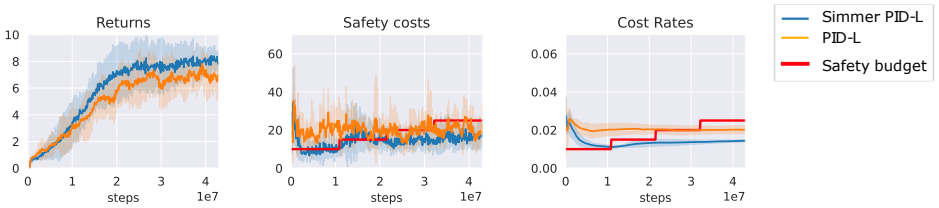

(a) CarPush1

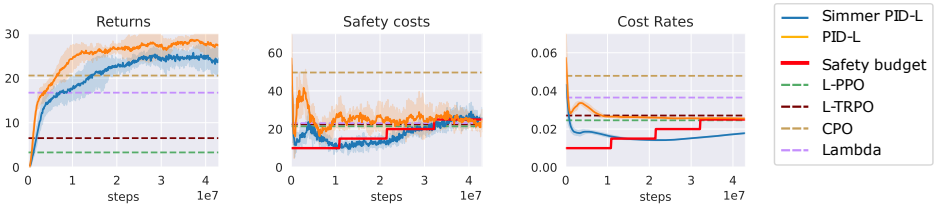

(b) CarGoal1

Figure 5: Results for the car robot in the goal and point environments. Simmer PID-L outperforms all baselines in terms of the cost rate, which is a measure of safety during training. At the same time, Simmer PID-L is competitive with PID-L in terms of the returns or safety costs, while significantly outperforming other baselines. The curves are means and shaded areas are standard deviations computed over five runs with random seeds. Note that we have baseline results only for the CarGoal1 environment.

additionally offers fewer constraint violations while training. These results suggest that simmering safe RL together with state augmentation delivers overall safer solutions with stable training curves.

3.4 Tests on safety gym benchmarks

We finally test our approach on more challenging environments from the safety gym benchmark: CarPush1 and CarGoal1 in Figures 5a, 5b, respectively, and PointPush1, PointGoal1, PointButton1, CarButton1 in Figures A4a, A4b, A4c and A5c respectively in Appendix. The hyper-parameters and the tuning specifics can be found in Appendix. Again we use a naïve schedule increasing the safety budget from 10 to 25 in 5 unit increments. We observe that Simmer PID-L and PID-L outperform

other baselines significantly in terms of the return while delivering safe policies. Simmer PID-L does perform very similarly to PID-L in terms of returns but generally outperforms PID-L in terms of safety costs and cost rates. We note that our results for PointGoal1 are consistent with the results from [41].

4 Conclusion

We augment the safety information into the state space and show how we can effectively use it to improve safe exploration. The safety state is nonnegative if and only if the constraint is satisfied and therefore it can serve as a distance toward constraint violations. We argue that the optimal policy must depend on this information to achieve safe performance. We validate this argument using intuitive examples, references to theoretical results, and experiments. We do not doubt that safety state augmentation is needed for effective safe RL.

Safety state augmentation and simmering show superior performance on pendulum swing-up and static point goal tasks for average constrained problems. We further tuned PID Lagrangian to show very strong performance on safety gym benchmarks, which has not been previously reported. Simmer PID Lagrangian shows competitive performance in terms of returns to PID Lagrangian and outperforms all baselines (including PID Lagrangian) in terms of the cost rate and the costs. Further, using state augmentation for the average constrained problems appears to be quite beneficial as well. Scheduling the safety budget can stabilize safe algorithms without state augmentation, but using both state augmentation and simmering improves performance and safety. We achieve this performance and safety boost at the expense of sample efficiency since we effectively learn a series of safe policies.

Simmering RL algorithms with probability one constraints can significantly reduce safety violations during training. We illustrate this feature by developing two algorithms for safe learning. The first one is based on a PI controller and allows for adjustment of a pre-defined learning schedule depending on the estimated training cost. The second one is based on online Q learning and learns to adjust the safety budget automatically. Both approaches have some advantages and limitations. PI simmering is more effective if there is a reasonable safety budget schedule, while Q simmering requires less prior knowledge to learn.

