

Provably Safe Reinforcement Learning: Conceptual Analysis, Survey, and Benchmarking

Anonymous authors

Paper under double-blind review

Abstract

Ensuring safety of reinforcement learning (RL) algorithms is crucial to unlock their potential for many real-world tasks. However, vanilla RL does not guarantee safety. In recent years, several methods have been proposed to provide formal safety guarantees for RL. Yet, there is no comprehensive comparison of these provably safe RL methods. We therefore introduce a categorization of existing provably safe RL methods, present the conceptual foundations for both continuous and discrete action spaces, and benchmark existing methods empirically. The methods are categorized based on how the action is adapted by the safety method: action replacement, action projection, and action masking. Our experiments on an inverted pendulum and a quadrotor stabilization task indicate that action replacement is one of the best performing approaches in all settings despite its comparatively simple formulation. Lastly, we give practical guidance on the selection of provably safe RL approaches depending on the safety specification, RL algorithm, and type of action space.

1 Introduction

Reinforcement learning (RL) contributes to many recent advancements in challenging research fields, such as robotics (El-Shamouty et al., 2020; Zhao et al., 2020), autonomous systems (Kiran et al., 2021; Ye et al., 2021), and games (Mnih et al., 2013; Silver et al., 2017). However, safety guarantees are crucial to train and deploy RL agents in the real world. Without them, it is unclear if the RL agent might harm humans or seriously damage the environment or itself. A vanilla RL agent typically explores unsafe or dangerous actions multiple times to learn how to achieve the highest possible reward. Consequently, we cannot train such systems in the real world for safety-critical tasks. Even if the agent learns in a controlled environment, such as a simulation where the execution of unsafe actions is of no concern, there are no safety guarantees for deploying the trained agent in the real world. Therefore, safe RL emerged that adapts the learning process such that the agent considers safety aspects next to performance during training and operation. In some areas, failures are disastrous, e.g., autonomous driving, human-robot collaboration, or space robotics. Therefore, these areas require strict safety guarantees. We refer to RL methods that provide these formal safety guarantees for training and operation as *provably safe RL*.

In this paper, we provide for the first time a consistent conceptual framework for provably safe RL in both continuous and discrete action spaces, a comprehensive literature survey, and a comparison between provably safe RL approaches on two benchmarks. The characteristic difference between provably safe RL approaches is how they adapt the actions of the agent. Therefore, we propose to classify them into three categories: *action replacement*, *action projection*, and *action masking*. Our experimental evaluation confirms that all three provably safe RL types guarantee safety and shows that action replacement exhibits the most robust performance across all tested environments and on the investigated RL algorithms.

Our contributions in this work are fourfold. First, we introduce a comprehensive classification of provably safe RL methods and their formal description. This categorization allows us to compare and benchmark the effects of choosing a specific type of action modification on the ability of agents to learn. Second, we are the first to formulate action masking for continuous action spaces. Third, we provide a comprehensive survey and show how previous provably safe RL works fall into the three categories. Finally, we validate the efficacy

and assess the performance of the provably safe RL methods on two common control benchmarks. This comparison provides insights into the strength and weaknesses of the different provably safe RL approaches and allows us to provide some advice on selecting the best-suited provably safe RL approach for a specific problem independent of the safety verification method used.

The remainder of this paper is structured as follows. First, we briefly overview the historical development of safe RL and relate provably safe RL to it in section 1.1. We describe preliminary concepts in section 2 and introduce our proposed categorization. We then show how the related provably safe RL literature fits our categorization in section 3. Section 4 compares the different provably safe RL categories experimentally on a two-dimensional (2D) quadrotor stabilization task. Section 5 discusses the results of our experimental evaluation and the practical considerations following them. Finally, we conclude this work in section 6.

1.1 Evolution towards provably safe RL

The notion of risk and safety in RL has existed since the 1990s (Heger, 1994). The reasons for combining safety and RL were to improve the convergence speed and focus the learning on relevant or safe regions. Thus, the field of safe RL started developing, and in 2015, García & Fernández (2015) were the first to cluster safe RL. They provide two high-level categories: approaches that modify the optimization criterion and approaches that modify the exploration with safety aspects. Provably safe RL approaches best fall in the category *teacher advice*, a subcategory of approaches modifying the exploration. However, the teacher advice category also includes approaches that provide no strict but probabilistic or non-formal safety guarantees. Thus, the notion of safety for the teacher advice category differs widely and includes many methodologies. Additionally, since 2015 significant advances in model-free RL and the increased applicability of deep RL amplified the need for formal guarantees in safe RL. A more recent work by Brunke et al. (2022) classifies safe learning approaches in robotics, including RL, by their safety constraints, which can be *soft*, *probabilistic*, and *hard*. We use their classification in the following to locate provably safe RL in the safe RL field.

Soft constraint approaches consider safety directly in their optimization objective. Here, the agent can explore all actions and states regardless of safety. Thus, these methods can be unsafe during training, especially in the beginning, but converge to a safer policy without formal safety guarantees after sufficient training steps. Most advances have been made in constrained RL (Altman, 1998; Achiam et al., 2017; Stooke et al., 2020), for which the policy aims to maximize the reward while satisfying user-defined specifications. The specifications can be formulated as constraint functions (Chow et al., 2018; Stooke et al., 2020; Yang et al., 2020; Marvi & Kiumarsi, 2021) or as temporal logic formulas (De Giacomo et al., 2021; Hasanbeig et al., 2020; 2019b;a). The main advantage of these methods is that no explicit model of the agent dynamics or the environment is required as the agent learns the safety aspects through experience. Thus, such safe RL methods have a high potential in non-critical settings, where unsafe actions do not cause major damage.

Probabilistic constraint approaches rely on probabilistic models or synthesize a model from sampled data. Here, the action and state space can be restricted based on probabilities. Nonetheless, unsafe actions are sometimes not detected and might occur occasionally. Several works (Turchetta et al., 2016; Berkenkamp et al., 2017; Mannucci et al., 2018) try to determine the maximal set of safe states by starting from an often user-defined conservative set and extending it with the gathered learning experience. Other methods (Könighofer et al., 2021; Thananjeyan et al., 2021; Dalal et al., 2018; Zanon & Gros, 2021; Yang et al., 2021; Gillula & Tomlin, 2013) are based on formulating probabilistic models that identify the probability of safety for an action. In general, approaches that rely on probabilistic methods are especially applicable if one cannot bind measurement errors, modeling errors, and disturbances by sets.

Provably safe RL defines hard constraints, which are fulfilled by integrating prior system knowledge into the learning process. Here, the agent only explores safe actions and only reaches states fulfilling the safety specifications. Provably safe RL already fulfills the given safety specifications during the learning process, which is essential when training or fine-tuning agents on safety-critical tasks in the physical world. Thus, we exclude approaches that only verify learned policies (Bastani et al., 2018; Schmidt et al., 2021) from our survey. We focus on model-free RL algorithms that do not explicitly learn or use a model of the system dynamics to optimize the policy. Generally, deploying learned controllers in the physical world became increasingly realistic in recent years, and thus, the need for provably safe RL grew, and more provably safe

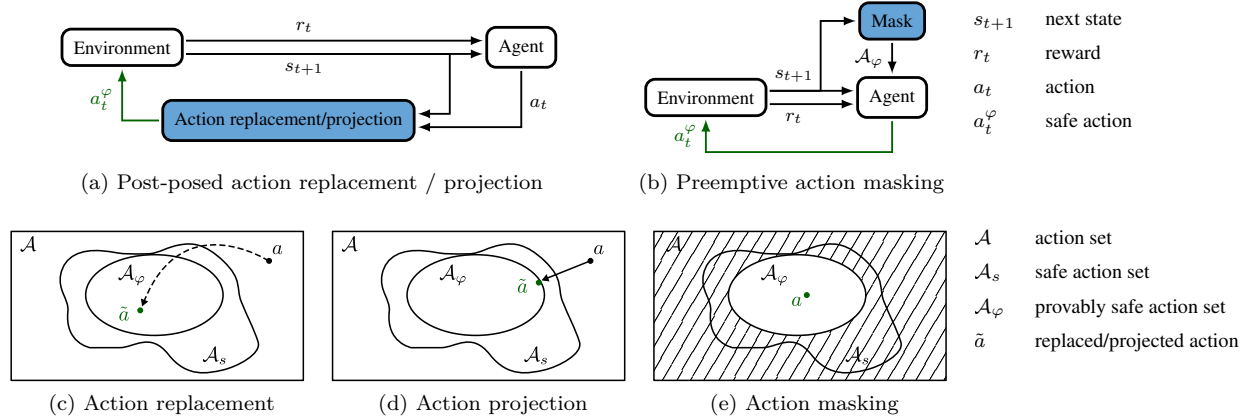


Figure 1: Structure of the three types of provably safe RL methods. The post-posed action replacement or projection methods (a) alter unsafe actions before sending them to the environment. In contrast, preemptive action masking approaches (b) allow the agent to only choose from the safe action space and therefore only output safe actions to the environment. Figures (c-e) highlight the differences of the three approaches in the action space. Here, action replacement (c) replaces unsafe actions with actions from the safe action space, action projection (d) projects unsafe actions to the closest safe action, and action masking (e) lets the agent choose solely from the safe action set.

RL approaches were developed. With this work, we aim to structure and provide practical insights into this growing field.

2 Conceptual Analysis

We introduce three provably safe RL classes by providing their formal notation in one comprehensive conceptual framework. This framework clarifies the differences between the three classes and eases the following literature review and benchmarking.

Markov decision process The RL agent learns on a Markov decision process (MDP) that is described by the tuple $(\mathbb{S}, \mathbb{A}, T, r, \gamma)$. Hereby, we assume that the set of states \mathbb{S} is fully observable with bounded precision. Partially observable MDPs can be handled using methods like particle filtering (Sunberg & Kochenderfer, 2018) and are not further discussed in this work. Both the action space \mathbb{A} and state space \mathbb{S} can be either continuous or discrete. $T(\mathbf{s}, \mathbf{a}, \mathbf{s}')$ is the transition function, which in the discrete case gives the probability that the transition from state \mathbf{s} to state \mathbf{s}' occurs by taking action \mathbf{a} . In the continuous case, $T(\mathbf{s}, \mathbf{a}, \mathbf{s}')$ denotes the probability density function of the transition. We assume that the transition function is stationary over time. For each transition, the agent receives a reward $r : \mathbb{S} \times \mathbb{A} \rightarrow \mathbb{R}$ from the environment. Finally, the discount factor $0 < \gamma < 1$ weights the relevance of future rewards.

Safety of a system For provably safe RL, it is required that the safety of states and actions are verifiable. Otherwise no formal claims about the safety of a system can be made. Thus, we first introduce the set of provably safe states $\mathbb{S}_{\varphi} \subseteq \mathbb{S}$ that contains all states in which all safety constraints are fulfilled. For verifying the safety of actions, we use a safety function $\varphi : \mathbb{S} \times \mathbb{A} \rightarrow \{0, 1\}$

$$\varphi(\mathbf{s}, \mathbf{a}) = \begin{cases} 1, & \text{if } (\mathbf{s}, \mathbf{a}) \text{ is verified safe} \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

It is clear that $\varphi(\mathbf{s}, \mathbf{a})$ requires some prior knowledge about the system dynamics. We can provide this knowledge with an abstract model that contains all safety relevant behaviors of the original model. These abstract safety models can get different inputs than the RL agent, and are in general less complex than the

underlying MDP. In systems where such a safety model is unavailable, provably safe RL is not applicable, and only non-provably safe approaches as discussed in section 1 can be used. Although $\varphi(\mathbf{s}, \mathbf{a})$ only verifies one action \mathbf{a} , it may take more than the next state into account, e.g., through an model predictive control (MPC) formulation.

Based on the safety function, we define a set of provably safe actions $\mathbb{A}_\varphi(\mathbf{s}) = \{\mathbf{a} | \varphi(\mathbf{s}, \mathbf{a}) = 1\}$ for a state \mathbf{s} , which is a subset of all possible safe actions $\mathbb{A}_\mathbf{s}(\mathbf{s})$, i.e., $\mathbb{A}_\varphi(\mathbf{s}) \subseteq \mathbb{A}_\mathbf{s}(\mathbf{s}) \subseteq \mathbb{A}^1$. All provably safe RL approaches rely on the availability of provably safe actions, so we require assumption 1.

Assumption 1 *There is at least one provably safe initial state $\mathbf{s}_0^\varphi \in \mathbb{S}_\varphi$ and for all safe states there exists at least one safe action $\forall \mathbf{s}^\varphi \in \mathbb{S}_\varphi \rightarrow \mathbb{A}_\varphi(\mathbf{s}^\varphi) \neq \emptyset$.*

With assumption 1, it is ensured that only provably safe states $\mathbf{s}^\varphi \in \mathbb{S}_\varphi$ can be reached when starting from any \mathbf{s}_0^φ and taking only provably safe actions thereafter.

Provably safe RL relies on model knowledge to provide safety guarantees, i.e., a conformant model that covers the safety-relevant system and environment dynamics. Hereby, the verification process can use an abstraction of the real system as long as it is conformant (Roehm et al., 2019) to the real system, i.e., it over-approximates both aleatoric and epistemic uncertainties, and covers all relevant safety aspects. This eases efficient verification, as the complexity of the abstraction is usually significantly lower than the complexity of the real system. In practice, we often weaken the specifications of provable safety to legal or passive safety. Hereby, inevitable safety violations caused by other agents are not considered to be the fault of the agent and are therefore not considered in this work. Examples of proving legal safety have been presented for autonomous driving (Pek et al., 2020) and in robotics (Bouraine et al., 2012).

There are multiple ways to ensure provable safety for RL systems, which we summarize in the three categories: action replacement, where the safety method replaces all unsafe actions from the agent with safe actions, action projection, which projects unsafe actions to the safe action space, and action masking, where the agent can only choose actions from the safe action space. We choose this categorization as it represents the three main approaches found in the literature to modify actions and thereby ensure safety for RL. Action replacement and action projection alter the action after it is outputted by the agent, so we refer to them as *post-posed* methods. Whereas, action masking only lets the agent choose from the safe action space; it is therefore a *preemptive* safety measure. Figure 1 displays the basic concept and structure of these methods. The following subsections describe the concept, mathematical formalization, required assumptions, and practical implications of the three approaches.

2.1 Action replacement

The first approach to ensure the safety of actions is to replace any unsafe action outputted by the agent with a safe action before its execution. The first step of action replacement is to evaluate the safety of the action $\mathbf{a} \in \mathbb{A}$ using $\varphi(\mathbf{s}, \mathbf{a})$. Conceptually, this is mostly done by over-approximating the set of states that are reachable by taking action \mathbf{a} in state \mathbf{s} , and then validating if the reachable set of states is a subset of \mathbb{S}_φ . We discuss the concrete verification methods used by previous action replacement works in section 3. If the action sampled from the policy $\pi(\mathbf{a}|\mathbf{s})$ is not verified as safe, it is replaced with a provably safe replacement action $\tilde{\mathbf{a}} = \psi(\mathbf{s})$, where $\psi : \mathbb{S} \rightarrow \mathbb{A}_\varphi$ is called replacement function. Following this procedure, it is guaranteed that only safe actions \mathbf{a}^φ with

$$\mathbf{a}^\varphi = \begin{cases} \mathbf{a} \sim \pi(\mathbf{a}|\mathbf{s}), & \text{if } \varphi(\mathbf{s}, \mathbf{a}) = 1 \\ \psi(\mathbf{s}), & \text{otherwise} \end{cases} \quad (2)$$

are executed. We discuss how this action replacement alters the MDP in the Appendix and additionally refer the interested reader to Hunt et al. (2021).

There are two general replacement functions found in literature, *sampling* and *failsafe*. In sampling, the replacement function $\psi_{\text{sample}}(\mathbf{s})$ uniformly samples a random action from $\mathbb{A}_\varphi(\mathbf{s})$. The other approach is

¹Please note that “taking no action” is commonly considered to be part of the action space, most often with the action $\mathbf{a} = [0, \dots, 0]^\top$.

to use a backup failsafe controller $\psi_{\text{failsafe}}(\mathbf{s})$ as replacement action, which could also stem from human feedback. In time-critical and complex scenarios, where building $\mathbb{A}_\varphi(\mathbf{s})$ online becomes too time-consuming, $\psi_{\text{failsafe}}(\mathbf{s})$ is the only available option.

