

Safe Deep Reinforcement Learning by Verifying Task-Level Properties

Enrico Marchesini*
Northeastern University

Boston (MA), USA
e.marchesini@northeastern.edu

Alessandro Farinelli
University of Verona

Verona, Italy
alessandro.farinelli@univr.it

Luca Marzari*

University of Verona
Verona, Italy
luca.marzari@univr.it

Christopher Amato
Northeastern University

Boston (MA), USA
c.amato@northeastern.edu

ABSTRACT

Cost functions are commonly employed in Safe Deep Reinforcement Learning (DRL). However, the cost is typically encoded as an indicator function due to the difficulty of quantifying the risk of policy decisions in the state space. Such an encoding requires the agent to visit numerous unsafe states to learn a cost-value function to drive the learning process toward safety. Hence, increasing the number of unsafe interactions and decreasing sample efficiency. In this paper, we investigate an alternative approach that uses domain knowledge to quantify the risk in the proximity of such states by defining a *violation* metric. This metric is computed by verifying task-level properties, shaped as input-output conditions, and it is used as a penalty to bias the policy away from unsafe states without learning an additional value function. We investigate the benefits of using the violation metric in standard Safe DRL benchmarks and robotic mapless navigation tasks. The navigation experiments bridge the gap between Safe DRL and robotics, introducing a framework that allows rapid testing on real robots. Our experiments show that policies trained with the violation penalty achieve higher performance over Safe DRL baselines and significantly reduce the number of visited unsafe states.

KEYWORDS

Deep Reinforcement Learning; Safety; Robot Navigation

ACM Reference Format:

Enrico Marchesini*, Luca Marzari*, Alessandro Farinelli, and Christopher Amato. 2023. Safe Deep Reinforcement Learning by Verifying Task-Level Properties. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 10 pages.

1 INTRODUCTION

Safe Deep Reinforcement Learning (DRL) approaches typically foster safety by limiting the accumulation of costs caused by unsafe interactions [12]. Defining informative cost functions, however, has the same issues as designing rewards [13] due to the difficulty of quantifying the risk around unsafe states. For this reason, recent

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

* indicates equal contribution.

works rely on indicator functions, where a positive value deems a state unsafe [22]. In detail, the cost refers to a state-action pair (s, a) , and it is backed up to propagate safety information and estimate a cost-value function. The cost metric and its value estimation have been used to drive the learning process towards safety using penalties [33, 43], cumulative or instantaneous constraints [21, 41]. However, the learned value functions have poor estimates and end up in local optima [14, 25], limiting their efficacy in fostering safety. We argue that the cost's sparse nature is another key issue that hinders safety and sample efficiency. A sparse definition of the cost requires the agent to visit unsafe states to learn good estimates from the sparse feedback. These issues are pivotal in a safety context where we aim to minimize the number of visited unsafe states.

In this direction, we note that the indicator costs do not carry information about areas around s where it is risky to perform the action that led to deeming s unsafe. For example, consider a navigation scenario where a policy chooses the robot's velocity, given its position. In this context, prior works trigger a positive cost when colliding in a state s (Figure 1 on the left) and need to visit similar unsafe interactions around s to learn the cost-value function that drives the learning process toward safety [7, 21, 41]. In contrast, it is possible to exploit high-level system specifications (e.g., robots' size and velocity) to define an area of size ϵ around s where performing the action a that led to collision would result in other unsafe interactions (Figure 1 on the right). We can then compute a safety value based on the policy's decisions, avoiding visiting such an *unsafe* area to learn a cost-value function. Hence, the idea is to replace indicator costs and the learning of cost-value functions by quantifying the states in the unsafe area where the policy chooses a and use it as a penalty to discourage unsafe decisions [1, 19]. Our hypothesis is that this procedure can significantly improve sample efficiency and reduce the number of visited hazardous states.

To this end, we propose an approximate *violation* metric as a penalty that uses system specifications to quantify how safe policy




Figure 1: Indicator cost function (left). Unsafe interactions, caused by the same action, around the unsafe state (right).

decisions are around s . Following recent literature [20, 26, 36], we encode whether a policy chooses specific actions (*outputs*) in a subspace of the state space (*inputs*) as input-output conditional statements, commonly referred to as task-level *properties* (or state-action mappings). Formal Verification (FV) approaches for Deep Neural Networks (DNNs) have been used to formally check all the states where the policy violates such properties [45]. In particular, [10] introduced a formal violation metric by provably quantifying how well the policy respects the properties. Such a formal violation naturally addresses the limitation of indicator cost functions but has not been previously investigated to foster safety in DRL, as FV has two main issues. (i) It is a *NP-Complete* problem [17] that makes it intractable to compute the formal violation metric at training time without a prohibitive overhead. (ii) The state-action mappings are hard-coded, which could be unfeasible in tasks with complex or unknown specifications. Against this background, we make the following contributions to the state-of-the-art:

- We replace FV with a sample-based approach that approximates the violation metric with forward propagations of the agent’s network on the sampled states. Such an approximation empirically shows a negligible error over the value computed with FV in a fraction of the computation time.
- We generate an additional state-action mapping when performing an unsafe interaction during the training, using a fixed-size area around the visited unsafe state and the action that led to such interaction.
- We show the advantages of using our approximate violation as a penalty in existing DRL algorithms and employ FV [10] on the trained policies to show that our approach allows learning safer behaviors (i.e., lower violations).