Our algorithmic development focused on probability one constraints for two reasons. First, since probability one constraints have to be imposed *on all trajectories* we can simply count the number of violations and do need to resort to empirical estimates of the mean cost (as we need in the average constrained problem). Second, it appears that in the average constrained case with "simple" environments (such as pendulum swing-up and static point goal), we can track the constraint quite efficiently after the initial (and unavoidable) burst in constraint violations. Thus our experiments suggest that PI Simmer and Q Simmer are simply redundant for the average constrained RL.

Augmenting the safety state and scheduling the safety budget does not solve all the problems in safe exploration. First, an initial burst of constraint violations is an inevitable reality of using our model-free approach. It appears that lowering the initial safety budget lowers the initial burst of constraint violations, but does not completely solve the problem. We do not see how to address this limitation without making further assumptions about the environment. Second, it is noticeable that the performance on the PointPush1 environment is quite noisy for both PID-L and Simmer PID-L. This is because some seeds achieve very good performance (return of 15) and some seeds do not (return of 5), while the learning curves appear to be stable. This suggests that the learning procedure is at a local maximum. In our point of view, this calls for a more sophisticated algorithm for exploration seeking maximal return. This may help achieve stable performance in this environment. Beyond these limitations, one can also list safety violations caused by variance and disturbances, lack of provable safety guarantees, etc. Studying these limitations is outside of the scope of this paper.

We have tested PI-Simmer and Q-Simmer on a rather simple in terms of the state and action spaces environment. We foresee that in a more complex setting the main problem would be learning the agent's behavior, but our algorithmic development would remain valid. Nevertheless, it would be interesting to extend this approach to a more complex setting.

References

[1] Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. In *International Conference on Machine Learning*, pages 22–31, 2017. (p. 2)

[2] Anayo K Akametalu, Jaime F Fisac, Jeremy H Gillula, Shahab Kaynama, Melanie N Zeilinger, and Claire J Tomlin. Reachability-based safe learning with Gaussian processes. In *IEEE Conference on Decision and Control*, pages 1424–1431, 2014. (p. 2)

[3] Eitan Altman. *Constrained Markov decision processes*, volume 7. CRC Press, 1999. (p. 1)

[4] Yarden As, Ilnura Usmanova, Sebastian Curi, and Andreas Krause. Constrained policy optimization via Bayesian world models. *arXiv preprint arXiv:2201.09802*, 2022. (p. 2), (p. 6)

[5] Karl Johan Åström. *Advanced PID control*, volume 461. The Instrumentation, Systems, and Automation Society, 2006. (p. 2)

[6] Felix Berkenkamp, Riccardo Moriconi, Angela P Schoellig, and Andreas Krause. Safe learning of regions of attraction for uncertain, nonlinear systems with Gaussian processes. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 4661–4666. IEEE, 2016. (p. 2)

[7] Felix Berkenkamp, Matteo Turchetta, Angela Schoellig, and Andreas Krause. Safe model-based reinforcement learning with stability guarantees. In *Advances in Neural Information Processing Systems*, pages 908–918, 2017. (p. 2)

[8] Homanga Bharadhwaj, Aviral Kumar, Nicholas Rhinehart, Sergey Levine, Florian Shkurti, and Animesh Garg. Conservative safety critics for exploration. *arXiv preprint arXiv:2010.14497*, 2020. (p. 1)

[9] Agustin Castellano, Hancheng Min, Juan Bazerque, and Enrique Mallada. Reinforcement learning with almost sure constraints. *arXiv preprint arXiv:2112.05198*, 2021. (p. 2)

[10] Richard Cheng, Gábor Orosz, Richard M Murray, and Joel W Burdick. End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks. *arXiv preprint arXiv:1903.08792*, 2019. (p. 2)

[11] Yinlam Chow, Mohammad Ghavamzadeh, Lucas Janson, and Marco Pavone. Risk-constrained reinforcement learning with percentile risk criteria. *The Journal of Machine Learning Research*, 18(1):6070–6120, 2017. (p. 2)

[12] Yinlam Chow, Ofir Nachum, Edgar Duenez-Guzman, and Mohammad Ghavamzadeh. A Lyapunov-based approach to safe reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 8092–8101, 2018. (p. 2)