When replacing an unsafe action, the agent’s training can be conducted with four possible learning tuples:

- *naive* - learning based on the action outputted by the policy network of the agent \mathbf{a} and the reward $r(\mathbf{s}, \mathbf{a}^\varphi)$ corresponding to the executed action $(\mathbf{s}, \mathbf{a}, \mathbf{s}', r)$. This ensures that the agent is updated according to its current policy. Learning with the original action should benefit on-policy learning, where the policy is updated based on experience collected using the most recent policy. However, utilizing the *naive* tuple for training is subject to state transitions that rely heavily on the replacement function $\psi(\mathbf{s})$, as demonstrated in the Appendix. This can be viewed as a form of noise on the observation-action relationship, thereby possibly hindering effective learning.
- *adaption penalty* - is *naive* with a penalty $r^*(\mathbf{s}, \mathbf{a}, \mathbf{a}^\varphi) = r(\mathbf{s}, \mathbf{a}^\varphi) + r_{\text{penalty}}(\mathbf{s}, \mathbf{a}, \mathbf{a}^\varphi)$ if an unsafe action was selected $(\mathbf{s}, \mathbf{a}, \mathbf{s}', r^*)$.
- *safe action* - learning based on the safe and possibly adapted action and the corresponding reward $(\mathbf{s}, \mathbf{a}^\varphi, \mathbf{s}', r)$. By using the *safe action* tuple, we are correctly rewarding the agent for the actual performed transition. However, this requires updating the agent with an action that did not stem from the agent’s current policy $\pi(\mathbf{a}|\mathbf{s})$. This is an expected behavior in off-policy learning, so it is assumed that the safe action tuple is a better fit for off-policy than for on-policy RL.
- *both* - in case the RL agent proposes an unsafe action, both the *adaption penalty* and the *safe action* tuples are used for learning, otherwise *naive* is used.

In all cases, the next state \mathbf{s}' and reward r are the true state and reward received from the environment after executing the safe action \mathbf{a}^φ .

2.2 Action projection

In contrast to action replacement, where the replacement action is not necessarily related to the agent’s action, action projection returns the closest provably safe action with respect to the original action and a distance function. For this, we define the optimization problem

$$\begin{aligned} \arg \min_{\tilde{\mathbf{a}}} \quad & \text{dist}(\mathbf{a}, \tilde{\mathbf{a}}) \\ \text{subject to} \quad & \varphi(\mathbf{s}, \tilde{\mathbf{a}}) = 1, \end{aligned} \tag{3}$$

where $\text{dist}(\cdot)$ describes an arbitrary distance function, e.g., any p -norm. Note, that it might not be possible to define such a distance function, especially in discrete action spaces. The constraints are often defined explicitly by constraint functions $f_i(\tilde{\mathbf{a}}, \mathbf{s}) \leq 0, \forall i \in 1, \dots, n$ that describe that the next state $\mathbf{s}' \in \mathbb{S}_\varphi$. The optimization problem in (3) minimizes the alteration of the actions while satisfying the safety constraints. Following assumption 1, we require that the optimizer always finds a solution for the problem in (3) if any exists.

The most prominent ways to formulate the safety constraints for action projection are based on control barrier functions (CBFs) or robust MPC. For the first method, the constraints are defined by CBFs (Wieland & Allgöwer, 2007) that translate state constraints to control input constraints. We formulate the CBFs according to Taylor et al. (2020) as it is an intuitive formulation for RL. Given is a nonlinear control-affine system

$$\dot{\mathbf{s}} = \mathbf{m}(\mathbf{s}) + \mathbf{b}(\mathbf{s}) \tilde{\mathbf{a}}, \tag{4}$$

where $\mathbf{s} \in \mathbb{S} \subseteq \mathbb{R}^N$ is the continuous state with N dimensions, and $\tilde{\mathbf{a}} \in \mathbb{A} \subset \mathbb{R}^M$ is the continuous control input with M dimensions and $\mathbf{m}(\mathbf{s})$ and $\mathbf{b}(\mathbf{s})$ are locally Lipschitz continuous functions. Then, the function h is a CBF if there exists an extended class \mathbb{K} function α such that (Taylor et al., 2020, Eq. 3)

$$\dot{h}(\mathbf{s}, \tilde{\mathbf{a}}) \geq -\alpha(h(\mathbf{s})). \tag{5}$$

The limitation to control-affine systems makes the formulation of the constrained optimization problem efficient, e.g., for a Euclidean norm as distance function, (3) results in a quadratic programming (QP). A downside of using CBFs is that $h(\mathbf{s})$ is not trivial to find, especially in environments with dynamic obstacles.

For the second common projection method, we formulate the optimization problem with MPC according to Wabersich & Zeilinger (2021). There, input constraints are formulated next to the system dynamics and safety constraints, which are usually given as state constraints. The constraints in (3) can then be formulated to find an input trajectory that steers the system from the current state \mathbf{s} to the safe terminal set $\mathbb{M} \subseteq \mathbb{S}_\varphi$ within the prediction horizon L (Wabersich & Zeilinger, 2021, Eq. 5):

$$\begin{aligned} & \arg \min_{\tilde{\mathbf{a}}} \quad \text{dist}(\mathbf{a}, \tilde{\mathbf{a}}) \\ & \text{subject to} \quad \mathbf{s}_{l+1} = \mathbf{g}(\mathbf{s}_l, \tilde{\mathbf{a}}_l, \mathbf{d}), \mathbf{s}_0 = \mathbf{s}, \\ & \quad \mathbf{s}_l \in \mathbb{S}_\varphi \quad \forall l \in \{1, \dots, L-1\}, \\ & \quad \mathbf{s}_L \in \mathbb{M}, \\ & \quad \mathbf{a}_l \in \mathbb{A} \quad \forall l \in \{1, \dots, L\}, \\ & \quad \tilde{\mathbf{a}} = \mathbf{a}_1, \end{aligned} \tag{6}$$

where \mathbf{s}_l and \mathbf{a}_l are the state and action l -steps ahead of the current time step, and $\mathbf{g}(\cdot)$ describes the system dynamics including disturbances \mathbf{d} . For the set \mathbb{M} , a controller exists that keeps the agent within this set for an infinite time. If the optimization problem is solvable, then $\tilde{\mathbf{a}}$ is executed. If it is not solvable, the control sequence from the previous state is used as a backup plan until the safe terminal set is reached or the optimization problem is solvable again (Schürmann et al., 2018). The MPC formulation allows one to easily integrate model and measurement uncertainties. However, for an environment with dynamic obstacles, the safe terminal set can be time-dependent and we are not aware of a straightforward integration that still guarantees assumption 1.

For both CBF and MPC approaches, nominal models are used to model the known system dynamics. The part of the system dynamics that are unknown can be modelled as disturbances

$$\dot{\mathbf{s}} = \mathbf{m}(\mathbf{s}) + \mathbf{b}(\mathbf{s}) \tilde{\mathbf{a}} + \mathbf{d}, \tag{7}$$

where safety can be ensured if \mathbf{d} is a time signal that is essentially bounded in time $t \geq 0$, $d(t) \leq \bar{d} > 0$ for every entry d of \mathbf{d} and where the bound is finite $\bar{d} < \infty$ (Taylor et al., 2021). To reduce conservatism, disturbances can be modelled as state and input depended $\mathbf{d}(\mathbf{s}, \tilde{\mathbf{a}})$ and the maximal disturbance can be learned from data as presented in Taylor et al. (2021) for CBF, and Hewing et al. (2020) for MPC models.

For provably safe RL using action projection, we can use the same learning tuples as discussed in section 2.1. Additionally, the penalty r^* in the *adaption penalty* tuple can be designed proportional to the projection distance $\text{dist}(\mathbf{a}, \tilde{\mathbf{a}})$, as proposed by Wang (2022).

2.3 Action masking

The two previous approaches modify unsafe actions from the agent. In action masking, we do not allow the agent to output an unsafe action in the first place (preemptive method). Hereby, a mask is added to the agent so that it can only choose from actions of the safe action set. In addition to assumption 1, action masking in practice requires an efficient function $\boldsymbol{\eta}(\mathbf{s})$ that determines a sufficiently large set of provably safe actions \mathbb{A}_φ for a given state \mathbf{s} . The policy function $\boldsymbol{\pi}$ is informed by the function $\boldsymbol{\eta}(\mathbf{s})$ and the action selection is adapted such that only actions from \mathbb{A}_φ can be selected:

$$\mathbf{a} \sim \boldsymbol{\pi}(\mathbf{a} | \boldsymbol{\eta}(\mathbf{s}), \mathbf{s}) \in \mathbb{A}_\varphi. \tag{8}$$

If $\boldsymbol{\eta}(\mathbf{s})$ can only verify one or a few actions efficiently, the agent cannot learn properly because the agent cannot explore different actions and find the optimal one among them. Ideally, the function $\boldsymbol{\eta}(\mathbf{s})$ achieves $\mathbb{A}_\varphi = \mathbb{A}_s$.

The action masking approaches for discrete and continuous action spaces are not easily transferable into each other, and will therefore be discussed separately in this subsection. For discrete actions, the safety of

each action is verified in each state using $\varphi(\mathbf{s}, \mathbf{a})$ and all verified safe actions are added to $\mathbb{A}_\varphi(\mathbf{s})$, i.e., the discrete action mask is an informed drop-out layer added at the end of the policy network. We define the resulting safe policy $\pi_m(\mathbf{a}|\mathbf{s})$ as

$$\pi_m(\mathbf{a}|\mathbf{s}) = \varphi(\mathbf{s}, \mathbf{a}) \frac{\pi(\mathbf{a}|\mathbf{s})}{\sum_{\mathbf{a}' \in \mathbb{A}_\varphi(\mathbf{s})} \pi(\mathbf{a}'|\mathbf{s})}. \quad (9)$$

The integration of masking in a specific learning algorithm is not trivial. The effects on policy optimization methods are discussed in Krasowski et al. (2020); Huang & Onta  n (2022). For RL algorithms that learn the Q-function, we exemplify the effects of discrete action masking for deep Q-network (DQN) (Mnih et al., 2013), which is most commonly used for Q-learning with discrete actions. During exploration with action masking, the agent samples its actions uniformly from \mathbb{A}_φ . When the agent exploits the Q-function, it chooses only the best action among \mathbb{A}_φ , i.e., $\arg \max_{\mathbf{a}} Q(\mathbf{s}, \mathbf{a})$. The temporal difference error for updating the Q-function $Q(\mathbf{s}, \mathbf{a})$ is (Mnih et al., 2013, Eq. 3)

$$r(\mathbf{s}, \mathbf{a}) + \gamma \max_{\mathbf{a}'} Q(\mathbf{s}', \mathbf{a}') - Q(\mathbf{s}, \mathbf{a}), \quad (10)$$

where the action in the next state is $\mathbf{a}' \in \mathbb{A}_\varphi$ in contrast to the vanilla temporal difference error where the maximum Q-value for the next state is searched among actions from \mathbb{A} . Using the adapted temporal difference error in (10), the learning updates are performed only with Q-values of actions relevant in the next state instead of the full action space.

We implement continuous action masking as a transformation of the action of agents to the provable safe action set. The action space \mathbb{A} is transformed into \mathbb{A}_φ by applying the transformation

$$\tilde{\mathbf{a}} = (\mathbf{a} - \min(\mathbb{A})) \frac{\max(\mathbb{A}_\varphi) - \min(\mathbb{A}_\varphi)}{\max(\mathbb{A}) - \min(\mathbb{A})} + \min(\mathbb{A}_\varphi) \quad (11)$$

to the actions \mathbf{a} , where \mathbb{A}_φ is assumed to be convex, $\min(\cdot)$ and $\max(\cdot)$ return a vector containing the minimal and maximal value of the given set in each dimension respectively, and all operations are evaluated element-wise. For example, given a two dimensional continuous action space $\mathbb{A} = [0, 1] \times [-1, 2]$, then $\min(\mathbb{A}) = [0, -1]^\top$. Note that this implementation approximates the safe action set with an interval set which can be conservative.

Since the action spaces for RL are defined a priori, there must always be a valid transformation from \mathbf{a} to $\tilde{\mathbf{a}}$ such that $\tilde{\mathbf{a}} \in \mathbb{A}_\varphi$ and such that the operation is deterministic and time-invariant for a specific state and action pair. In contrast to action replacement, the action is only transformed rather than replaced with a potentially non-deterministic safe action. In this way the agent can explore the full action space \mathbb{A} since (11) ensures that the action is transformed to the safe mapping of it. If one simply clips the action to the safe action space, the policy distribution is altered so that all probabilities for actions outside the clip range are projected to the safe action on the boundary of the safe action space. Thus, the probability distribution differs from the distribution of the current policy, which can be problematic for on-policy RL. Additionally, actions on the bound of the provably safe action space become disproportional more likely. In contrast, the action transformation we propose in (11) conserves the original policy and is conceptually similar to action normalization, which is commonly used in RL (Sutton & Barto, 2018, Ch. 16.8).

Action masking is always paired with the *naive* or *adaption penalty* learning tuple. Using the *safe action* tuple with the transformed action $\tilde{\mathbf{a}}$ leads to inconsistencies in training. For example, the agent chooses in state \mathbf{s}_1 the action \mathbf{a}_1 which is transformed to $\tilde{\mathbf{a}}_1$. If the agent visits state \mathbf{s}_1 again, and chooses $\mathbf{a}_2 = \tilde{\mathbf{a}}_1$, the executed action becomes $\tilde{\mathbf{a}}_2 \neq \mathbf{a}_2$ due to the transformation.

3 Literature Review

In this section, we summarize previous works in provably safe RL and assign them to the proposed categories. To identify the related literature, we used the search string `TITLE-ABS("reinforcement learning") AND TITLE(learning) AND [TITLE(safe*) OR TITLE(verif*) OR TITLE(formal*) OR TITLE(shield*)]`

AND LIMIT-TO(LANGUAGE,"English") for the Scopus² and IEEEExplore³ search engine, which led to 620 papers already removing duplicates. Then, we screened papers by title and abstracts to identify 160 seemingly relevant papers. After close inspection, we identified 47 of these 160 papers as provably safe RLworks. We give a condensed overview of all application-independent provably safe RL works in Table 1, and cluster all 47 provably safe RL works in Table 2 by their application.

Table 1: Comparison of application-independent provably safe RL approaches.

Reference	Verification Method	Space		Learning Tuple	Environment
		State	Action		
Action replacement					
Akametalu et al. (2014)	HJI ¹ reachability analysis	cont.	cont.	special RL alg.	1D quadrotor, cart-pole
Fisac et al. (2019)	HJI reachability analysis	cont.	cont.	N/A	1D quadrotor
Alshiekh et al. (2018)	model checking of automaton constructed from LTL ²	disc.	disc.	<i>safe action</i>	Grid world
Könighofer et al. (2020)	model checking of automaton constructed from LTL	disc.	disc.	<i>adaption penalty</i>	ACC ³
Anderson et al. (2020)	robust control invariant set	cont.	cont.	N/A	pendulum, reach-avoid, others
Hunt et al. (2021)	theorem proofing of dL ⁴ formulas	disc.	disc.	<i>naive</i>	Grid world
Bastani (2021)	MPC ⁵	cont.	cont.	deployment only	bicycle, cart-pole
Shao et al. (2021)	Set-based reachability analysis	cont.	cont.	<i>naive</i>	3D quadrotor, highway driving
Selim et al. (2022b)	Set-based reachability analysis	cont.	cont.	<i>adaption penalty</i>	3D quadrotor, mobile robot
Action projection					
Pham et al. (2018)	verification of affine action constraints	cont.	cont.	<i>adaption penalty</i>	manipulator
Cheng et al. (2019)	CBF ⁶	cont.	cont.	<i>safe action</i>	ACC, pendulum
Li et al. (2019a)	LTL and CBF	cont.	cont.	<i>adaption penalty</i>	manipulator
Gros et al. (2020)	MPC	cont.	cont.	<i>naive</i>	2D LTI ⁷ system
Wabersich & Zeilinger (2021)	MPC	cont.	cont.	<i>naive</i>	3D quadrotor, pendulum
Marvi & Kiumarsi (2022)	CBF	cont.	cont.	<i>adaption penalty</i>	2D LTI system
Selim et al. (2022a)	Set-based reachability analysis	cont.	cont.	<i>naive</i>	3D quadrotor, mobile robot
Kochdumper et al. (2023)	Set-based reachability analysis	cont.	cont.	<i>adaption penalty</i>	3D quadrotor
Action masking					
Fulton & Platzer (2018)	theorem proofing of dL formulas	cont.	disc.	<i>naive</i>	ACC
Fulton & Platzer (2019)	theorem proofing of dL formulas	cont.	disc.	<i>naive</i>	ACC
Huang & Ontañón (2022)	verification of affine action equality constraints	disc.	disc.	<i>naive</i>	Grid world
This study	Set-based reachability analysis	cont.	cont.	<i>naive</i>	pendulum, 2D quadrotor

Abbreviations: ¹Hamilton-Jacobi-Isaacs (HJI), ²linear temporal logic (LTL), ³adaptive cruise control (ACC), ⁴differential dynamic logic ($d\mathcal{L}$), ⁵model predictive control (MPC), ⁶control barrier function (CBF), ⁷linear time-invariant (LTI).