Our empirical evaluation considers a set of Unity [15] robotic mapless navigation tasks [29, 42, 51]. In contrast to Safe DRL benchmarks (e.g., SafeMuJoCo [21]), our scenarios allow transferring policies directly on the robot to foster the development of Safe DRL approaches in realistic applications. We also evaluate violation-based approaches in standard SafeMuJoCo tasks. In all scenarios, we compare with unconstrained DRL baselines augmented with a cost penalty [39, 47], and constrained DRL [41]. Our evaluation shows that cost-based algorithms have higher costs and higher violations, confirming the lack of information provided by indicator cost functions. In contrast, the approximate violation-based penalty drastically reduces unsafe behaviors (i.e., lower cost and violation) while preserving good returns during and after training.

2 PRELIMINARIES

A DRL problem is typically modeled as a Markov Decision Process (MDP), described by a tuple $\langle S, A, P, r, \gamma \rangle$ where S is the state space, A is the action space, $P : S \times A \rightarrow S$ is the state transition function, $r : S \times A \rightarrow \mathbb{R}$ is a reward function, and $\gamma \in [0, 1)$ is the discount factor. In particular, given a policy $\pi \in \Pi := \{\pi(a|s) : s \in S, a \in A\}$, the agent aims to maximize the expected discounted return for each trajectory $\tau = (s_0, a_0, r_0, \dots)$:

$$\max_{\pi \in \Pi} J_r^\pi := \mathbb{E}_{\tau \sim \pi} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \right] \quad (1)$$

A common approach to fostering safety is adding a penalty to the reward to learn how to avoid unsafe interactions characterized by lower payoffs [12]. Otherwise, MDPs can be extended to Constrained MDPs (CMDPs) to incorporate a set of constraints C defined on $C_{0,\dots,c} : S \times A \rightarrow \mathbb{R}$ cost functions (where c is the number of constraints) and their thresholds $t_{0,\dots,c}$ [3]. A C_i -return is defined as $J_{C_i}^\pi := \mathbb{E}_{\tau \sim \pi} [\sum_{t=0}^{\infty} \gamma^t C_i(s_t, a_t)]$. Constraint-satisfying (feasible) policies Π_C , and optimal policies π^* are thus defined as:

$$\Pi_C := \{\pi \in \Pi : J_{C_i}^\pi \leq t_i, \forall i \in [0, \dots, c]\}, \quad \pi^* = \max_{\pi \in \Pi_C} J_r^\pi \quad (2)$$

Constrained DRL algorithms aim at maximizing the expected return J_r^π while maintaining costs under hard-coded thresholds \mathbf{t} :

$$\max_{\pi \in \Pi} J_r^\pi \quad \text{s.t.} \quad J_C^\pi \leq \mathbf{t} \quad (3)$$

Such a constrained optimization problem is usually transformed into an equivalent unconstrained one using the Lagrangian method [35, 38, 41], resembling a penalty-based approach. Recently, Tessler et al. [43] also argued that constrained DRL has significant limitations as it requires a parametrization of the policy (i.e., it does not work with value-based DRL), a propagation of the constraint violation over trajectories and works with limited types of constraints (e.g., cumulative, instantaneous). For these reasons, the authors show the efficacy of integrating a penalty signal into the reward function. This motivates our choice of using penalty-based approaches to evaluate the benefits of incorporating the proposed approximate violation as a penalty to foster safety.

2.1 Properties and Violation

From FV literature [20], a property P is hard-coded using task-level knowledge as a pre-condition R and a post-condition Q (i.e., $P := \langle R, Q \rangle$). In a DRL context, R is the domain of the property (i.e., the area around s), and Q is the co-domain (i.e., an action). Broadly speaking, given a DNN \mathcal{N} with y_1, \dots, y_n outputs, R is defined with an interval ϵ_i for each input i of \mathcal{N} , which we refer to as ϵ -area, and Q represents different desiderata on the output [4]. For example, in a value-based setup where the policy selects the action with the highest value, the post-condition Q of a property designed for safety models *never select the action corresponding to the output* $y_k \in y_1, \dots, y_n$. Formally, Q checks the following inequality:

$$y_k < y_i \quad \forall i \in [1, n] - \{k\} \quad (4)$$

that extends to continuous actions as shown in prior work [20, 45].

Recent works introduced a *violation* metric to quantify the number of violations in the domain of the property using verification techniques [10, 31]. In more detail, the violation is defined as the ratio between the size of $R' \subseteq R$ where the post-condition is violated (i.e., the inequality does not hold and y_k is selected) and R . Such a provable violation carries the task-level information of the properties and quantifies how often a property is violated. Hence, when using properties to model safety specifications, which we refer to as safety properties, the violation represents a locally-aware cost function.

3 METHODS

We aim to investigate the benefits of combining a reward and a violation-based penalty value into an MDP. Following prior penalty-based approaches [12, 43], we maximize the following objective:

$$\max_{\pi \in \Pi} J_{r,C}^{\pi} := \mathbb{E}_{\tau \sim \pi} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) - Z(\cdot) \right] \quad (5)$$

where $Z(\cdot)$ is a generic penalty function. For example, a violation-penalty is $Z_{\pi}(s_t \pm \epsilon)$, indicating that the violation depends on the policy decisions in a proximity ϵ of the state (i.e., the ϵ -area, or R). Equation 5 has two core benefits over other Safe DRL methods based on constraints: (i) penalty objectives potentially maintain the same optimal policy of the underlying MDP as they do not constrain exploration nor reduce the space of feasible policies as constrained approaches [34].¹ (ii) Constrained DRL typically estimates an advantage function to propagate cost information, hindering their application with values that strictly depend on the current policy. Moreover, this requires visiting unsafe states to learn effective estimates for the sparse cost values. In contrast, penalty-based methods do not require a separate advantage estimate. In addition, various DRL algorithms, such as Proximal Policy Optimization (PPO) [39], provide significant empirical evidence of the benefits of using penalties instead of constraints [43].