[13] Alexander I Cowen-Rivers, Daniel Palenicek, Vincent Moens, Mohammed Amin Abdullah, Aivar Sootla, Jun Wang, and Haitham Bou-Ammar. SAMBA: Safe model-based & active reinforcement learning. *Machine Learning*, pages 1–31, 2022. (p. 2), (p. 5)

[14] Gal Dalal, Krishnamurthy Dvijotham, Matej Vecerik, Todd Hester, Cosmin Paduraru, and Yuval Tassa. Safe exploration in continuous action spaces. *arXiv preprint arXiv:1801.08757*, 2018. (p. 2)

[15] AN Daryin and AB Kurzhanski. Nonlinear control synthesis under double constraints. *IFAC Proceedings Volumes*, 38(1):247–252, 2005. (p. 2)

[16] S. Dean, S. Tu, N. Matni, and B. Recht. Safely learning to control the constrained linear quadratic regulator. In *American Control Conference*, pages 5582–5588, 2019. (p. 2)

[17] Jonas Degrave, Federico Felici, Jonas Buchli, Michael Neunert, Brendan Tracey, Francesco Carpanese, Timo Ewalds, Roland Hafner, Abbas Abdolmaleki, Diego de Las Casas, et al. Magnetic control of tokamak plasmas through deep reinforcement learning. *Nature*, 602(7897):414–419, 2022. (p. 1)

[18] Benjamin Eysenbach, Shixiang Gu, Julian Ibarz, and Sergey Levine. Leave no trace: Learning to reset for safe and autonomous reinforcement learning. In *International Conference on Learning Representations*, 2018. (p. 1)

[19] J. F. Fisac, A. K. Akametalu, M. N. Zeilinger, S. Kaynama, J. Gillula, and C. J. Tomlin. A general safety framework for learning-based control in uncertain robotic systems. *IEEE Transactions on Automatic Control*, 64(7):2737–2752, 2019. (p. 2)

[20] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pages 1861–1870, 2018. (p. 2)

[21] Nils Jansen, Bettina Könighofer, Sebastian Junges, AC Serban, and Roderick Bloem. Safe reinforcement learning using probabilistic shields. 2020. (p. 2)

[22] Chi Jin, Akshay Krishnamurthy, Max Simchowitz, and Tiancheng Yu. Reward-free exploration for reinforcement learning. In *International Conference on Machine Learning*, pages 4870–4879. PMLR, 2020. (p. 1)

[23] Gregory Kahn, Adam Villaflor, Vitchyr Pong, Pieter Abbeel, and Sergey Levine. Uncertainty-aware reinforcement learning for collision avoidance. *arXiv preprint arXiv:1702.01182*, 2017. (p. 2)

[24] Sanket Kamthe and Marc Deisenroth. Data-efficient reinforcement learning with probabilistic model predictive control. In *International Conference on Artificial Intelligence and Statistics*, pages 1701–1710. PMLR, 2018. (p. 2)

[25] Torsten Koller, Felix Berkenkamp, Matteo Turchetta, and Andreas Krause. Learning-based model predictive control for safe exploration. In *IEEE Conference on Decision and Control*, pages 6059–6066, 2018. (p. 2)

[26] Zuxin Liu, Zhepeng Cen, Vladislav Isenbaev, Wei Liu, Zhiwei Steven Wu, Bo Li, and Ding Zhao. Constrained variational policy optimization for safe reinforcement learning. *arXiv preprint arXiv:2201.11927*, 2022. (p. 2)

[27] Yecheng Jason Ma, Andrew Shen, Osbert Bastani, and Jayaraman Dinesh. Conservative and adaptive penalty for model-based safe reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 5404–5412, 2022. (p. 2)

[28] David Mguni, Joel Jennings, Taher Jafferjee, Aivar Sootla, Yaodong Yang, Changmin Yu, Usman Islam, Ziyan Wang, and Jun Wang. DESTA: A framework for safe reinforcement learning with markov games of intervention. *arXiv preprint arXiv:2110.14468*, 2021. (p. 2)

[29] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013. (p. 1)

[30] Motoya Ohnishi, Li Wang, Gennaro Notomista, and Magnus Egerstedt. Barrier-certified adaptive reinforcement learning with applications to brushbot navigation. *IEEE Transactions on Robotics*, 35(5):1186–1205, 2019. (p. 2)