Action replacement One of the earliest provably safe RL works are Alshiekh et al. (2018), which construct a so-called safety shield from linear temporal logic formulas. Hereby, they construct the verification function $\varphi(\mathbf{s}, \mathbf{a})$ by converting the linear temporal logic formulas into an automaton, on which they then perform model checking. They propose to save online computation time by computing $\varphi(\mathbf{s}, \mathbf{a})$ offline prior to training. The agent outputs n ranked actions, which are all checked for safety using $\varphi(\mathbf{s}, \mathbf{a})$. The shield then either executes the highest-ranked safe action or replaces the action with $\psi_{\text{sample}}(\mathbf{s})$ if none of the n actions is safe. They update the agent with the *safe action* learning tuple but also propose that *both* can be used to obtain additional training information. In a similar fashion as Alshiekh et al. (2018), Könighofer et al. (2020) show that both probabilistic and deterministic shields increase the sample efficiency and performance for both action replacement and masking methods.

²scopus.com

³ieeexplore.ieee.org

Table 2: Overview of applications in provably safe RL.

	Action Replacement	Action Projection	Action Masking
Aerial Vehicles	^{††} Akametalu et al. (2014); Shyamsundar et al. (2017); ^{††} Fisac et al. (2019); Anderson et al. (2020); [†] Harris & Schaub (2020); [†] Shao et al. (2021); [†] Selim et al. (2022b); [†] Nazmy et al. (2022)	[†] Wabersich & Zeilinger (2021); [†] Selim et al. (2022a); [†] Kochdumper et al. (2023)	N/A
Autonomous Driving	[†] Chen et al. (2020); Könighofer et al. (2020); [†] Shao et al. (2021); [†] Chen et al. (2022a); [†] Lee & Kwon (2022); ^{††} Wang et al. (2023); [†] Evans et al. (2023)	Cheng et al. (2019); [†] Wang (2022); [†] Hailemichael et al. (2022b); [†] Hailemichael et al. (2022a); ^{††} Kochdumper et al. (2023)	Fulton & Platzer (2018); [†] Mirchevska et al. (2018); Fulton & Platzer (2019); [†] Krasowski et al. (2020); Brosowsky et al. (2021); [†] Krasowski et al. (2022)
Power Systems	[†] Ceusters et al. (2023)	[†] Eichelbeck et al. (2022); [†] Chen et al. (2022b); [†] Zhang et al. (2023); [†] Yu et al. (2023)	[†] Tabas & Zhang (2022)
Robotic Manipulation	[†] Thumm & Althoff (2022)	^{††} Pham et al. (2018); ^{††} Li et al. (2019a)	N/A
Control Benchmarks	Akametalu et al. (2014); Anderson et al. (2020); Bastani (2021); Shao et al. (2021)	Cheng et al. (2019); Gros et al. (2020); Wabersich & Zeilinger (2021); Marvi & Kiumarsi (2022)	N/A
Grid World Games	Alshiekh et al. (2018); Hunt et al. (2021)	N/A	Huang & Ontañón (2022)
Miscellaneous	<i>Active suspension:</i> Li et al. (2019b); <i>Computation:</i> [†] Wang et al. (2022); <i>Mobile robot:</i> [†] Selim et al. (2022b)	<i>Mobile robot:</i> [†] Selim et al. (2022a); <i>Engine emission:</i> [†] Norouzi et al. (2023)	<i>Computation:</i> Seetanadi et al. (2020); <i>Traffic signal:</i> [†] Müller & Sabatelli (2022)

Note: Studies using high-fidelity simulators are marked with [†], and ^{††} indicates physical experiments. Additionally, papers occur multiple times in case they demonstrate their approach for different applications.

The authors in Akametalu et al. (2014) propose action replacement based on Hamilton-Jacobi-Isaacs reachability analysis, which was later extended by Fisac et al. (2019) to a general safe RL framework. They determine \mathbb{S}_φ using Hamilton-Jacobi-Isaacs reachability analysis given bounded system disturbances. The safety is then guaranteed by replacing any learned action on the border of \mathbb{S}_φ with $\psi_{\text{failsafe}}(\mathbf{s})$ stemming from an Hamilton-Jacobi-Isaacs optimal controller to guide the system back inside the safe set. The authors argue that replacing the unsafe action with the action that maximizes the distance to the unsafe set increases performance in uncertain real-world environments in comparison to action projection. However, Hamilton-Jacobi-Isaacs approaches suffer from the curse of dimensionality and are therefore only feasible for systems with specific characteristics (Herbert et al., 2021).

Shao et al. (2021) use a continuous reachability-based trajectory safeguard based on set-based reachability analysis for $\varphi(\mathbf{s}, \mathbf{a})$. If the agent’s action is unsafe, n new actions are sampled randomly in the vicinity of the agent’s action, and the closest safe action to the original (unsafe) action is executed. If none of the n new actions is safe, a failsafe action $\psi_{\text{failsafe}}(\mathbf{s})$ is executed. They train their agent on the *naive* learning tuple. Selim et al. (2022b) also verify the safety of actions with set-based reachability analysis. They propose an informed replacement for $\psi(\mathbf{s})$ such that the reachable set of the controlled system are pushed away from the unsafe set $\mathbb{S} \setminus \mathbb{S}_\varphi$. The authors further propose a method to account for unknown system dynamics using a so-called black-box reachability analysis. They use the *adaption penalty* learning tuple and showcase that their approach achieves provable safety on three use cases.

Hunt et al. (2021) build the verification function $\varphi(\mathbf{s}, \mathbf{a})$ using theorem proofing of differential dynamic logic formulas. Using $\varphi(\mathbf{s}, \mathbf{a})$, they determine $\mathbb{A}_\varphi(\mathbf{s})$ for discrete action spaces, and use $\psi_{\text{failsafe}}(\mathbf{s})$ for replacement. They further show how provably safe end-to-end learning can be accomplished using controller and model

monitors. They train the RL agent using the *naive* learning tuple on a drone example. The work of Anderson et al. (2020) proposes to define \mathbb{S}_φ as a robust control invariant set and construct $\varphi(\mathbf{s}, \mathbf{a})$ from the single-step discrete system dynamics. A further notable work is Bastani (2021) that proposes a model predictive shield alongside the trained policy, which uses $\psi_{\text{failsafe}}(\mathbf{s})$ for the replacement. Finally, Saunders et al. (2018) propose to use a human in the loop for the $\varphi(\mathbf{s}, \mathbf{a})$ and $\psi_{\text{failsafe}}(\mathbf{s})$ functions in slow-paced environments.

The most popular method by publications is action replacement. This is also visible from the large variety of application-specific approaches, e.g., for aerial vehicles (Shyamsundar et al., 2017; Harris & Schaub, 2020; Nazmy et al., 2022), autonomous driving (Chen et al., 2020; 2022a; Lee & Kwon, 2022; Wang et al., 2023; Evans et al., 2023), power systems (Ceusters et al., 2023), robotic manipulation (Thumm & Althoff, 2022), active suspension system (Li et al., 2019b), and data traffic engineering for computation (Wang et al., 2022).

Action projection Research on action projection is usually conducted on continuous action and state spaces. The main differentiating factor between studies in this category is the specification of the optimization problem for the projection. To begin with, the work of Pham et al. (2018) guarantees safety using a differentiable constrained optimization layer, called OptLayer. Unfortunately, their approach is restricted to QP problems, so the system model and constraints have to be linear. Despite these limitations, they show the effectiveness of their approach on a collision avoidance task with a simple robotic manipulator.

Cheng et al. (2019) specify the safety constraints $\varphi(\mathbf{s}, \mathbf{a}) = 1$ of the optimization problem via CBFs. To solve the optimization problem more efficiently, they add a neural network to the approach in section 2.2, which approximates the correction of the CBF controllers on this action. The action is then shifted by the approximated value prior to optimization. This improves the implementation efficiency while still guaranteeing safety, as the action is often already safe after the shift and no optimization problem needs to be solved. Their safe learning with CBFs shows faster convergence speed in comparison to vanilla RL when learning on a pendulum and a car following task. Li et al. (2019a) propose a method to construct a continuous CBF from an automaton, which is defined by linear temporal logic formulas. They further construct a guiding reward from the given automaton to improve the learning performance. The proposed approach is capable of learning a high-dimensional cooperative manipulation task in a safe manner. The authors in Marvi & Kiumarsi (2022) define a different problem, where the system model is assumed to be deterministic but unknown. They learn an optimal controller and the system dynamics in an iterative fashion, while decreasing the conservativeness of their CBF in each iteration. The approach is provably safe for linear time-invariant (LTI) systems without disturbances.

Gros et al. (2020) and Wabersich & Zeilinger (2021) implement the optimization problem as a robust MPC problem as defined in (6). Thus, safety is specified through safe sets in the state space. Gros et al. (2020) mainly discuss how the learning update has to be adapted if action projection is used for different RL algorithms. For Q-learning, they find that no adaption of the learning algorithm is necessary when the *naive* tuple is used. For policy gradient methods, they argue that the projection must also be included in the gradient for stable learning (Gros et al., 2020, Sec. 3). One downside of the robust MPC formulation of Gros et al. (2020) and Wabersich & Zeilinger (2021) is that dynamic constraints originating from moving obstacles or persons in the environment are not trivial to integrate. They approximate the dynamics of the system with a Gaussian process (GP), so that absolute safety is not given. However, they could guarantee absolute safety if they would assume bounded dynamics of the environment as the aforementioned approaches do. Gros et al. (2020) evaluate their approach on a simple 2D LTI system and Wabersich & Zeilinger (2021) show the efficacy of their approach on a pendulum and a quadrotor task.

Contrary to their previous work in Selim et al. (2022b), Selim et al. (2022a) propose to solve an optimization problem to find the closest safe action instead of using an informed replacement. They again use set-based reachability analysis to construct $\varphi(\mathbf{s}, \mathbf{a})$. They test their approach on a quadrotor and mobile robot benchmark. Kochdumper et al. (2023) utilize set-based reachability analysis to verify actions in $\varphi(\mathbf{s}, \mathbf{a})$. They formulate the projection for the generator’s parameterization of the action space zonotope and arrive at a mixed-integer quadratic problem with polynomial constraints. Their approach achieves provable safety for non-linear systems with bounded disturbances and they demonstrate their approach on two quadrotor tasks, autonomous driving on highways and on a physical fltenth car.

Next to the conceptual approaches, action projection algorithms are also specifically proposed for many cyber-physical systems such as autonomous driving (Wang, 2022; Hailemichael et al., 2022a;b), power systems (Eichelbeck et al., 2022; Chen et al., 2022b; Zhang et al., 2023; Yu et al., 2023), and engine emission control (Norouzi et al., 2023).

Action masking To the best of our knowledge, all existing literature considers action masking only for discrete action spaces. The work Huang & Onta  n (2022) analyzes the effect of discrete action masking on the policy gradient algorithm in RL, but they assume that \mathbb{A}_s is known, which is only the case in game and grid world environments. The two main works investigating action masking are Fulton & Platzer (2018) and Fulton & Platzer (2019). They construct controller and model monitors based on theorem proving of differential dynamic logic specifications of hybrid systems. The controller monitor is used to build the mask $\eta(s)$. The model monitor verifies if the underlying system model is correct based on previous transitions. In each state, the agent can choose from the set of actions that the controller monitor verified as safe. To identify the correct system can be challenging, thus an approach to automatically generate candidates is introduced as well. Their approach is provably safe if the initial model is correct (Fulton & Platzer, 2018) or multiple models are given, from which at least one is correct (Fulton & Platzer, 2019) for all times. They validate their provably safe action masking on adaptive cruise control tasks. Additionally, we want to highlight that our (this) work is the first to introduce action masking for continuous action spaces as described in section 2.3, and is therefore included in Table 1.

In addition to the aforementioned works, there are works investigating action masking for the specific application of autonomous driving (Mirchevska et al., 2018; Krasowski et al., 2020; Brosowsky et al., 2021; Krasowski et al., 2022), power systems (Tabas & Zhang, 2022), adaptive routing for computation (Seetanadi et al., 2020), and urban traffic signal control (M  ller & Sabatelli, 2022).

4 Experimental Comparison

We compare four provably safe RL implementations. Action masking, action projection with CBFs, and action replacement with $\psi_{\text{sample}}(s)$ and $\psi_{\text{failsafe}}(s)$. We run all experiments for discrete and continuous action spaces. For action replacement and action projection, we compare the four possible learning tuples introduced in section 2.1. When action replacement is used with a failsafe controller, we omit *safe action* and *both* because, especially in the discrete action space, the failsafe controller might use an action that is not in the action space. When action masking is used, only safe actions can be sampled, i.e., only the *na  ve* tuple is meaningful. For on-policy learning, proximal policy optimization (PPO) (Schulman et al., 2017) is used. For off-policy learning, we use the Twin Delayed Deep Deterministic policy gradient algorithm (TD3) (Fujimoto et al., 2018) for the continuous action space and DQN (Mnih et al., 2013) for the discrete. The learning algorithms are adapted from Stable Baselines3 (Raffin et al., 2021). Every configuration is evaluated on five random seeds.

4.1 Environments

We compare the provably safe RL approaches on an inverted pendulum and a 2D quadrotor stabilization task.

Inverted Pendulum The state of the pendulum is defined as $s = [\theta, \dot{\theta}]$, and follows the dynamics

$$\dot{s} = \begin{pmatrix} \dot{\theta} \\ \frac{g}{l} \sin(\theta) + \frac{1}{ml^2} a \end{pmatrix}, \quad (12)$$

where a is the one-dimensional action, g is gravity, m is the mass of the pendulum, l its length, and friction and damping are ignored. We discretize the dynamics using the Euler method. The input lies within $[a_{\min}, a_{\max}] = [-30\text{rad s}^{-1}, 30\text{rad s}^{-1}]$. The desired equilibrium state is $s^* = [0, 0]$. The observation and reward are the same as for the *OpenAI Gym Pendulum-V0*⁴ environment.