3.1 Approximate Violation

A violation metric computed on a safety property quantifies the number of unsafe policy decisions over an area of the state space around the state s . Such a local component and the task-level knowledge inherited by the safety properties naturally address the indicator cost functions’ lack of information. However, the *NP-Completeness* of FV [17] makes the provable violation computation intractable during the training due to the significant overhead of verification tools [18]. Hence, we address the computational demands of FV by proposing a novel sample-based method to approximate the violation.

Given a DNN \mathcal{N} that parameterizes a policy, a list of properties $\mathbf{P} = \langle \mathbf{R}, \mathbf{Q} \rangle$, and the current state s , we aim at checking whether $y_k < y_i \forall i \in [1, n] - \{k\}$ (where $\mathbf{y} := [y_1, \dots, y_n]$ are the outputs of \mathcal{N} given s as input). The general flow of our method is presented in Algorithm 1: first, we embed the condition (4) in the network architecture by concatenating a new output layer that implements $y_i - y_k \forall i \in [1, n]$ (line 1). We refer to this augmented network as \mathcal{N}' . If s is contained in one or more pre-condition (i.e., it is deemed risky according to the properties), we consider such a subset of properties $\mathbf{P}' \subseteq \mathbf{P}$ to compute the approximate violation (line 2). Hence, we randomly sample a set of states I from the pre-conditions $\mathbf{R} \in \mathbf{P}'$ (line 3). Finally, after propagating I through \mathcal{N}' , we enumerate the outputs $\mathcal{N}'(I) \leq 0$, which are the ones that do not satisfy the post-conditions (line 4). Finally, our approximate violation is the ratio between the number of such outputs over the total sampled points (line 5), which closely resembles how the formal violation is computed in [10].

Similarly to the formal violation, our approximation can be interpreted as a locally-aware cost because it is computed using the

¹This is not the case for navigation tasks as safe policies avoid obstacles.

Algorithm 1 Computing the Approximate Violation

- Given** \mathcal{N} with outputs y_1, \dots, y_n , the current state s , properties $\langle \mathbf{R}, \mathbf{Q} \rangle$, and size m of states to sample.
- 1: $\mathcal{N}' \leftarrow$ add a layer with n outputs to \mathcal{N} that implements (4)
 - 2: $\mathbf{P}' \leftarrow \langle \mathbf{R}, \mathbf{Q} \rangle$ if $s \cap R \neq \emptyset, \forall \langle \mathbf{R}, \mathbf{Q} \rangle \in \langle \mathbf{R}, \mathbf{Q} \rangle \triangleright$ i.e., s is unsafe
 - 3: $\mathbf{I} \leftarrow$ Sample m points from $p[\mathbf{R}] \forall p \in \mathbf{P}'$
 - 4: $violation = \text{Count } \max[\mathcal{N}'(I)] \leq 0 \forall I \in \mathbf{I} \triangleright$ considering y_k from $p[\mathbf{Q}] \forall p \in \mathbf{P}'$
 - 5: **return** $violation / (m |\mathbf{P}'|)$

information in a local region around a state. Moreover, our approximate violation:

- Includes safety information of areas of interest due to the state-action mappings (i.e., properties).
- Approximates how often a property violation might occur, having a similar role to Lagrangian multipliers but without requiring additional gradient steps or value estimators.
- It does not require additional environment interactions, drastically reducing the number of visited unsafe states.

Finally, prior verification works only rely on hard-coded properties [10, 46, 48]. Still, it is not uncommon to experience an unsafe state not included in the pre-conditions due to design issues, i.e., $s \cap R = \emptyset \forall R \in \mathbf{R}$. Hence, we generate an additional property upon experiencing an unsafe state s using a fixed-size area around s as R and the performed action as Q . However, hard-coded properties are still crucial as there could be corner cases with more than one action to avoid.

3.1.1 Visual Example. We further detail the approximate violation computation using a visual example. In particular, Figure 2 shows an illustrative example of a DNN and a property Pr (following Section 4.2 formalization). For the sake of simplicity, we show the process assuming to use only $m = 1$ sample from the property pre-condition. Following Algorithm 1, we show \mathcal{N}' on the right of Figure 2, consisting of a new output that implements $y_i - y_k \forall i \in [1, n]$, where y_k with $k = 1$ is the node that represents the action we are interested in avoiding. Considering the example in Figure 1, x_0, \dots, x_3 are the current position, orientation, and distance from

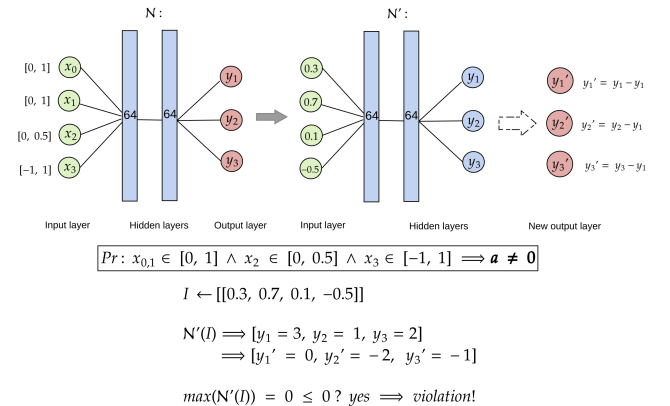


Figure 2: Example of computing the approximate violation.

the goal for the robot. We want to avoid the action with index 0 (i.e., y_1), a forward movement at velocity \vec{v} . We sample the $m = 1$ point from the property pre-condition, obtaining $i \in I$ that is then forward propagated through the network $N'(I)$. Such propagation returns $y_1 = 3, y_2 = 1, y_3 = 2$ in the original output layer, and $y'_1 = 0, y'_2 = -2, y'_3 = -1$ in the output of N' . Finally, we enumerate the states $\in I$ where the maximum of $N'(I)$ is less than or equal to 0, which means that y_1 will be chosen, leading to a violation.