[31] Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *International Conference on Machine Learning (ICML)*, pages 2778–2787, 2017. (p. 1)

[32] Laurent Pfeiffer. Two approaches to stochastic optimal control problems with a final-time expectation constraint. *Applied Mathematics & Optimization*, 77(2):377–404, 2018. (p. 3)

[33] Kyriakos Polymenakos, Nikitas Rontsis, Alessandro Abate, and Stephen Roberts. SafePILCO: A software tool for safe and data-efficient policy synthesis. In Marco Gribaudo, David N. Jansen, and Anne Remke, editors, *Quantitative Evaluation of Systems*, pages 18–26. Springer International Publishing, 2020. (p. 2)

[34] Alex Ray, Joshua Achiam, and Dario Amodei. Benchmarking safe exploration in deep reinforcement learning, 2019. (p. 2), (p. 6), (p. 13)

[35] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International Conference on Machine Learning*, pages 1889–1897, 2015. (p. 2)

[36] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017. (p. 2)

[37] Philipp Schwartenbeck, Thomas FitzGerald, Ray Dolan, and Karl Friston. Exploration, novelty, surprise, and free energy minimization. *Frontiers in psychology*, 4:710, 2013. (p. 1)

[38] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016. (p. 1)

[39] Thiago D Simão, Nils Jansen, and Matthijs TJ Spaan. Alwayssafe: Reinforcement learning without safety constraint violations during training. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2021. (p. 2)

[40] Aivar Sootla, Alexander I. Cowen-Rivers, Taher Jafferjee, Ziyan Wang, David Mguni, Jun Wang, and Haitham Bou-Ammar. SAUTE RL: Almost surely safe reinforcement learning using state augmentation, 2022. (p. 2), (p. 3), (p. 6)

[41] Adam Stooke, Joshua Achiam, and Pieter Abbeel. Responsive safety in reinforcement learning by pid lagrangian methods. In *International Conference on Machine Learning*, pages 9133–9143. PMLR, 2020. (p. 2), (p. 6), (p. 7), (p. 9), (p. 13)

[42] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018. (p. 1)

[43] MuZero Applied Team. Muzero's first step from research into the real world. In *DeepMind blog*. Deepmind, Feb 2022. https://www.youtube.com/watch?v=ZTnySFuJGiM. (p. 1)

[44] Matteo Turchetta, Felix Berkenkamp, and Andreas Krause. Safe exploration in finite Markov decision processes with Gaussian processes. In *Advances in Neural Information Processing Systems*, pages 4312–4320, 2016. (p. 2)

[45] Matteo Turchetta, Andrey Kolobov, Shital Shah, Andreas Krause, and Alekh Agarwal. Safe reinforcement learning via curriculum induction. *arXiv preprint arXiv:2006.12136*, 2020. (p. 1), (p. 2)

[46] Richard B Vinter. *Optimal control*. Springer, 2010. (p. 3)

[47] Akifumi Wachi, Yanan Sui, Yisong Yue, and Masahiro Ono. Safe exploration and optimization of constrained MDPs using Gaussian processes. In *AAAI Conference on Artificial Intelligence*, 2018. (p. 2)

[48] Qisong Yang, Thiago D Simão, Simon H Tindemans, and Matthijs TJ Spaan. WCSAC: Worst-case soft actor critic for safety-constrained reinforcement learning. In *AAAI Conference on Artificial Intelligence.*, 2021. (p. 2), (p. 5), (p. 7)

[49] Jesse Zhang, Brian Cheung, Chelsea Finn, Sergey Levine, and Dinesh Jayaraman. Cautious adaptation for reinforcement learning in safety-critical settings. In *International Conference on Machine Learning*, pages 11055–11065. PMLR, 2020. (p. 2)

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] We reported the error bars over all the runs across different seeds
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] We use a PC equipped with 512GB of RAM, two Intel Xeon E5 CPUs, and four 16GB NVIDIA Tesla V100 GPUs
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] We use the code from safety starter agents, safety gym by open AI [34], PID Lagrangian method [41].
   (b) Did you mention the license of the assets? [No] All the libraries are under MIT Licence.
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]