⁴gymnasium.farama.org/environments/classic_control/pendulum/

2D quadrotor The quadrotor can only fly in the vertical x - z -plane and rotate around the y -axis with angle θ . The state of the system is defined as $\mathbf{s} = [x, z, \dot{x}, \dot{z}, \theta, \dot{\theta}]$ and the action as $\mathbf{a} = [u_1, u_2]$. The system dynamics are

$$\dot{\mathbf{s}} = \begin{pmatrix} \dot{x} \\ \dot{z} \\ u_1 k \sin(\theta) + w_1 \\ -g + u_1 k \cos(\theta) + w_2 \\ \dot{\theta} \\ -d_0 \theta - d_1 \dot{\theta} + n_0 u_2 \end{pmatrix} \quad (13)$$

where the noise is sampled uniformly from a compact noise set $\mathbf{w} = [w_1, w_2] \sim \mathbb{W} \subset \mathbb{R}^2$, and k , d_0 , d_1 , and n_0 are constant system parameters (see Table 4). We linearize the dynamics with a first-order Taylor expansion at the equilibrium point $\mathbf{s}^* = [0, 1, 0, 0, 0, 0]$ and solve the differential equations for a time-discrete system. The input ranges from $\mathbf{a}_{\min} = [-1.5 + \frac{g}{K}, -\frac{\pi}{12}]$ to $\mathbf{a}_{\max} = [1.5 + \frac{g}{K}, \frac{\pi}{12}]$. The reward is defined as $r = \exp\left(-\|\mathbf{s} - \mathbf{s}^*\| - 0.01 \sum_{i=1}^2 \frac{a_t[i] - a_{\min}[i]}{a_{\max}[i] - a_{\min}[i]}\right)$.

4.2 Computation of the Safe State Set

To obtain a possibly large set of provably safe states \mathbb{S}_φ and a failsafe controller for our environments, we use the scalable approach for computing robust control invariant sets presented in Schäfer et al. (2023): for every state $\mathbf{s}_t \in \mathbb{S}_\varphi$, there exists a failsafe action $\tilde{\mathbf{a}}_t \in \mathbb{A}$ so that $\mathbf{s}_{t+1} = \mathbf{g}(\mathbf{s}_t, \tilde{\mathbf{a}}_t) \in \mathbb{S}_\varphi$. Since $\mathbb{S}_\varphi \subseteq \mathbb{S}$, satisfaction of our safety specifications $\mathbf{s} \in \mathbb{S}$ and $\tilde{\mathbf{a}} \in \mathbb{A}$ for every future point in time follows by induction. Using the algorithm in Schäfer et al. (2023) provides us with an explicit representation of \mathbb{S}_φ with maximized volume, which enables a fair comparison of our provably safe RL implementations.

To retrieve \mathbb{A}_φ from \mathbb{S}_φ at a given state \mathbf{s} , we first convert \mathbb{S}_φ from a so-called generator representation used in Schäfer et al. (2023) into half-space representation, i.e., $\mathbb{S}_\varphi = \{\mathbf{s} | \mathbf{C}\mathbf{s} \leq \mathbf{d}\}$ using the open-source toolbox CORA (Althoff, 2015). We then formulate the dynamics of the system in the form of (4) with an additional noise term $\mathbf{g}(\mathbf{s}, \mathbf{a}) = \mathbf{g}_a(\mathbf{s})\mathbf{a} + \mathbf{g}_u(\mathbf{s}) + \mathbf{w}$. The noise \mathbf{w} is assumed to stem from a compact set $\mathbf{w} \sim \mathbb{W}$. Finally, we reformulate the constraints as

$$\mathbf{C}\mathbf{g}_a(\mathbf{s})\mathbf{a} \leq \mathbf{d} - \mathbf{C}\mathbf{g}_u(\mathbf{s}) - \mathbf{C}\mathbf{w}. \quad (14)$$

This results in a polytope with $\dim(\mathbf{d}) \times 2^{\dim(\mathbb{W})}$ halfspaces. Finally, we use `pypoman`⁵ to obtain the halfspace-representation of \mathbb{A}_φ .

In theory, the approach of Schäfer et al. (2023) allows us to calculate \mathbb{S}_φ for nonlinear systems with up to 20 dimensions, such as chemical reactors. However, the conversion to halfspace representation is computational too complex for higher dimensional systems. Therefore, we plan to develop generator-based versions of our provably safe RL methods in future work to mitigate these shortcomings.

4.3 Results

The differences between the provably safe RL algorithms during training is shown in figure 2 for discrete PPO on the 2D quadrotor task. We selected the 2D quadrotor stabilization task as it allows for a better differentiation between provably safe RL approaches than the inverted pendulum task. The difference between the three provably safe RL categories is clearest for the discrete PPO results, but we report similar trends on all other algorithms as well. The complete set of results for all algorithms and environments can be found in the Appendix.

The safety violations for the baselines in figure 2b show that safety can only be guaranteed when using provably safe RL. Between the baselines, TD3 converges significantly faster than all other algorithms.

The comparison between the three provably safe RL categories action replacement (sample), projection, and masking in figure 2c shows that replacement is the most stable and high-performing method. Action

⁵Available at <https://github.com/stephane-caron/pypoman>

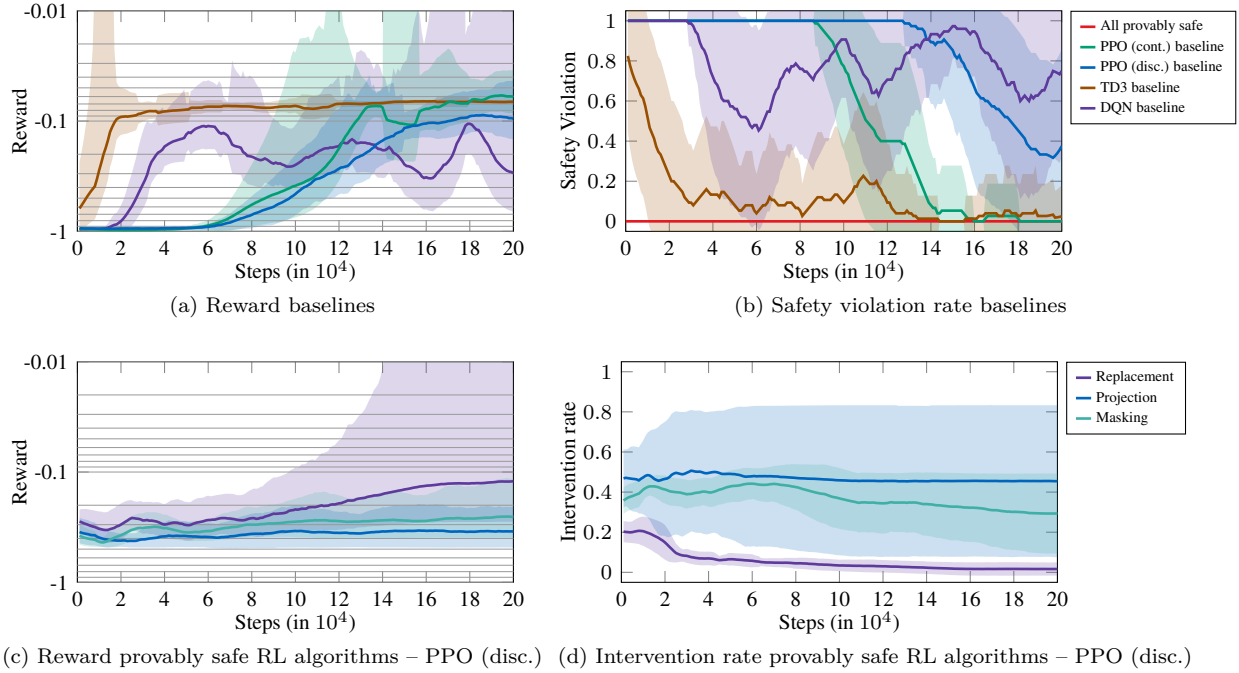


Figure 2: Training curves for the 2D quadrotor benchmark. Top: Comparison of benchmark algorithms TD3, DQN, PPO continuous and discrete. Bottom: Comparison of the provably safe RL algorithms using discrete PPO with the *naïve* tuple. The left column depicts the reward. The right column shows the safety violations for the baselines, and the safety intervention rate for the provably safe RL algorithms.

projection and masking perform similarly. To compare how active the provably safe RL methods are, we introduce the intervention rate metric. For action replacement and projection, this indicates the average rate of RL steps in which the action was altered by the safety function. For action masking, the intervention rate compares the average volume of the provably safe action set over an episode with the volume of the provably safe action set at the equilibrium point of the system, e.g., $V_{\mathbb{A}_\varphi, \text{episode}}/V_{\mathbb{A}_\varphi, \text{equi}}$. Figure 2d shows that action replacement converges to almost zero safety interventions, whereas projection and masking always depend on the safety method for the considered use-case. In general, we report that a lower intervention rate often coincides with a higher reward.

We evaluate the impact of different learning tuples on the overall performance in figure 3. A notable result of our evaluation is that using the adapted action \mathbf{a}^φ , i.e., in the *safe action* and *both* tuple, is inferior to using the action of the agent when using PPO. We expect that this also transfers to other on-policy algorithms as they assume that the action used for training is sampled from the current policy. Using an action altered by the safety method breaks this paradigm. For the off-policy algorithms TD3 and DQN, shown in the Appendix, the *safe action* and *both* tuple did not improve the performance of the *naïve* tuple. The addition of an *adaptation penalty* reduced the intervention rate and improved the overall performance of the replacement and projection methods. Especially in the case of action replacement with a failsafe action, the *adaptation penalty* improved the performance significantly.

5 Discussion

The theory and experiments confirm that provably safe RL methods are always safe, while the baselines still violate the safety results even after the reward converged. Subsequently, we discuss four statements, which result from the experiments.

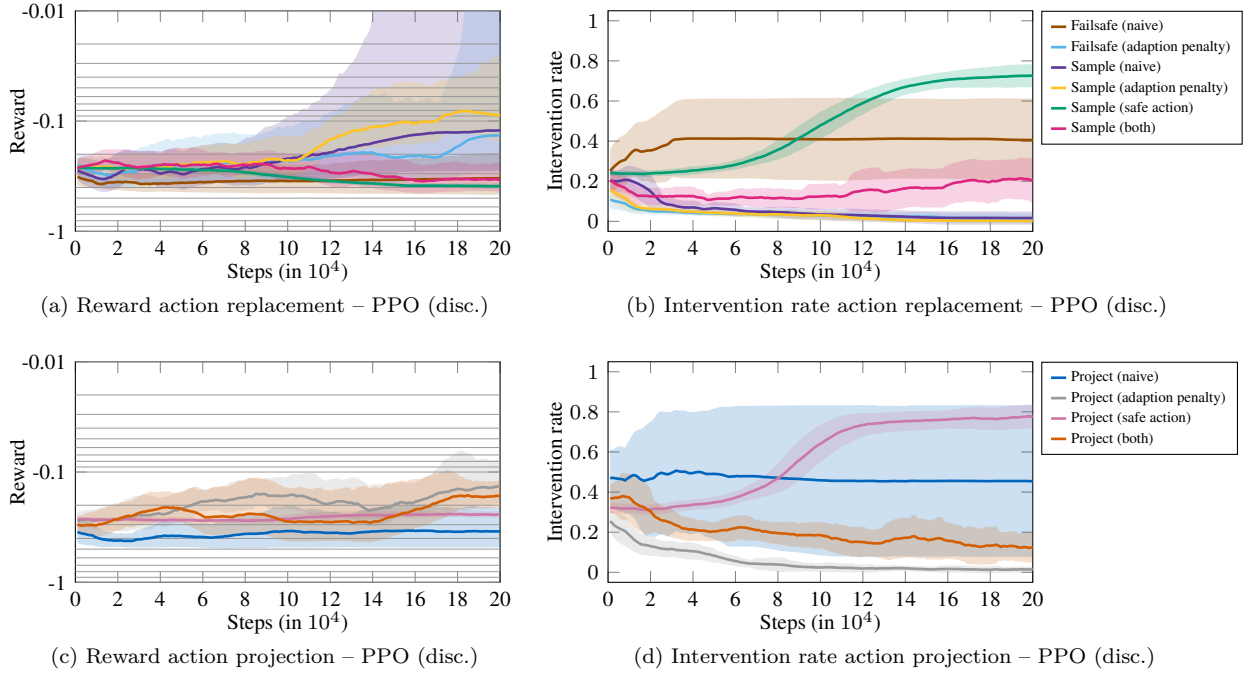


Figure 3: Evaluation of the training tuples for the 2D quadrotor using continuous PPO. The left column depicts the reward and the right column the safety intervention rate. The top row shows the different learning tuples for action replacement and the bottom row for action projection.

First, we want to summarize our experience of the three provably safe RL classes to give the reader an intuition when to choose which method. Replacement is the easiest method to implement for continuous action spaces. It shows very good performance and low intervention rates. Hereby, using a failsafe action for replacement with an adaption penalty is simple to implement and usually among the best performing methods in our experiments and is thereby recommended if available. Still, the random sampling safe actions might outperform the failsafe action. So if sampling from the safe action space is easily available, e.g., in discrete action spaces, a failsafe controller can be omitted. Projection tends to be problematic due to floating point errors and other small numerical errors, resulting in infeasible optimization problems. We, therefore, have to use methods like the provably safe MPC in Schürmann et al. (2018), or use a failsafe controller if the optimization problem is not solvable. Together with the higher intervention rates than action replacement and the complex implementation, we would usually not recommend to use action projection. However, if one already has a CBF or MPC formulation, it might be the most suitable solution. Action masking is particularly easy to implement for discrete action spaces and shows good performance for that setting. However, action masking can be difficult to implement or very restrictive for continuous action spaces, where the safe action space significantly diverges from an axis-aligned box.

Second, the *naive* learning tuple usually is among the best performing tuples while no adaption of the RL algorithms is necessary. If the intervention rate should be minimized and sometimes also to improve performance, the *adaption penalty* tuple can be used. This adaption penalty r^* is easy to integrate in the reward function, but it might require careful reward tuning for more complex environments than ours. For the PPO case, using the safe action \tilde{a} in the training tuple (i.e., configurations *both* and *safe action*) should be avoided since we observed that they are much less stable across different hyperparameters while at the same time lead to worse performance. We expect that this statement is also valid for other on-policy RL algorithms.

Third, provably safe RL converges similarly fast or faster than the baselines, while the performance at the beginning of the training is better for provably safe RL agents. There are two factors influencing the convergence. On the one hand, the baseline agents need to explore more actions including actions that

do not lead to reaching the goal. On the other hand, safe agents do not get as much information about the environment as only safe states are explored. It should be mentioned that the safety methods create a computational overhead such that convergence speed is not necessarily aligned with the elapsed real time.

Finally, the computational complexity of the three approaches highly depends on the scenario-specific implementation. For action projection, the main implementation challenge is to guarantee that the optimization problem is always feasible. If the optimization problem can be formulated as a QP problem, then the computational complexity is polynomial as shown by Vavasis (2001). Contrary, the computational complexity of action replacement and action masking highly depends on the algorithm that identifies the safety of actions. For discrete action masking, the computational complexity scales linearly with the total number of actions $\mathcal{O}(|\mathbb{A}|)$. While for action replacement, we only need a single safe action, so in the ideal case, e.g., using a failsafe controller, the computational complexity is constant in regard to the total number of actions $\mathcal{O}(1)$. If the action replacement approach needs to determine the entire set of safe actions, they have the same computational complexity as action masking. The computational complexity for identifying the continuous safe action space depends on the task-specific implementation. One possibility is to compute the safe action space using set-based reachability analysis, where we want to point the interested reader to Althoff (2015) for different implementations.

6 Conclusion

In conclusion, we define a categorization for provably safe RL methods that structures the literature from a learning perspective. We present these provably safe RL methods from a conceptual perspective and define necessary assumptions. Our proposed categorization into action replacement, action projection, and action masking supports researchers in comparing their works and provide valuable insights into the selection process of provably safe RL methods. The comparison of the four implementations of provably safe RL on a 2D quadrotor and an inverted pendulum stabilization benchmark provides further insights on the best-suited method for different tasks.

We envision several promising avenues for further research that can help to bring provably safe RL research to broader adoption in real-world applications. First, the computation time for some approaches is too high to be real-world applicable. Thus, more computationally efficient algorithms need to be developed. Second, the variety of possible safety specifications for provably safe RL should be increased. Currently, mainly stabilization and reach-avoid specifications are prominent but real-world safety is more complex, e.g., traffic rules such as waiting at a red light and safely but quickly moving at a green light. Third, the expert knowledge necessary to create provably safe RL should be reduced through more modular and automatic approaches where fewer engineering decisions are necessary and more parameters can be tuned automatically. Finally, provably safe RL approaches should be intertwined with safe RL approaches with probabilistic guarantees since, in the real world, it is likely that many system parameters and disturbances are mainly bounded but rarely exceed these bounds. In this case, the safest possible action should be taken. With these advances provably safe RL could bring the best of RL and formal specifications together towards RL methods that require as little expert knowledge as necessary and provide formal guarantees for complex safety specifications in order to achieve reliable and trustworthy cyber-physical systems.