3.2 Limitations

The violation requires hard-coded properties, which are challenging to design when considering agents with unknown dynamics. Our generated property does not consider scenarios where multiple actions are unsafe for the same state, so it remains unclear how to collect and refine properties during the training to model safe behaviors. Hence, as in FV literature [20], we assume having access to task-level knowledge to design the hard-coded properties. In safety-critical contexts, this assumption typically holds. Considering different input types (e.g., images) is conceptually feasible but would require further research and empirical evaluation. To this end, model-based DRL would allow using the model to design the unsafe area. Finally, it is unclear how to provide guarantees in model-free Safe DRL approaches, including our work, constrained DRL [2, 21, 49], and several other approaches summarized in Garcia and Fernández [12]. As discussed in Ilyas et al. [14], using DNNs for approximating policies and values makes the method diverge from the underlying theoretical framework. Nonetheless, we start addressing such key issues by employing existing FV approaches to check the trained policy decisions over the properties of interest.

4 EXPERIMENTS

First, we introduce a set of Turtlebot3-based safety mapless navigation tasks to enable rapid testing of policies in realistic contexts. Our environments rely on Unity [15] as it allows rapid prototyping, Gym compatibility, and interface with the Robotic Operating System (ROS). Mapless navigation is a well-known problem in model-free DRL [42, 51], prior Safe DRL literature [27, 37], and multi-agent DRL [29, 30]. While standard navigation planners use a map of the environment and exhaustive information from the sensors, DRL setups consider more challenging conditions, such as not having a map and relying only on sparse local sensing. We use a similar encoding to prior work [6, 28, 32, 42, 51]: 11 sparse laser scans with a limited range and two values for the target relative position (i.e., distance and heading) as observations. Discrete actions encode angular and linear velocities to reduce training times while maintaining good navigation skills [28].² At step t , the agent receives a reward:

$$r_t = \begin{cases} 1 & \text{if goal reached} \\ \Delta(d_{t-1}, d_t) - \beta & \text{otherwise} \end{cases} \quad (6)$$

The agent thus obtains a dense reward given by the distance (d) difference (Δ) from the goal in two consecutive steps, with a per-step $-\beta$ to incentive shorter paths. Each collision returns a positive cost signal that can be used to compute the desired penalty

²Our environments also support continuous actions and different domain randomization of the tasks and physical properties through the Unity editor.

(e.g., violation), enabling a straightforward application to different penalty-based objectives (as in Equation 5) or constraints.

We introduce four training and one testing environment with different obstacles, namely *Fixed_obs_{T, NT}*, *Dynamic_obs_{T, NT}*, and *Evaluation_NT* depicted in Figure 3. Such a variety of conditions serve to provide different settings for evaluating Safe DRL algorithms in robotic navigation. The environments inherit several characteristics from known benchmarks such as SafetyGym [37] and differ from each other as the obstacles that can be *Fixed* (parallelepiped-shaped static objects) or *Dynamic* (cylindrical-shaped objects that move to random positions at a constant velocity). Moreover, obstacles can be *Terminal* (T) if they end an episode upon collision, or *Non-Terminal* (NT) if the agent can navigate through them.

4.1 Environment Descriptions

Our scenarios share a $4m \times 4m$ size ($6m \times 6m$ for the testing one), randomly generated obstacle-free goals, and a timeout at 500 steps. A list of environments and their other main features follows:

- *Fixed_obs_T* has fixed terminal (T) obstacles. With terminal, we intend that the episode ends upon collision.
- *Fixed_obs_NT* differs from the previous one for the non-terminal (NT) obstacles. The environment returns a signal upon each collision that can be used to model cost functions or other penalties. Non-terminal obstacles are visible to the lidars but non-tangible, i.e., the Turtlebot3 can pass through them. This class of obstacles represents the main challenge to designing safe DRL solutions, as the robot could get more positive rewards by crossing an obstacle at the expense of a higher number of unsafe behaviors.
- *Dynamics_obs_T* has cylindrical-shaped dynamic terminal (T) obstacles. Such obstacles move toward random positions at a constant velocity, representing a harder challenge. The obstacles can travel on the robot's goal, so the agent must learn a wider variety of behaviors (e.g., react to an approaching obstacle, stand still to wait for the goal to clear).
- *Dynamics_obs_NT*: differs from the previous one for the non-terminal (NT) obstacles.
- *Evaluation_NT*: we use this evaluation environment to test the generalization abilities of trained policies to new situations. This scenario is wider and contains both fixed and dynamic non-terminal obstacles of different shapes.




Figure 3: Fixed, Dynamic, Evaluation tasks with different obstacles. Terminal obstacles (T) reset the environment upon collision. Non-terminal ones (NT) allow the robot to cross them, experiencing more unsafe states. The evaluation environment has fixed and dynamic non-terminal obstacles.

Table 1: Average violation (%), and computation time for properties $p_{\uparrow, \leftarrow, \rightarrow}$ calculated using ProVe[10], and our approximation with 100, 1.000, and 10000 samples.