Acknowledgments

Anonymous for double-blind peer review.

References

- Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained Policy Optimization. In *Proc. of the Int. Conf. on Machine Learning (ICML)*, pp. 22–31, 2017.
- Anayo K. Akametalu, Shahab Kaynama, Jaime F. Fisac, Melanie N. Zeilinger, Jeremy H. Gillula, and Claire J. Tomlin. Reachability-based safe learning with Gaussian processes. In *Proc. of the IEEE Conf. on Decision and Control (CDC)*, pp. 1424–1431, 2014.

- Mohammed Alshiekh, Roderick Bloem, Rüdiger Ehlers, Bettina Könighofer, Scott Niekum, and Ufuk Topcu. Safe reinforcement learning via shielding. In *Proc. of the AAAI Conf. on Artificial Intelligence*, pp. 2669–2678, 2018.
- Matthias Althoff. An Introduction to CORA 2015. In *Proc. of the Workshop on Applied Verification for Continuous and Hybrid Systems*, pp. 120–151, 2015.
- Eitan Altman. Constrained Markov decision processes with total cost criteria: Lagrangian approach and dual linear program. *Mathematical Methods of Operations Research*, 48(3):387–417, 1998.
- Greg Anderson, Abhinav Verma, Isil Dillig, and Swarat Chaudhuri. Neurosymbolic Reinforcement Learning with Formally Verified Exploration. In *Advances in Neural Information Processing Systems*, volume 33, pp. 6172–6183, 2020.
- Osbert Bastani. Safe Reinforcement Learning with Nonlinear Dynamics via Model Predictive Shielding. In *Proc. of the American Control Conf. (ACC)*, pp. 3488–3494, 2021.
- Osbert Bastani, Yewen Pu, and Armando Solar-Lezama. Verifiable Reinforcement Learning via Policy Extraction. In *Proc. of the Int. Conf. on Neural Information Processing Systems (NIPS)*, pp. 2499–2509, 2018.
- Felix Berkenkamp, Angela P Schoellig, Matteo Turchetta, and Andreas Krause. Safe Model-based Reinforcement Learning with Stability Guarantees. In *Proc. of the Int. Conf. on Neural Information Processing Systems (NIPS)*, pp. 908–918, 2017.
- Sara Bouraine, Thierry Fraichard, and Hassen Salhi. Provably safe navigation for mobile robots with limited field-of-views in unknown dynamic environments. In *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, pp. 174–179, 2012.
- Mathis Brosowsky, Florian Keck, Jakob Ketterer, Simon Isele, Daniel Slieter, and Marius Zöllner. Safe Deep Reinforcement Learning for Adaptive Cruise Control by Imposing State-Specific Safe Sets. In *Proc. of the IEEE Intelligent Vehicles Symposium (IV)*, pp. 488–495, 2021.
- Lukas Brunke, Melissa Greeff, Adam W. Hall, Zhaocong Yuan, Siqi Zhou, Jacopo Panerati, and Angela P. Schoellig. Safe Learning in Robotics: From Learning-Based Control to Safe Reinforcement Learning. *Annual Review of Control, Robotics, and Autonomous Systems*, 5:411–444, 2022.
- Glenn Ceusters, Luis Ramirez Camargo, Rüdiger Franke, Ann Nowé, and Maarten Messagie. Safe reinforcement learning for multi-energy management systems with known constraint functions. *Energy and AI*, 12, 2023.
- Dong Chen, Longsheng Jiang, Yue Wang, and Zhaojian Li. Autonomous Driving using Safe Reinforcement Learning by Incorporating a Regret-based Human Lane-Changing Decision Model. In *Proc. of the American Control Conf. (ACC)*, pp. 4355–4361, 2020.
- Shengduo Chen, Yaowei Sun, Dachuan Li, Qiang Wang, Qi Hao, and Joseph Sifakis. Runtime Safety Assurance for Learning-enabled Control of Autonomous Driving Vehicles. In *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, pp. 8978–8984, 2022a.
- Yize Chen, Yuanyuan Shi, Daniel Arnold, and Sean Peisert. SAVER: Safe Learning-Based Controller for Real-Time Voltage Regulation. In *Proc. of the IEEE Power and Energy Society General Meeting (PESGM)*, pp. 1–5, 2022b.
- 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. In *Proc. of the AAAI Conf. on Artificial Intelligence*, pp. 3387–3395, 2019.
- Yinlam Chow, Ofir Nachum, Edgar Duenez-Guzman, and Mohammad Ghavamzadeh. A Lyapunov-Based Approach to Safe Reinforcement Learning. In *Proc. of the Int. Conf. on Neural Information Processing Systems (NIPS)*, pp. 8103–8112, 2018.

- Gal Dalal, Krishnamurthy Dvijotham, Matej Vecerik, Todd Hester, Cosmin Paduraru, and Yuval Tassa. Safe exploration in continuous action spaces. *arXiv*, abs/1801.0, 2018.
- Giuseppe De Giacomo, Luca Iocchi, Marco Favorito, and Fabio Patrizi. Foundations for Restraining Bolts: Reinforcement Learning with LTLf/LDLf Restraining Specifications. In *Proc. of the Int. Conf. on Automated Planning and Scheduling (ICAPS)*, pp. 128–136, 2021.
- Michael Eichelbeck, Hannah Markgraf, and Matthias Althoff. Contingency-constrained economic dispatch with safe reinforcement learning. In *Proc. of the IEEE Int. Conf. on Machine Learning and Applications (ICMLA)*, pp. 597–602, 2022.
- Mohamed El-Shamouty, Xinyang Wu, Shanqi Yang, Marcel Albus, and Marco F. Huber. Towards Safe Human-Robot Collaboration Using Deep Reinforcement Learning. In *Proc. of the Int. Conf. on Robotics and Automation (ICRA)*, pp. 4899–4905, 2020.
- Benjamin D Evans, Hendrik W Jordaan, and Herman A Engelbrecht. Safe reinforcement learning for high-speed autonomous racing. *Cognitive Robotics*, 3:107 – 126, 2023.
- Jaime F. Fisac, Anayo K. Akametalu, Melanie N. Zeilinger, Shahab Kaynama, Jeremy Gillula, and Claire J. Tomlin. A General Safety Framework for Learning-Based Control in Uncertain Robotic Systems. *IEEE Transactions on Automatic Control*, 64(7):2737–2752, 2019.
- Scott Fujimoto, Herke Van Hoof, and David Meger. Addressing Function Approximation Error in Actor-Critic Methods. In *Proc. of the Int. Conf. on Machine Learning (ICML)*, pp. 2587–2601, 2018.
- Nathan Fulton and André Platzer. Safe Reinforcement Learning via Formal Methods: Toward Safe Control Through Proof and Learning. In *Proc. of the AAAI Conf. on Artificial Intelligence (AAAI)*, pp. 6485–6492, 2018.
- Nathan Fulton and André Platzer. Verifiably Safe Off-Model Reinforcement Learning. In *Proc. of the Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, pp. 413–430, 2019.
- Javier García and Fernando Fernández. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16(42):1437–1480, 2015.
- Jeremy H. Gillula and Claire J. Tomlin. Reducing Conservativeness in Safety Guarantees by Learning Disturbances Online: Iterated Guaranteed Safe Online Learning. *Robotics: Science and Systems*, 8(1): 81–88, 2013.
- Sebastien Gros, Mario Zanon, and Alberto Bemporad. Safe Reinforcement Learning via Projection on a Safe Set: How to Achieve Optimality? *IFAC-PapersOnLine*, 53(2):8076–8081, 2020.
- Habtamu Hailemichael, Beshah Ayalew, Lindsey Kerbel, Andrej Ivanco, and Keith Loiselle. Safety Filtering for Reinforcement Learning-based Adaptive Cruise Control. *IFAC-PapersOnLine*, 55(24):149–154, 2022a.
- Habtamu Hailemichael, Beshah Ayalew, Lindsey Kerbel, Andrej Ivanco, and Keith Loiselle. Safe Reinforcement Learning for an Energy-Efficient Driver Assistance System. In *IFAC-PapersOnLine*, volume 55:37, pp. 615 – 620, 2022b.
- Andrew Harris and Hanspeter Schaub. Spacecraft command and control with safety guarantees using shielded deep reinforcement learning. In *AIAA Scitech 2020 Forum*, volume 1, 2020.
- Mohammadhosein Hasanbeig, Yiannis Kantaros, Alessandro Abate, Daniel Kroening, George J. Pappas, and Insup Lee. Reinforcement Learning for Temporal Logic Control Synthesis with Probabilistic Satisfaction Guarantees. In *Proc. of the IEEE Conf. on Decision and Control (CDC)*, pp. 5338–5343, 2019a.
- Mohammadhosein Hasanbeig, Daniel Kroening, and Alessandro Abate. Towards verifiable and safe model-free reinforcement learning. In *Proc. of the Workshop on Artificial Intelligence and Formal Verification, Logic, Automata, and Synthesis*, pp. 1–9, 2019b.

- Mohammadhosein Hasanbeig, Alessandro Abate, and Daniel Kroening. Cautious reinforcement learning with logical constraints. In *Proc. of the Int. Conf. on Autonomous Agents and MultiAgent Systems (AAMAS)*, pp. 483–491, 2020.
- Matthias Heger. Consideration of Risk in Reinforcement Learning. In *Proc. of the Int. Conf. on Machine Learning (ICML)*, pp. 105–111, 1994.
- Sylvia Herbert, Jason J. Choi, Suvansh Sanjeev, Marsalis Gibson, Koushil Sreenath, and Claire J. Tomlin. Scalable Learning of Safety Guarantees for Autonomous Systems using Hamilton-Jacobi Reachability. In *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, pp. 5914–5920, 2021.
- Lukas Hewing, Kim P. Wabersich, Marcel Menner, and Melanie N. Zeilinger. Learning-Based Model Predictive Control: Toward Safe Learning in Control. *Annual Review of Control, Robotics, and Autonomous Systems*, 3(1):269–296, 2020.
- Shengyi Huang and Santiago Ontañón. A Closer Look at Invalid Action Masking in Policy Gradient Algorithms. *The Int. Florida Artificial Intelligence Research Society Conf. Proc. (FLAIRS)*, 35, 2022.
- Nathan Hunt, Nathan Fulton, Sara Magliacane, Trong Nghia Hoang, Subhro Das, and Armando Solar-Lezama. Verifiably safe exploration for end-to-end reinforcement learning. In *Proc. of the Int. Conf. on Hybrid Systems: Computation and Control (HSCC)*, pp. 1–11, 2021.
- B. Ravi Kiran, Ibrahim Sobh, Victor Talpaert, Patrick Mannion, Ahmad A. Al Sallab, Senthil Yogamani, and Patrick Perez. Deep Reinforcement Learning for Autonomous Driving: A Survey. *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–18, 2021.
- Niklas Kochdumper, Hanna Krasowski, Xiao Wang, Stanley Bak, and Matthias Althoff. Provably Safe Reinforcement Learning via Action Projection Using Reachability Analysis and Polynomial Zonotopes. *IEEE Open Journal of Control Systems*, 2:79–92, 2023.
- Bettina Könighofer, Florian Lorber, Nils Jansen, and Roderick Bloem. Shield Synthesis for Reinforcement Learning. In *Leveraging Applications of Formal Methods, Verification and Validation: Verification Principles*, pp. 290–306, 2020.
- Bettina Könighofer, Julian Rudolf, Alexander Palmisano, Martin Tappler, and Roderick Bloem. Online Shielding for Stochastic Systems. In *Proc. of the NASA Formal Methods (NFM)*, pp. 231–248, 2021.
- Hanna Krasowski, Xiao Wang, and Matthias Althoff. Safe Reinforcement Learning for Autonomous Lane Changing Using Set-Based Prediction. In *Proc. of the IEEE Int. Conf. on Intelligent Transportation Systems (ITSC)*, pp. 1–7, 2020.
- Hanna Krasowski, Yinqiang Zhang, and Matthias Althoff. Safe Reinforcement Learning for Urban Driving using Invariably Safe Braking Sets. In *Proc. of the IEEE Int. Conf. on Intelligent Transportation Systems (ITSC)*, pp. 2407–2414, 2022.
- Dongsu Lee and Minhae Kwon. ADAS-RL: Safety learning approach for stable autonomous driving. *ICT Express*, 8(3):479 – 483, 2022.
- Xiao Li, Zachary Serlin, Guang Yang, and Calin Belta. A formal methods approach to interpretable reinforcement learning for robotic planning. *Science Robotics*, 4(37), 2019a.
- Zhaojian Li, Tianshu Chu, and Uroš Kalabić. Dynamics-Enabled Safe Deep Reinforcement Learning: Case Study on Active Suspension Control. In *Proc. of the IEEE Conf. on Control Technology and Applications (CCTA)*, pp. 585–591, 2019b.
- Tommaso Mannucci, Erik-Jan van Kampen, Cornelis de Visser, and Qiping Chu. Safe Exploration Algorithms for Reinforcement Learning Controllers. *IEEE Transactions on Neural Networks and Learning Systems*, 29(4):1069–1081, 2018.

- Zahra Marvi and Bahare Kiumarsi. Safe reinforcement learning: A control barrier function optimization approach. *Int. Journal of Robust and Nonlinear Control*, 31(6):1923–1940, 2021.
- Zahra Marvi and Bahare Kiumarsi. Reinforcement Learning With Safety and Stability Guarantees During Exploration For Linear Systems. *IEEE Open Journal of Control Systems*, 1:322–334, 2022.
- Branka Mirchevska, Christian Pek, Moritz Werling, Matthias Althoff, and Joschka Boedecker. High-level Decision Making for Safe and Reasonable Autonomous Lane Changing using Reinforcement Learning. In *Proc. of the Int. Conf. on Intelligent Transportation Systems (ITSC)*, pp. 2156–2162, 2018.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing Atari with Deep Reinforcement Learning. *arXiv*, abs/1312.5, 2013.
- Arthur Müller and Matthia Sabatelli. Safe and Psychologically Pleasant Traffic Signal Control with Reinforcement Learning using Action Masking. *Proc. of the IEEE Conf. on Intelligent Transportation Systems (ITSC)*, pp. 951–958, 2022.
- Islam Nazmy, Andrew Harris, Morteza Lahijanian, and Hanspeter Schaub. Shielded Deep Reinforcement Learning for Multi-Sensor Spacecraft Imaging. In *Proc. of the American Control Conf. (ACC)*, pp. 1808–1813, 2022.
- Armin Norouzi, Saeid Shahpouri, David Gordon, Mahdi Shahbakhti, and Charles Robert Koch. Safe deep reinforcement learning in diesel engine emission control. *Proc. of the Institution of Mechanical Engineers. Part I: Journal of Systems and Control Engineering*, 2023.
- Christian Pek, Stefanie Manzing, Markus Koschi, and Matthias Althoff. Using online verification to prevent autonomous vehicles from causing accidents. *Nature Machine Intelligence*, 2(9):518–528, 2020.
- Tu-Hoa Pham, Giovanni De Magistris, and Ryuki Tachibana. OptLayer - Practical Constrained Optimization for Deep Reinforcement Learning in the Real World. In *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, pp. 6236–6243, 2018.
- Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah Dormann. Stable-Baselines3: Reliable Reinforcement Learning Implementations. *Journal of Machine Learning Research*, 22(268):1–8, 2021.
- Hendrik Roehm, Jens Oehlerking, Matthias Woehrle, and Matthias Althoff. Model conformance for cyber-physical systems: A survey. *ACM Transactions on Cyber-Physical Systems*, 3(3):1–26, 2019.
- William Saunders, Girish Sastry, Andreas Stuhlmüller, and Owain Evans. Trial without Error: Towards Safe Reinforcement Learning via Human Intervention. In *Proc. of the Int. Conf. on Autonomous Agents and MultiAgent Systems (AAMAS)*, pp. 2067 – 2069, 2018.
- Lukas Schäfer, Felix Gruber, and Matthias Althoff. Scalable Computation of Robust Control Invariant Sets of Nonlinear Systems. *IEEE Transactions on Automatic Control*, pp. 1–15, 2023.
- Lukas M Schmidt, Georgios D Kontes, Axel Plinge, and Christopher Mutschler. Can you trust your autonomous car? Interpretable and verifiably safe reinforcement learning. In *Proc. of the IEEE Intelligent Vehicles Symposium (IV)*, pp. 171–178, 2021.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal Policy Optimization Algorithms. *arXiv*, abs/1707.0, 2017.
- Bastian Schürmann, Niklas Kochdumper, and Matthias Althoff. Reachset Model Predictive Control for Disturbed Nonlinear Systems. In *Proc. of the IEEE Conf. on Decision and Control (CDC)*, pp. 3463–3470, 2018.
- Gautham Nayak Seetanadi, Karl-Erik Årzén, and Martina Maggio. Adaptive routing with guaranteed delay bounds using safe reinforcement learning. In *ACM Int. Conf. Proc. Series*, pp. 149 – 160, 2020.

- Mahmoud Selim, Amr Alanwar, M Watheq El-Kharashi, Hazem M Abbas, and Karl H Johansson. Safe Reinforcement Learning using Data-Driven Predictive Control. In *Proc. of the Int. Conf. on Communications, Signal Processing, and their Applications (ICCSPA)*, pp. 1–6, 2022a.
- Mahmoud Selim, Amr Alanwar, Shreyas Kousik, Grace Gao, Marco Pavone, and Karl H. Johansson. Safe Reinforcement Learning Using Black-Box Reachability Analysis. *IEEE Robotics and Automation Letters*, 7(4):10665–10672, 2022b.
- Yifei Simon Shao, Chao Chen, Shreyas Kousik, and Ram Vasudevan. Reachability-Based Trajectory Safe-guard (RTS): A Safe and Fast Reinforcement Learning Safety Layer for Continuous Control. *IEEE Robotics and Automation Letters*, 6(2):3663–3670, 2021.
- Suhas Shyamsundar, Tommaso Mannucci, and Erik-Jan Van Kampen. Reinforcement learning based algorithm with Safety Handling and Risk Perception. In *IEEE Symposium Series on Computational Intelligence (SSCI)*, 2017.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George Van Den Driessche, Thore Graepel, and Demis Hassabis. Mastering the game of Go without human knowledge. *Nature*, 550(7676):354–359, 2017.
- Adam Stooke, Joshua Achiam, and Pieter Abbeel. Responsive safety in reinforcement learning by PID lagrangian methods. In *Proc. of the Int. Conf. on Machine Learning (ICML)*, pp. 9133–9143, 2020.
- Zachary Sunberg and Mykel Kochenderfer. Online Algorithms for POMDPs with Continuous State, Action, and Observation Spaces. *Proc. of the Int. Conf. on Automated Planning and Scheduling (ICAPS)*, 28(1): 259–263, 2018.
- Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. A Bradford Book, 2nd edition, 2018.
- Daniel Tabas and Baosen Zhang. Computationally Efficient Safe Reinforcement Learning for Power Systems. In *Proc. of the American Control Conf. (ACC)*, pp. 3303–3310, 2022.
- Andrew Taylor, Andrew Singletary, Yisong Yue, and Aaron Ames. Learning for Safety-Critical Control with Control Barrier Functions. In *Proc. of the Conf. on Learning for Dynamics and Control*, pp. 708–717, 2020.
- Andrew J Taylor, Andrew Singletary, Yisong Yue, and Aaron D Ames. A Control Barrier Perspective on Episodic Learning via Projection-to-State Safety. *IEEE Control Systems Letters*, 5(3):1019–1024, 2021.
- Brijen Thananjeyan, Ashwin Balakrishna, Suraj Nair, Michael Luo, Krishnan Srinivasan, Minh Hwang, Joseph E Gonzalez, Julian Ibarz, Chelsea Finn, and Ken Goldberg. Recovery RL: Safe Reinforcement Learning With Learned Recovery Zones. *IEEE Robotics and Automation Letters*, 6(3):4915–4922, 2021.
- Jakob Thumm and Matthias Althoff. Provably Safe Deep Reinforcement Learning for Robotic Manipulation in Human Environments. In *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, pp. 6344–6350, 2022.
- Matteo Turchetta, Felix Berkenkamp, and Andreas Krause. Safe exploration in finite Markov decision processes with Gaussian processes. In *Proc. of the Int. Conf. on Neural Information Processing Systems (NIPS)*, pp. 4312–4320, 2016.
- Stephen A Vavasis. *Complexity Theory: Quadratic Programming*. Springer, Boston, MA., 2001.
- Kim Peter Wabersich and Melanie N Zeilinger. A predictive safety filter for learning-based control of constrained nonlinear dynamical systems. *Automatica*, 129(1):109597–109614, 2021.

- Chengyu Wang, Luhan Wang, Zhaoming Lu, Xinghe Chu, Zhengrui Shi, Jiayin Deng, Tianyang Su, Guochu Shou, and Xiangming Wen. SRL-TR2: A Safe Reinforcement Learning Based TRajectory TRacker Framework. *IEEE Transactions on Intelligent Transportation Systems*, 24(6):5765–5780, 2023.
- Linghao Wang, Miao Wang, and Yujun Zhang. A Safe Training Approach for Deep Reinforcement Learning-based Traffic Engineering. In *Proc. of the IEEE Int. Conf. on Communications (ICC)*, pp. 1450–1455. IEEE, 2022.
- Xiao Wang. Ensuring Safety of Learning-Based Motion Planners Using Control Barrier Functions. *IEEE Robotics and Automation Letters*, 7(2):4773–4780, 2022.
- Peter Wieland and Frank Allgöwer. Constructive safety using control barrier functions. *IFAC Proc. Volumes*, 40(12):462–467, 2007.
- Chenchen Yang, Jing Liu, Haiying Sun, Junfeng Sun, Xiang Chen, and Lipeng Zhang. Safe Reinforcement Learning for CPSs via Formal Modeling and Verification. In *Proc. of the Int. Joint Conf. on Neural Networks (IJCNN)*, pp. 1–8, 2021.
- Tsung-Yen Yang, Justinian Rosca, Karthik Narasimhan, and Peter J Ramadge. Projection-based constrained policy optimization. In *Proc. of the Int. Conf. on Learning Representations (ICLR)*, pp. 1–21, 2020.
- Fei Ye, Shen Zhang, Pin Wang, and Ching Yao Chan. A survey of deep reinforcement learning algorithms for motion planning and control of autonomous vehicles. In *Proc. of the IEEE Intelligent Vehicles Symposium (IV)*, pp. 1073–1080, 2021.
- Peipei Yu, Hongcai Zhang, Yonghua Song, Hongxun Hui, and Ge Chen. District Cooling System Control for Providing Operating Reserve Based on Safe Deep Reinforcement Learning. *IEEE Transactions on Power Systems*, pp. 1–13, 2023.
- Mario Zanon and Sebastien Gros. Safe Reinforcement Learning Using Robust MPC. *IEEE Transactions on Automatic Control*, 66(8):3638–3652, 2021.
- Jin Zhang, Yuxiang Guan, Liang Che, and Mohammad Shahidehpour. EV Charging Command Fast Allocation Approach based on Deep Reinforcement Learning with Safety Modules. *IEEE Transactions on Smart Grid*, 2023.
- Wenshuai Zhao, Jorge Pena Queralta, and Tomi Westerlund. Sim-to-Real Transfer in Deep Reinforcement Learning for Robotics: A Survey. In *Proc. of the IEEE Symposium Series on Computational Intelligence (SSCI)*, pp. 737–744, 2020.

A Appendix

MDP modification with action replacement

Action replacement alters the MDP on which the agent learns. Hunt et al. (2021) discuss this modification for discrete action spaces and uniformly sampling from the safe action space. We generalize this discussion to use any replacement function and also include continuous action spaces. We define $\psi(\mathbf{s})$ so that it randomly samples the replacement action $\tilde{\mathbf{a}}$ according to a replacement policy $\pi_r(\tilde{\mathbf{a}}|\mathbf{s})$ with $\sum_{\tilde{\mathbf{a}} \in \mathbb{A}_\varphi(\mathbf{s})} \pi_r(\tilde{\mathbf{a}}|\mathbf{s}) = 1$ for the discrete case, and $\int_{\mathbb{A}_\varphi(\mathbf{s})} \pi_r(\tilde{\mathbf{a}}|\mathbf{s}) d\tilde{\mathbf{a}} = 1$ for the continuous case, and $\pi_r(\tilde{\mathbf{a}}|\mathbf{s}) \geq 0 \forall \tilde{\mathbf{a}} \in \mathbb{A}_\varphi(\mathbf{s})$. In the example of uniform sampling from $\mathbb{A}_\varphi(\mathbf{s})$, the replacement policy is $\pi_r(\tilde{\mathbf{a}}|\mathbf{s}) = 1/|\mathbb{A}_\varphi(\mathbf{s})|$. By replacing unsafe actions, the transition function of the MDP changes to

$$T_\varphi(\mathbf{s}, \mathbf{a}, \mathbf{s}') = \begin{cases} T(\mathbf{s}, \mathbf{a}, \mathbf{s}'), & \text{if } \varphi(\mathbf{s}, \mathbf{a}) = 1 \\ T_r(\mathbf{s}, \mathbf{s}'), & \text{otherwise,} \end{cases} \quad (15)$$

$$T_r(\mathbf{s}, \mathbf{s}') = \sum_{\tilde{\mathbf{a}} \in \mathbb{A}_\varphi(\mathbf{s})} \pi_r(\tilde{\mathbf{a}}|\mathbf{s}) T(\mathbf{s}, \tilde{\mathbf{a}}, \mathbf{s}'). \quad (16)$$

The reward function of the MDP changes accordingly to

$$r_\varphi(\mathbf{s}, \mathbf{a}) = \begin{cases} r(\mathbf{s}, \mathbf{a}), & \text{if } \varphi(\mathbf{s}, \mathbf{a}) = 1 \\ r_r(\mathbf{s}), & \text{otherwise,} \end{cases} \quad (17)$$

$$r_r(\mathbf{s}) = \sum_{\tilde{\mathbf{a}} \in \mathbb{A}_\varphi(\mathbf{s})} \pi_r(\tilde{\mathbf{a}}|\mathbf{s}) r(\mathbf{s}, \tilde{\mathbf{a}}). \quad (18)$$

In the continuous case, we get $T_r(\mathbf{s}, \mathbf{s}')$ by marginalizing the transition probability density function over $\mathbb{A}_\varphi(\mathbf{s})$

$$T_r(\mathbf{s}, \mathbf{s}') = \int_{\mathbb{A}_\varphi(\mathbf{s})} \pi_r(\tilde{\mathbf{a}}|\mathbf{s}) T(\mathbf{s}, \tilde{\mathbf{a}}, \mathbf{s}') d\tilde{\mathbf{a}}, \quad (19)$$

and $r_r(\mathbf{s})$ analogously

$$r_r(\mathbf{s}) = \int_{\mathbb{A}_\varphi(\mathbf{s})} \pi_r(\tilde{\mathbf{a}}|\mathbf{s}) r(\mathbf{s}, \tilde{\mathbf{a}}) d\tilde{\mathbf{a}}. \quad (20)$$

Environment Parameters

We give an overview of all environment-specific parameters in Table 3 and Table 4.

Table 3: Pendulum environment parameters.

Parameter	Value
Gravity g	9.81 m s^{-2}
Mass m	1 kg
Length l	1 m

Table 4: 2D quadrotor environment parameters.

Parameter	Value
Gravity g	9.81 m s^{-2}
k	1 l/kg
d_0	70
d_1	17
n_0	55
\mathbb{W}	$[[[-0.1, 0.1], [-0.1, 0.1]]]$

Hyperparameters for learning algorithms

We specify the hyperparameters for all learning algorithms (see Table 5 for PPO, Table 6 for TD3, and Table 7 for DQN) that are different from the Stable Baselines3 (Raffin et al., 2021) default values. Additionally, the code for the experiments is submitted as supplementary material to further increase the transparency.

Table 5: Hyperparameters for PPO.

Parameter	Pendulum	2D quadrotor
Learning rate	$1\text{E} - 4$	$5\text{E} - 5$
Discount factor γ	0.98	0.999
Steps per update	2048	512
Optimization epochs	20	30
Minibatch size	16	128
Max gradient clipping	0.9	0.5
Entropy coefficient	$1\text{E} - 3$	$2\text{E} - 6$
Value function coefficient	0.045	0.5
Clipping range	0.3	0.1
GAE λ	0.8	0.92
Activation function	ReLU	ReLU
Hidden layers	2	2
Neurons per layer	32	64
Training steps	60k	100k

Table 6: Hyperparameters for TD3.

Parameter	Pendulum	2D quadrotor
Learning rate	$3.5\text{E} - 3$	$2\text{E} - 3$
Replay buffer size	1E4	1E5
Discount factor γ	0.98	0.98
Initial exploration steps	10E3	100
Steps between model updates	256	5
Gradient steps per model update	256	10
Minibatch size per gradient step	512	512
Soft update coefficient τ	$5\text{E} - 3$	$5\text{E} - 3$
Gaussian smoothing noise σ	0.2	0.12
Activation function	ReLU	ReLU
Hidden layers	2	2
Neurons per layer	32	64
Training steps	60k	100k

Table 7: Hyperparameters for DQN.

Parameter	Pendulum	2D quadrotor
Learning rate	$2\text{E} - 3$	$1\text{E} - 4$
Replay buffer size	5E4	1E6
Discount factor γ	0.95	0.999 99
Initial exploration steps	500	100
Steps between model updates	8	2
Gradient steps per model update	4	4
Minibatch size per gradient step	512	64
Maximum for gradient clipping	10	100
Update frequency target network	1E3	1E3
Initial exploration probability ϵ	1.0	0.137
Linear interpolation steps of ϵ	6E3	1E4
Final exploration probability ϵ	0.1	0.004
Activation function	Tanh	Tanh
Hidden layers	2	2
Neurons per layer	32	64
Training steps	60k	100k