Property	ProVe	Estimation 100	Estimation 1k	Estimation 10k
p_{\uparrow}	81.48 ± 1.2	81.0 ± 0.6	81.1 ± 0.8	81.17 ± 0.5
p_{\leftarrow}	73.9 ± 0.8	73.5 ± 0.2	73.63 ± 0.2	73.65 ± 0.3
p_{\rightarrow}	74.2 ± 0.3	73.6 ± 0.1	73.67 ± 0.1	73.68 ± 0.1
Mean violation:	76.53	76.00	76.13	76.17
Mean computation time:	≈2m37s	≈0.053s	≈ 0.056s	≈0.060s

4.2 Properties for Mapless Navigation

Our properties shape rational, safe behaviors for mapless navigation and are used to compute the approximate violation. Moreover, we consider an online generated property described in Section 3.1. A natural language description of the main hard-coded properties follows:

- p_{\uparrow} : *There is an obstacle close in front \Rightarrow Do not go forward*
- p_{\leftarrow} : *There is an obstacle close to the left \Rightarrow Do not turn left*
- p_{\rightarrow} : *There is an obstacle close to the right \Rightarrow Do not turn right*

We use the maximum agent velocity to determine the size of the area around the (unsafe) states of interest (i.e., the ϵ -area). For example, a formal definition for p_{\uparrow} is:

$$p_{\uparrow} : x_0, \dots, x_4 \in [0, 1] \wedge x_5 \in [0, 0.05] \wedge x_6, \dots, x_{10} \in [0, 1] \wedge x_{11}, x_{12} \in [-1, 1] \implies a \neq 4$$

where x_0, \dots, x_{10} are the 11 lidar values, x_{11}, x_{12} is the relative position of the goal (i.e., distance and heading), and action $a = 4$ corresponds to a forward movement. Crucially, each input x_i potentially considers a different interval to model the area of interest. Hence, p_{\uparrow} checks the policy’s decisions when there is an obstacle close to the front (i.e., $x_5 \in [0, 0.05]$) under any possible target position ($x_{11}, x_{12} \in [-1, 1]$).³ The approximate violation computed over these properties thus contains information about a specific unsafe situation under a general goal configuration (i.e., $x_{11}, x_{12} \in [-1, 1]$).

To assess how good our approximate violation is over the provable one, we compare it with the violation computed by a FV framework. In particular, the formal violation and our approximate one are computed using the above properties. They are averaged over the same ten models collected at random steps during the training. Table 1 shows the violation values for each property $p_{\uparrow, \leftarrow, \rightarrow}$ computed by ProVe [10], and our approximation using 100, 1000, and 10000 samples. Our approximation shows an average 0.69% error over the formal violation even when using only 100 samples. Such an error further decreases to 0.47% by using 10000 samples. By exploiting parallelism and batch computation of modern Deep Learning Frameworks, the increase in computation time for the approximate violation with 100 or 10000 samples is comparable. Conversely, as discussed in Section 3.1, ProVe’s average computation time is orders of magnitude higher with respect to our approach (i.e., 0.06 over 157 seconds). However, our approximate methodology does not formally guarantee the policy behaviors due to its

³We measured that a minimum normalized distance of 0.05 is required for the robot to turn at max speed and avoid obstacles in front.

sample-based nature. For this reason, the next section uses standard FV frameworks on the trained policies to show the provable violation results.

4.3 Empirical Evaluation

Our evaluation aims at showing the following:

- The benefits of integrating reward and safety specifications into a single MDP.
- The advantages of using our violation over indicator cost functions. To assess our claims, we plot the following values averaged over the last 1000 steps: (i) success (i.e., the number of goals reached), (ii) cost, (iii) and the violation at different stages of the training.

Data are collected on an i7-9700k and consider the mean and standard deviation of ten independent runs [9]. We consider the cost and violation penalty objective (5) in a value-based (Dueling Double Deep Q-Network (DuelDDQN) [47]) and policy-based (PPO [39]) baselines, referring to the resultant algorithms as `DuelDDQN_{cost, violation}` and `PPO_{cost, violation}`. We compare with Lagrangian PPO (LPPO) [41] as it is a widely adopted Constrained DRL baseline and achieves state-of-the-art performance in similar SafetyGym navigation-based tasks.⁴ According to the literature, the value-based and policy-gradient baselines should achieve the highest rewards and costs (having no penalty information) [37]. Conversely, LPPO should show a significant trade-off between average cost and reward or fail at maintaining the cost threshold when set to low values [21, 27, 37]. In contrast, we expect the penalty-based methods to achieve promising returns while significantly reducing the cost and the number of violations during the training.

Terminal Results. Our results are shown in the Appendix. The information carried by the violation results in a significant performance advantage, maintaining similar or higher successes over non-violation-based approaches. Moreover, the policy-based algorithms show superior performance over the value-based implementations.