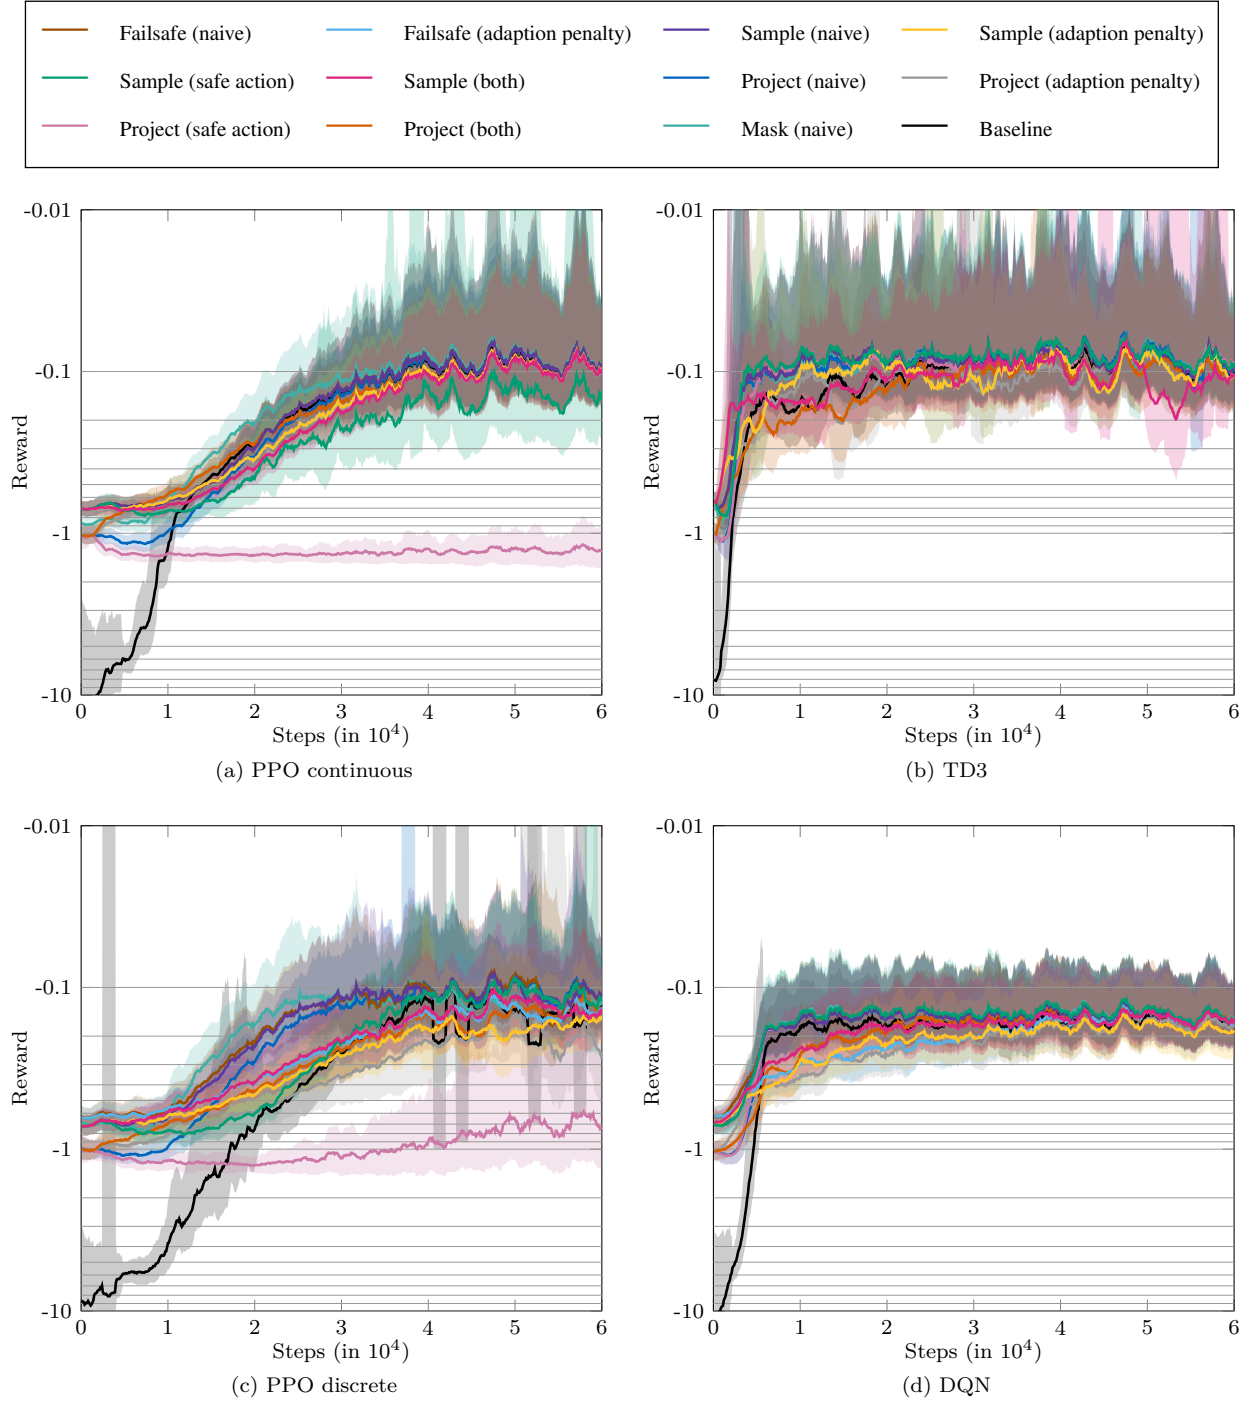


Figure 4: **Pendulum:** Average reward and standard deviation per training step for TD3, DQN, PPO discrete, and PPO continuous. For each configuration, five training runs with different random seeds were conducted. Each subplot contains all implemented variants. Note that for better comparability the reward for the *adaption penalty* variants is still r and the adaption penalty r^* is not included in the curves.

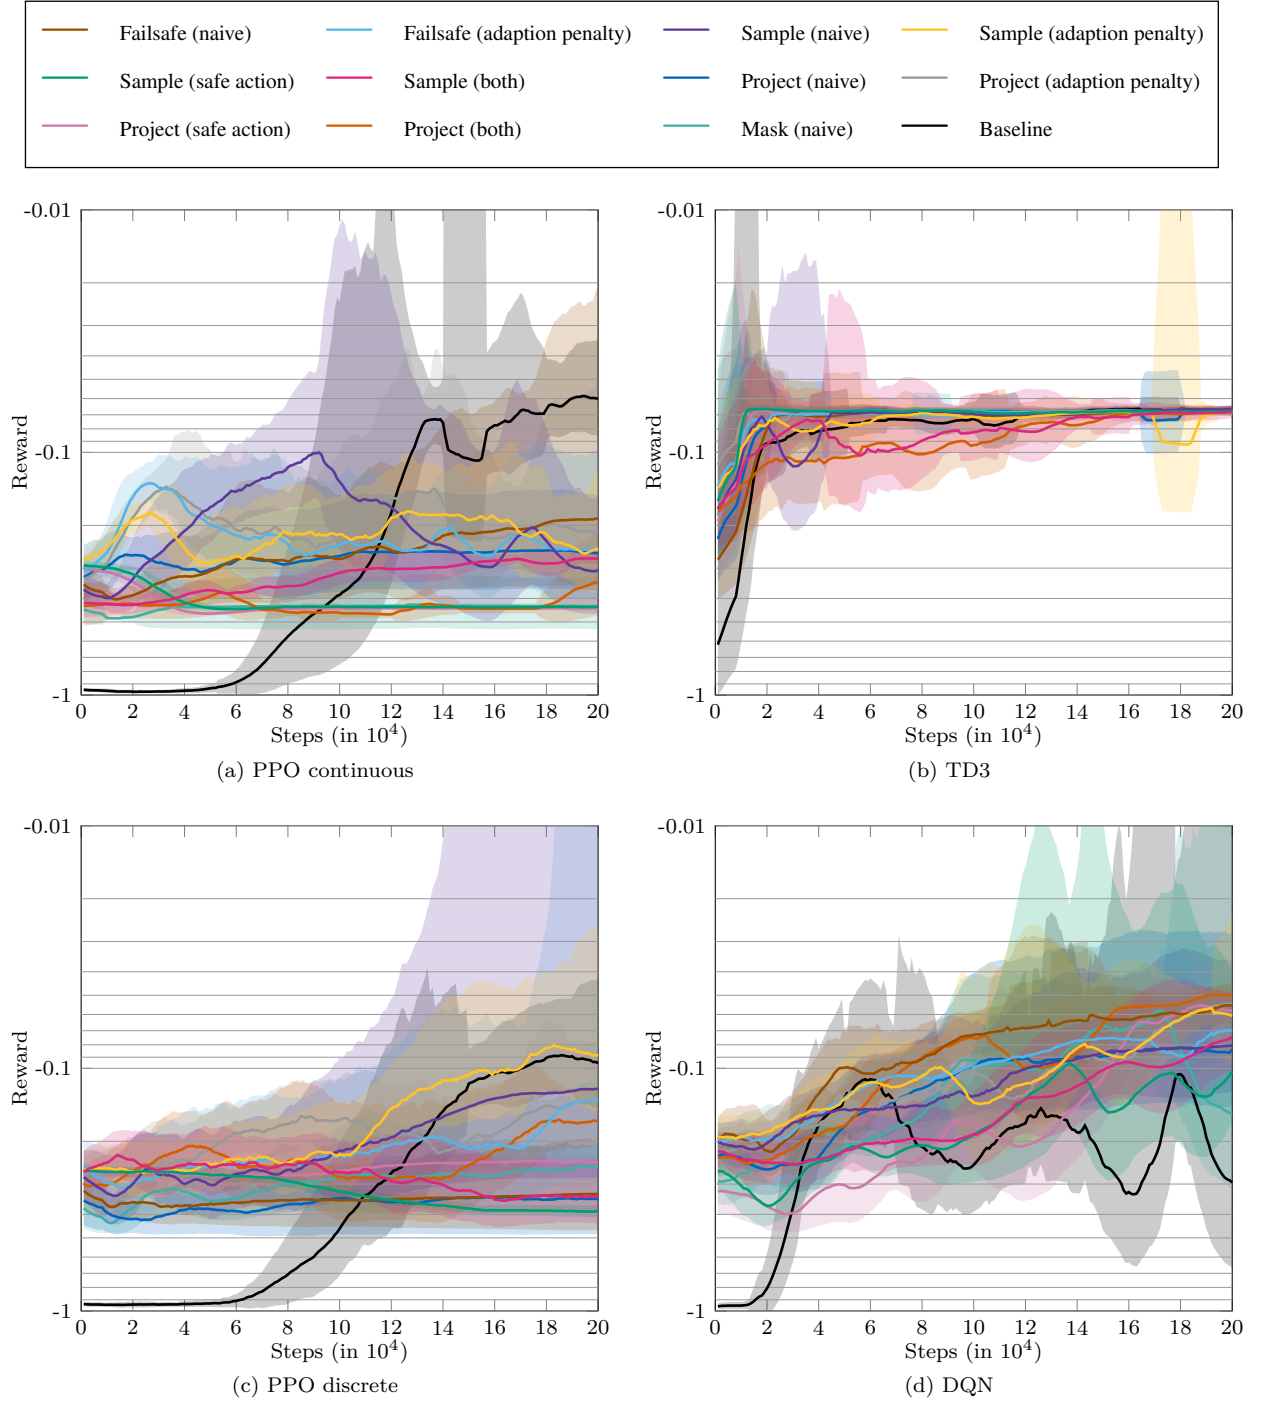


Figure 5: **2D quadrotor**: Average reward and standard deviation per training step for TD3, DQN, PPO discrete, and PPO continuous. For each configuration, five training runs with different random seeds were conducted. Each subplot contains all implemented variants. Note that for better comparability the reward for the *adaption penalty* variants is still r and the adaption penalty r^* is not included in the curves.

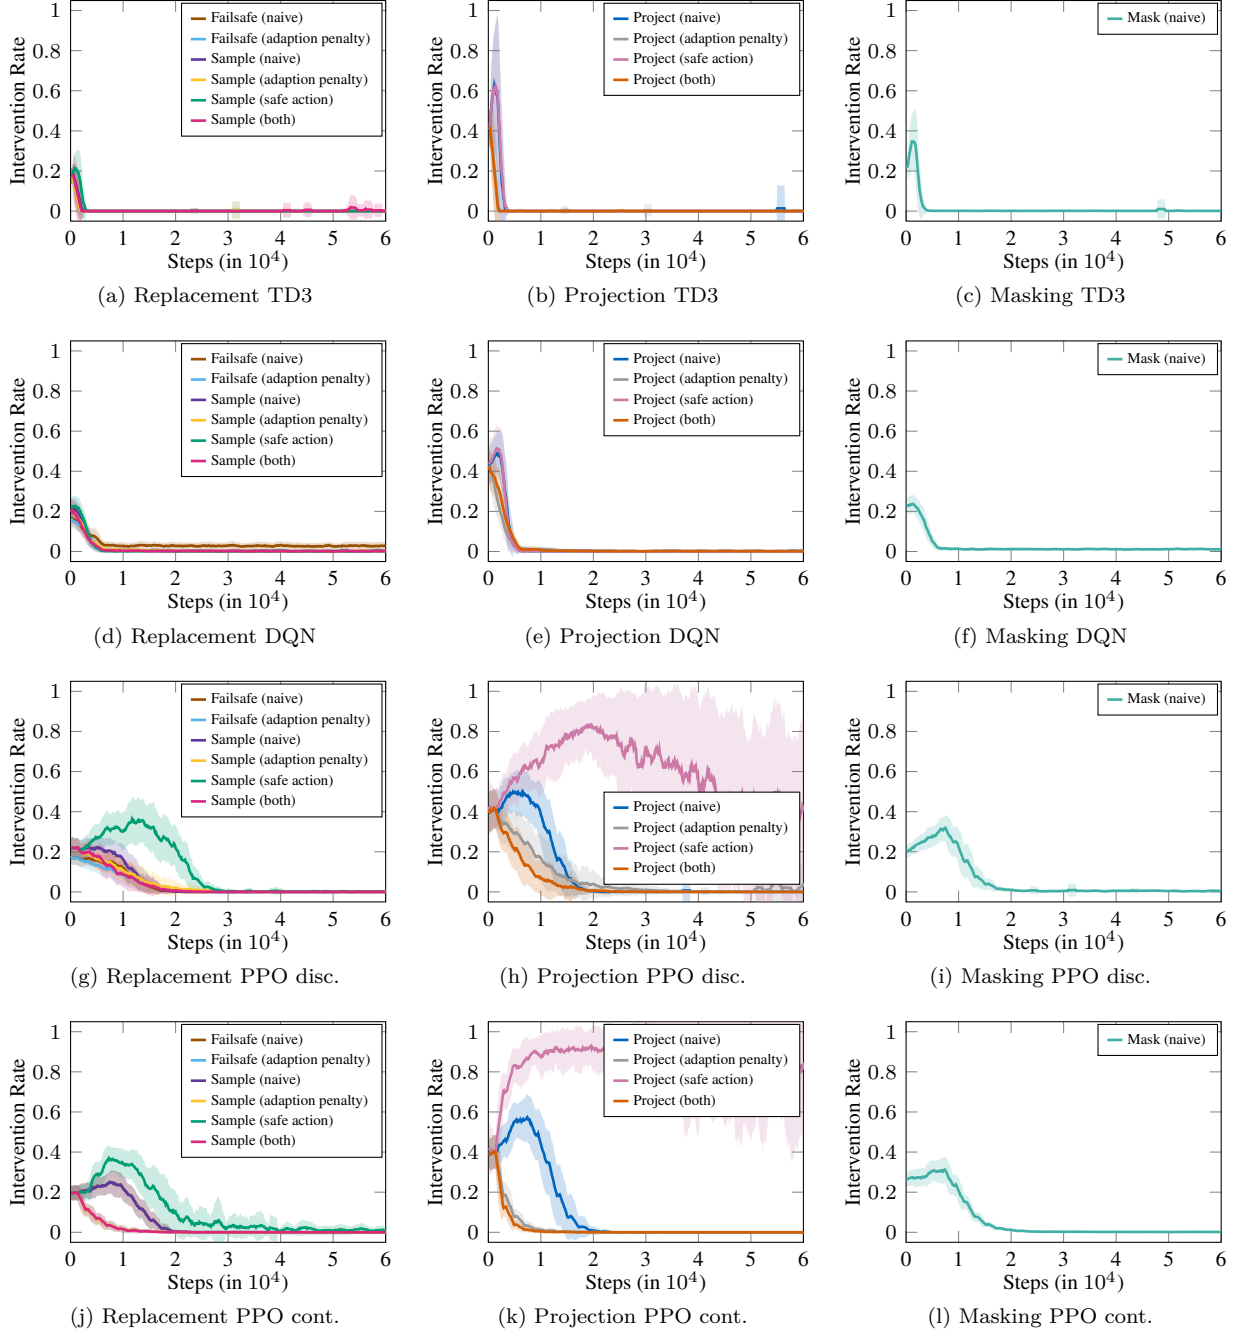


Figure 6: **Pendulum**: Intervention rate for TD3, DQN, PPO discrete, and PPO continuous. Fig. (a) - (c) show the intervention rate of TD3 agents, (d) - (f) for DQN agents, (g) - (i) for discrete PPO agents, and (j) - (l) for continuous PPO agents.