Non-Terminal Results. Given the higher performance of policy-based algorithms, the following experiments omit value-based ones. Figure 4 show the results of (NT) tasks. As in the previous evaluation, PPO achieves a higher number of successes and a higher cost. In contrast, LPPO satisfies the constraint most of the time but achieves the lowest successes and does not learn effective navigation behaviors in *Dynamic_obs_NT*, confirming the performance

⁴For a fair comparison, we set the cost threshold of LPPO to the average cost obtained by `PPO_cost`. The Appendices detail our hyper-parameters and are accessible at the following link: shorturl.at/crSX6

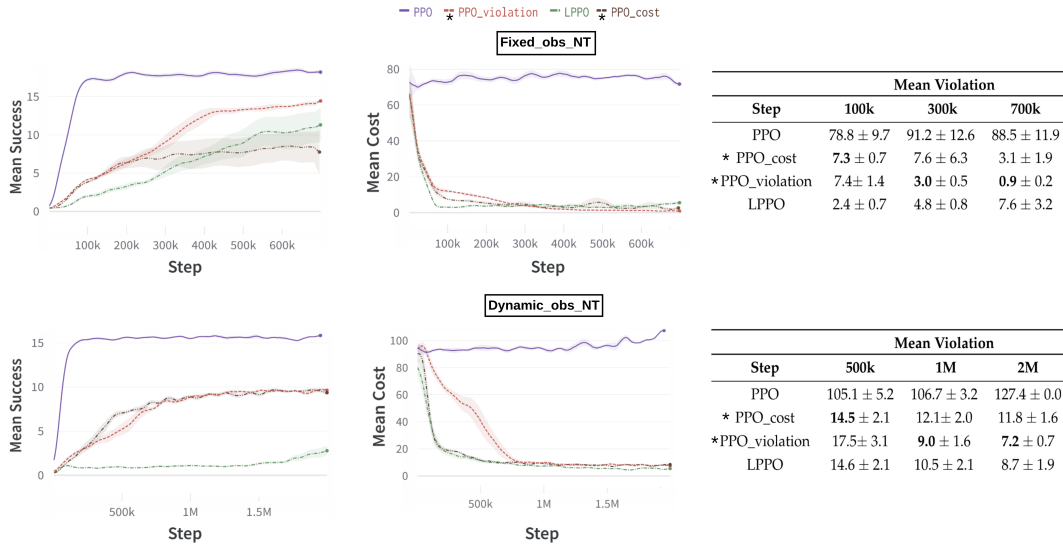


Figure 4: The two rows show average success, cost, and violation in the (NT) environments for PPO, the cost and violation penalty versions PPO_{cost, violation}, and LPPO. * indicates penalty-based algorithms.

Table 2: Average formal violation for the models at convergence in NT environment. Our violation-based penalty algorithm is the safest overall, as a lower violation value translates into fewer collisions.

	Fixed_obs_NT				Dynamic_obs_NT			
	PPO	*PPO_cost	*PPO_violation	LPPO	PPO	*PPO_cost	*PPO_violation	LPPO
p_{\uparrow}	0.25±0.07	0.27±0.15	0.23±0.09	0.19±0.0	0.24±0.0	0.32±0.23	0.22±0.15	0.08±0.07
p_{\leftarrow}	0.49±0.12	0.5±0.14	0.46±0.1	0.43±0.35	0.36±0.0	0.04±0.02	0.05±0.04	0.11±0.08
p_{\rightarrow}	0.5±0.16	0.44±0.33	0.42±0.14	0.55±0.12	0.4±0.0	0.08±0.08	0.04±0.04	0.18±0.16
Sum:	1.24	1.21	1.11	1.21	1.0	0.44	0.31	0.37

trade-off in complex scenarios [27, 37]. Moreover, PPO_violation achieves better or comparable successes and cost values over PPO_cost but significantly reduces the violations during the training. At convergence, PPO_violation shows a $\approx 2.2\%$ and $\approx 4.6\%$ improvement over the cost counterparts, corresponding to 1320 and 2760 fewer unsafe policy decisions. In general, non-terminal tasks allow experiencing more unsafe situations in a trajectory, making the performance advantage in terms of the safety of violation-based algorithms more evident.

In addition, we use FV [10] on the trained models of NT environments to provably guarantee the number of unsafe policy decisions over our safety properties. Table 2 shows the average violations of the main properties for the models at the convergence of each training seed, which confirms that policies trained with PPO_violation achieve lower violations, i.e., perform fewer collisions. The Appendix shows the same results for the T environments.

Evaluation Results. Table 3 reports the average success, cost, and violation for the best model at the convergence of each training seed in the evaluation task of Figure 3. The evaluation in a previously unseen scenario is used to test the generalization skills of the trained policies. In our experiments, the violation-based algorithm confirms superior navigation skills, achieving more success while

Table 3: Performance of the best-trained models in the testing environment *Evaluation_NT*.

	Mean Success	Mean Cost	Mean Violation
PPO	10.0 ± 0.4	52.0 ± 4.8	63.7 ± 11.3
PPO_cost	6.4 ± 1.8	25.5 ± 8.3	29.2 ± 9.5
PPO_violation	7.7 ± 0.2	17.9 ± 1.2	18.8 ± 1.7
LPPO	2.5 ± 0.3	18.2 ± 2.5	19.5 ± 0.9

being safer than LPPO and the cost counterpart.

Real-Robot Testing. The core motivation for introducing our Safe DRL environments is the lack of benchmarks that allow testing the policies directly on real robots. In contrast, Unity environments enable the transfer of policies trained in simulation on ROS-enabled platforms. We report a video with the real-world experiments of the policies trained in our environments here: shorturl.at/ijmFV, while Figure 5 shows our actual setup with the Turtlebot3.

4.4 Additional Experiments

We performed several additional experiments in the *Fixed_obs_NT* task to highlight the impact of the different components presented



Figure 5: Overview of real-world experiments.

in this paper using: (i) different sizes for the pre-conditions, (ii) the online property.⁵ (iii) We show the results of our violation-based penalty method in standard safe locomotion benchmarks to further confirm the performance improvement of the proposed approach.

Different Sizes for Pre-Conditions. Figure 6 shows the results of two PPO_violation versions: one that uses a constant ϵ for the size of the pre-conditions, and one that has different ϵ values for each input. The latter considers all the possible target positions as in previous experiments. As detailed in Section 4.2, using wider ranges for the inputs that shape environment configurations allows the violation to contain details about the unsafe behavior in general target initialization, resulting in higher performance. Crucially, the PPO_violation with a constant ϵ also returns better performance over PPO_cost, confirming the importance of having locally-aware information in the penalty value. This is particularly important as it may not be possible to shape detailed pre-conditions in setups where accurate task-level knowledge is lacking.

Online Property. As detailed in Section 3.1, we generate an additional property upon experiencing an unsafe state s because it is not uncommon to experience an unsafe transition not included in the pre-conditions (due to the limitation of hard-coding properties). Figure 8 shows the size of the set P' during the training for PPO_violation under two implementations. The first adds an online-generated property, and the second uses only the hard-coded properties. Results for the latter show an average size of $P' < 1$, confirming our hypothesis and the limitations of the properties' design of Section 3.2. In contrast, the generated property implementation ensures having at least one input-output mapping for each unsafe state. This allows the violation to correctly shape information about undesired situations, which biases the policy toward safer regions. Moreover, the growth in the size of P' over the training indicates that the policy experiences unsafe states in rare corner cases captured by the intersection of multiple properties (i.e., some complex situations require not choosing more than one action to avoid collisions).

Standard Safe DRL Tasks. We performed additional experiments in the standard Safe-DroneRun task to consider a different simulated robot, task, and safety specifications. Our goal is to confirm further our framework's results and the benefits of using the

violation as a penalty in a different domain known in the literature.⁶ In more detail, we consider the same hyperparameters of our navigation experiments. Moreover, given the challenges of hand-writing properties for the drone, we rely only on the online generated property, considering a fixed ϵ -area of size 0.05 where we want to avoid a similar action that led to the unsafe state, up to decimal precision. The following results consider the average reward, cost, and violation collected over ten runs with different random seeds.

Figure 7 shows the results Safe-DroneRun. These results confirm the behavior of previous experiments (we omitted the PPO results to maintain the plot scale for better visualization), where LPPO struggles to keep the cost threshold set to 20 and results in lower performance compared to the penalty-based approaches. In contrast, the violation-based PPO maintains the best trade-off between reward and cost.

5 RELATED WORK

Garcia and Fernández [12] presents an exhaustive taxonomy of the main families of approaches for Safe DRL, analyzing the pros and cons of each category. For example, model-based DRL approaches have been investigated in constrained and unconstrained settings [16, 52]. However, having access to or approximating a model is not always possible. Similarly, using barrier functions effectively fosters safety but requires an accurate system model [44].

In contrast, we focus on model-free learning. In this context, shielding approaches typically synthesize a shield (i.e., an automaton) to enforce safety specifications. However, this is usually unfeasible in complex setups due to the exponential growth in the size of the automaton. Hence, most DRL shielding approaches rely on simple grid-world domains [5, 11]. Although providing safety guarantees, it is unclear how to scale shielding approaches in complex, realistic applications. In contrast, constrained DRL has been used as a natural way to address Safe DRL [2, 8, 41, 49, 53]. In detail, CPO [2] is characterized by strongly constrained satisfaction at the expense of possibly infeasible updates that requires demanding second-order recovery steps. Similarly, PCPO [49] uses second-order derivatives and has mixed improvements over CPO [53]. Moreover, Lyapunov-based algorithms [8] combine a projection step with action-layer interventions to address safety. However, the cardinality of Lyapunov constraints equals the number of states, resulting in a significant implementation effort. Despite the variety of constrained literature, we compare with Lagrangian methods [41] as they reduce the complexity of prior approaches and show promising constraints satisfaction. However, constrained DRL has several drawbacks. For example, incorrect threshold tuning leads to algorithms being too permissive or, conversely, too restrictive [12]. Moreover, the guarantees of such approaches rely on strong assumptions that can not be satisfied in DRL, such as having an optimal policy. Constrained DRL is thus not devoid of short-term fatal behaviors as it can fail at satisfying the constraints [27]. Moreover, constraints naturally limit exploration, causing getting stuck in local optima or failing to learn desired behaviors properly [13, 24].

⁵A regularization term can modulate the importance of the penalty over the training. We perform additional experiments with it in the Appendix.

⁶We refer to the original works for more details about the environments and the shaped cost functions [2], github.com/SvenGronauer/Bullet-Safety-Gym




Figure 6: Average performance of penalty-based PPO using different pre-condition shapes: (i) fixed size (i.e., with a constant ϵ), (ii) and one that considers the whole domain for the target’s coordinates (i.e., different values for ϵ . For example, $x_{11,12} \in [-1, 1]$).




Figure 7: Average reward, cost, and violation in Safe-DroneRun for PPO, PPO_{cost, violation}, LPPO. * indicates penalty-based algorithms.




Figure 8: Mean size of P' over the training for PPO_violation with an online generated property (red), and with only the hard-coded properties (green).

We note that our Equation 5 falls under the category of reward engineering, which has been proved effective by several works [1, 12, 19, 40]. For example, IPO [21] uses a penalty function based on constraints, giving a zero penalty when constraints are satisfied and a negative infinity upon violation. However, tuning both the barrier parameter and the constraint threshold has sub-optimal solutions over Lagrangian methods [27]. Finally, Statistical Verification (SV) has been recently employed on learning systems [50] to deal with the computational demands of FV. In these approaches, desired specifications are defined as Signal Temporal Logic [23], which closely resembles the properties used by FV literature. However, a

distinctive feature of our method is the use of the violation value, which can not be directly computed by using SV.