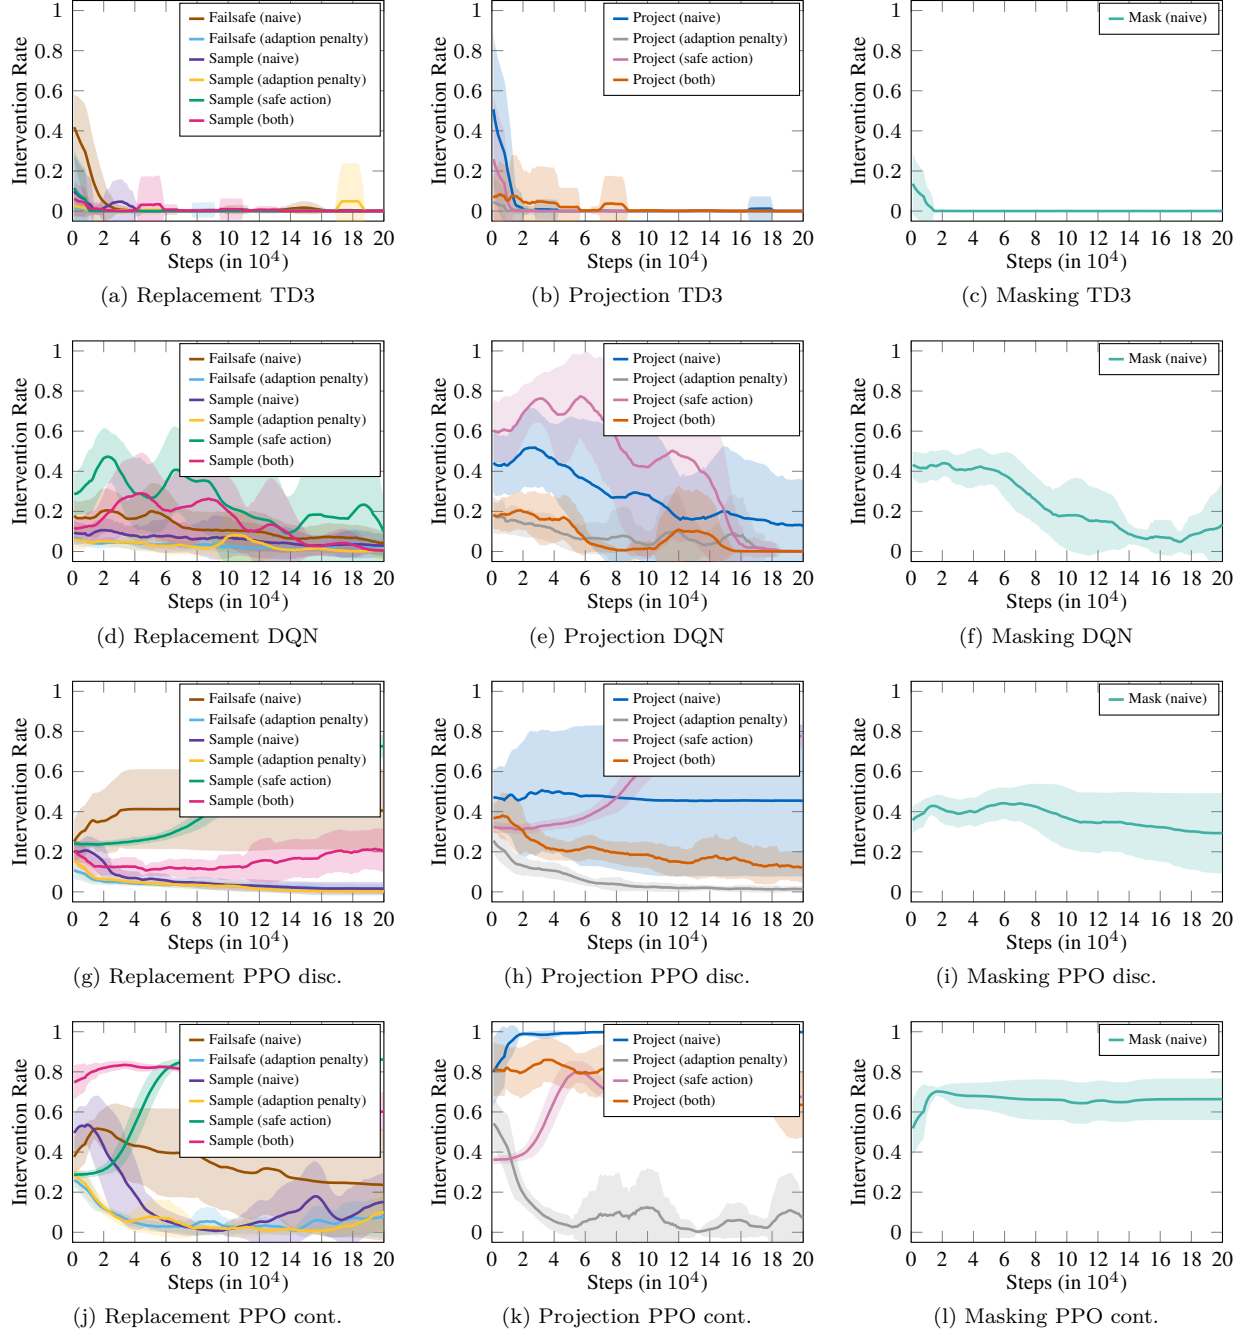


Figure 7: **2D quadrotor**: Intervention rate for TD3, DQN, PPO discrete, and PPO continuous. Fig. (a) - (c) show the intervention rate of TD3 agents, (d) - (f) for DQN agents, (g) - (i) for discrete PPO agents, and (j) - (l) for continuous PPO agents.

Table 8: Mean and standard deviation of 15 pendulum deployment episodes.

Approach	Reward		Intervention Rate		Safety Violation	
	MEAN	STD. DEV.	MEAN	STD. DEV.	MEAN	STD. DEV.
PPO (continuous)						
PROJECTION (SAFEACTION)	-1.14	0.42	0.76	0.29	0.000	0.000
PROJECTION (ADAPTIONPENALTY)	-0.07	0.07	0.00	0.00	0.000	0.000
PROJECTION (BOTH)	-0.07	0.07	0.00	0.00	0.000	0.000
PROJECTION (NAIVE)	-0.06	0.07	0.00	0.00	0.000	0.000
SAMPLE (SAFEACTION)	-0.13	0.16	0.01	0.02	0.000	0.000
SAMPLE (ADAPTIONPENALTY)	-0.07	0.07	0.00	0.00	0.000	0.000
SAMPLE (BOTH)	-0.07	0.07	0.00	0.00	0.000	0.000
SAMPLE (NAIVE)	-0.07	0.07	0.00	0.00	0.000	0.000
FAILSAFE (ADAPTIONPENALTY)	-0.07	0.07	0.00	0.00	0.000	0.000
FAILSAFE (NAIVE)	-0.07	0.07	0.00	0.00	0.000	0.000
BASELINE (NAIVE)	-0.06	0.07	—	—	0.000	0.000
MASKING (ADAPTIONPENALTY)	-0.07	0.07	0.00	0.00	0.000	0.000
MASKING (NAIVE)	-0.07	0.07	0.00	0.00	0.000	0.000
PPO (discrete)						
PROJECTION (SAFEACTION)	-0.52	0.55	0.28	0.42	0.000	0.000
PROJECTION (ADAPTIONPENALTY)	-0.14	0.12	0.00	0.00	0.000	0.000
PROJECTION (BOTH)	-0.09	0.09	0.00	0.00	0.000	0.000
PROJECTION (NAIVE)	-0.15	0.27	0.03	0.10	0.000	0.000
SAMPLE (SAFEACTION)	-0.09	0.08	0.00	0.00	0.000	0.000
SAMPLE (ADAPTIONPENALTY)	-0.13	0.16	0.00	0.00	0.000	0.000
SAMPLE (BOTH)	-0.10	0.08	0.00	0.00	0.000	0.000
SAMPLE (NAIVE)	-0.09	0.08	0.00	0.00	0.000	0.000
FAILSAFE (ADAPTIONPENALTY)	-0.09	0.07	0.00	0.00	0.000	0.000
FAILSAFE (NAIVE)	-0.08	0.07	0.00	0.00	0.000	0.000
BASELINE (NAIVE)	-0.08	0.07	—	—	0.000	0.000
MASKING (ADAPTIONPENALTY)	-0.07	0.07	0.00	0.00	0.000	0.000
MASKING (NAIVE)	-0.07	0.07	0.00	0.00	0.000	0.000
TD3						
PROJECTION (SAFEACTION)	-0.07	0.07	0.00	0.00	0.000	0.000
PROJECTION (ADAPTIONPENALTY)	-0.08	0.08	0.00	0.00	0.000	0.000
PROJECTION (BOTH)	-0.07	0.07	0.00	0.00	0.000	0.000
PROJECTION (NAIVE)	-0.07	0.07	0.00	0.00	0.000	0.000
SAMPLE (SAFEACTION)	-0.07	0.07	0.00	0.00	0.000	0.000
SAMPLE (ADAPTIONPENALTY)	-0.09	0.07	0.00	0.00	0.000	0.000
SAMPLE (BOTH)	-0.09	0.07	0.00	0.00	0.000	0.000
SAMPLE (NAIVE)	-0.07	0.07	0.00	0.00	0.000	0.000
FAILSAFE (ADAPTIONPENALTY)	-0.09	0.07	0.00	0.00	0.000	0.000
FAILSAFE (NAIVE)	-0.07	0.07	0.00	0.00	0.000	0.000
BASELINE (NAIVE)	-0.07	0.07	—	—	0.000	0.000
MASKING (ADAPTIONPENALTY)	-0.07	0.07	0.00	0.00	0.000	0.000
MASKING (NAIVE)	-0.07	0.07	0.00	0.00	0.000	0.000
DQN						
PROJECTION (SAFEACTION)	-0.07	0.07	0.00	0.00	0.000	0.000
PROJECTION (ADAPTIONPENALTY)	-0.07	0.07	0.00	0.00	0.000	0.000
PROJECTION (BOTH)	-0.07	0.08	0.00	0.00	0.000	0.000
PROJECTION (NAIVE)	-0.07	0.07	0.00	0.00	0.000	0.000
SAMPLE (SAFEACTION)	-0.07	0.07	0.00	0.00	0.000	0.000
SAMPLE (ADAPTIONPENALTY)	-0.09	0.07	0.00	0.00	0.000	0.000
SAMPLE (BOTH)	-0.07	0.08	0.00	0.00	0.000	0.000
SAMPLE (NAIVE)	-0.07	0.07	0.00	0.00	0.000	0.000
FAILSAFE (ADAPTIONPENALTY)	-0.07	0.07	0.00	0.00	0.000	0.000
FAILSAFE (NAIVE)	-0.07	0.07	0.00	0.00	0.000	0.000
BASELINE (NAIVE)	-0.07	0.07	—	—	0.000	0.000
MASKING (ADAPTIONPENALTY)	-0.07	0.07	0.00	0.00	0.000	0.000
MASKING (NAIVE)	-0.07	0.07	0.00	0.00	0.000	0.000

Table 9: Mean and standard deviation of 15 2D Quadrotor deployment episodes.

Approach	Reward		Intervention Rate		Safety Violation	
	MEAN	STD. DEV.	MEAN	STD. DEV.	MEAN	STD. DEV.
PPO (continuous)						
PROJECTION (SAFEACTION)	-0.44	0.00	0.68	0.00	0.000	0.000
PROJECTION (ADAPTIONPENALTY)	-0.33	0.07	0.44	0.05	0.000	0.000
PROJECTION (BOTH)	-0.39	0.07	0.57	0.13	0.000	0.000
PROJECTION (NAIVE)	-0.31	0.10	0.47	0.25	0.000	0.000
SAMPLE (SAFEACTION)	-0.43	0.00	0.86	0.00	0.000	0.000
SAMPLE (ADAPTIONPENALTY)	-0.28	0.12	0.10	0.12	0.000	0.000
SAMPLE (BOTH)	-0.36	0.03	0.41	0.20	0.000	0.000
SAMPLE (NAIVE)	-0.39	0.06	0.23	0.13	0.000	0.000
FAILSAFE (ADAPTIONPENALTY)	-0.31	0.11	0.08	0.06	0.000	0.000
FAILSAFE (NAIVE)	-0.27	0.11	0.27	0.29	0.000	0.000
BASELINE (NAIVE)	-0.86	0.01	—	—	0.941	0.009
MASKING (ADAPTIONPENALTY)	-0.43	0.09	0.57	0.28	0.000	0.000
MASKING (NAIVE)	-0.43	0.09	0.57	0.28	0.000	0.000
PPO (discrete)						
PROJECTION (SAFEACTION)	-0.44	0.01	0.74	0.21	0.000	0.000
PROJECTION (ADAPTIONPENALTY)	-0.24	0.13	0.02	0.02	0.000	0.000
PROJECTION (BOTH)	-0.35	0.11	0.16	0.14	0.000	0.000
PROJECTION (NAIVE)	-0.34	0.14	0.46	0.37	0.000	0.000
SAMPLE (SAFEACTION)	-0.42	0.02	0.82	0.01	0.000	0.000
SAMPLE (ADAPTIONPENALTY)	-0.11	0.10	0.00	0.00	0.000	0.000
SAMPLE (BOTH)	-0.38	0.09	0.32	0.18	0.000	0.000
SAMPLE (NAIVE)	-0.13	0.16	0.02	0.03	0.000	0.000
FAILSAFE (ADAPTIONPENALTY)	-0.19	0.15	0.01	0.03	0.000	0.000
FAILSAFE (NAIVE)	-0.34	0.13	0.41	0.21	0.000	0.000
BASELINE (NAIVE)	-0.11	0.03	—	—	0.000	0.000
MASKING (ADAPTIONPENALTY)	-0.25	0.14	0.28	0.21	0.000	0.000
MASKING (NAIVE)	-0.25	0.14	0.28	0.21	0.000	0.000
TD3						
PROJECTION (SAFEACTION)	-0.21	0.03	0.28	0.10	0.000	0.000
PROJECTION (ADAPTIONPENALTY)	-0.22	0.03	0.26	0.10	0.000	0.000
PROJECTION (BOTH)	-0.21	0.04	0.28	0.12	0.000	0.000
PROJECTION (NAIVE)	-0.25	0.05	0.26	0.17	0.000	0.000
SAMPLE (SAFEACTION)	-0.18	0.03	0.07	0.01	0.000	0.000
SAMPLE (ADAPTIONPENALTY)	-0.21	0.04	0.13	0.05	0.000	0.000
SAMPLE (BOTH)	-0.21	0.06	0.06	0.03	0.000	0.000
SAMPLE (NAIVE)	-0.19	0.04	0.12	0.09	0.000	0.000
FAILSAFE (ADAPTIONPENALTY)	-0.20	0.03	0.05	0.02	0.000	0.000
FAILSAFE (NAIVE)	-0.26	0.09	0.05	0.02	0.000	0.000
BASELINE (NAIVE)	-0.90	0.05	—	—	0.945	0.018
MASKING (ADAPTIONPENALTY)	-0.16	0.03	0.03	0.03	0.000	0.000
MASKING (NAIVE)	-0.16	0.03	0.03	0.03	0.000	0.000
DQN						
PROJECTION (SAFEACTION)	-0.06	0.02	0.00	0.01	0.000	0.000
PROJECTION (ADAPTIONPENALTY)	-0.05	0.00	0.00	0.00	0.000	0.000
PROJECTION (BOTH)	-0.05	0.01	0.00	0.00	0.000	0.000
PROJECTION (NAIVE)	-0.09	0.06	0.12	0.24	0.000	0.000
SAMPLE (SAFEACTION)	-0.07	0.01	0.01	0.02	0.000	0.000
SAMPLE (ADAPTIONPENALTY)	-0.06	0.03	0.00	0.00	0.000	0.000
SAMPLE (BOTH)	-0.07	0.02	0.00	0.00	0.000	0.000
SAMPLE (NAIVE)	-0.05	0.01	0.00	0.00	0.000	0.000
FAILSAFE (ADAPTIONPENALTY)	-0.06	0.02	0.02	0.04	0.000	0.000
FAILSAFE (NAIVE)	-0.07	0.03	0.10	0.10	0.000	0.000
BASELINE (NAIVE)	-0.24	0.37	—	—	0.199	0.398
MASKING (ADAPTIONPENALTY)	-0.15	0.16	0.14	0.26	0.000	0.000
MASKING (NAIVE)	-0.15	0.16	0.14	0.26	0.000	0.000