6 DISCUSSION

We present an unconstrained DRL framework that leverages local violations of input-output conditions to foster safety. We discussed the limitations of using cost functions as in Safe DRL [12] presenting: (i) a sample-based approach to approximate a violation metric and use it as a penalty in DRL algorithms. Such a violation introduces task-level safety specifications into the optimization, addressing the cost’s lack of information. (ii) The influence of generating properties to cope with the limitations of hard-coded conditions. (iii) We argued the importance of developing real-world environments for broader applications of DRL in practical scenarios. To this end, we presented an initial suite of robotic navigation tasks that allow rapid testing on ROS-based robots.

This work paves the way for several research directions, including extending our task suite to create a general safe DRL benchmark. Such extension is possible due to the rapid prototyping benefits of Unity [15]. Moreover, studying the effects of different time horizons would be interesting to separate the importance given to rewards and safety values. It would also be interesting to design a shield to avoid unsafe behaviors at deployment, leveraging the information from FV. Finally, our insights on the violation could be used to model desired behaviors in single and multi-agent applications to improve performance and sample efficiency.

7 ACKNOWLEDGEMENTS

This work was partially funded by the Army Research Office under award number W911NF20-1-0265.

REFERENCES

- [1] David Abel, Will Dabney, Anna Harutyunyan, Mark K. Ho, Michael L. Littman, Doina Precup, and Satinder Singh. 2021. On the Expressivity of Markov Reward. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- [2] Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. 2017. Constrained Policy Optimization. In *International Conference on Machine Learning (ICML)*.
- [3] Eitan Altman. 1999. Constrained Markov Decision Processes. In *CRC Press*.
- [4] Guy Amir, Davide Corsi, Raz Yershalmi, Luca Marzari, David Harel, Alessandro Farinelli, and Guy Katz. 2022. Verifying Learning-Based Robotic Navigation Systems. *arXiv* (2022).
- [5] Steven Carr, Nils Jansen, Sebastian Junges, and Ufuk Topcu. 2022. Safe Reinforcement Learning via Shielding for POMDPs. In *arXiv*.
- [6] Hao-Tien Lewis Chiang, Aleksandra Faust, Marek Fiser, and Anthony Francis. 2019. Learning Navigation Behaviors End to End. *RA-L* (2019).
- [7] Yinlam Chow, Ofir Nachum, Edgar Duenez-Guzman, and Mohammad Ghavamzadeh. 2018. A Lyapunov-based Approach to Safe Reinforcement Learning. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- [8] Yinlam Chow, Ofir Nachum, Aleksandra Faust, Mohammad Ghavamzadeh, and Edgar A. Duenez-Guzman. 2019. Lyapunov-based Safe Policy Optimization for Continuous Control. In *International Conference on Machine Learning (ICML)*.
- [9] Cédric Colas, Olivier Sigaud, and Pierre-Yves Oudeyer. 2019. A Hitchhiker’s Guide to Statistical Comparisons of Reinforcement Learning Algorithms. In *International Conference on Learning Representations (ICLR) Workshop on Reproducibility*.
- [10] Davide Corsi, Enrico Marchesini, and Alessandro Farinelli. 2021. Formal Verification of Neural Networks for Safety-Critical Tasks in Deep Reinforcement Learning. In *Conference on Uncertainty in Artificial Intelligence (UAI)*.
- [11] Ingy Elsayed-Aly, Suda Bharadwaj, Christopher Amato, Rüdiger Ehlers, Ufuk Topcu, and Lu Feng. 2021. Safe Multi-Agent Reinforcement Learning via Shielding. In *AAMAS*.
- [12] Javier Garcia and Fernando Fernández. 2015. A comprehensive survey on safe reinforcement learning. In *Journal of Machine Learning Research (JMLR)*.
- [13] Zhang-Wei Hong, Tzu-Yun Shann, Shih-Yang Su, Yi-Hsiang Chang, and Chun-Yi Lee. 2018. Diversity-Driven Exploration Strategy for Deep Reinforcement Learning. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- [14] Andrew Ilyas, Logan Engstrom, Shibani Santurkar, Dimitris Tsipras, Firdaus Janoos, Larry Rudolph, and Aleksander Madry. 2020. A Closer Look at Deep Policy Gradients. In *International Conference on Learning Representations (ICLR)*.
- [15] Arthur Juliani, Vincent-Pierre Berges, Esh Vckay, Yuan Gao, Hunter Henry, Marwan Mattar, and Danny Lange. 2018. Unity: A Platform for Intelligent Agents. In *arXiv*.
- [16] Parv Kapoor, Anand Balakrishnan, and Jyotirmoy V. Deshmukh. 2020. Model-based Reinforcement Learning from Signal Temporal Logic Specifications. In *arXiv*.
- [17] Guy Katz, Clark Barrett, David L Dill, Kyle Julian, and Mykel J Kochenderfer. 2017. Reluplex: An efficient SMT solver for verifying deep neural networks. In *International conference on computer aided verification*. Springer, 97–117.
- [18] Guy Katz, Derek A Huang, Duligur Ibeling, Kyle Julian, Christopher Lazarus, Rachel Lim, Parth Shah, Shantanu Thakoor, Haoze Wu, Aleksandar Zeljic, et al. 2019. The marabou framework for verification and analysis of deep neural networks. In *International Conference on Computer Aided Verification*.
- [19] Zachary C. Lipton, Jianfeng Gao, Lihong Li, Jianshu Chen, and Li Deng. 2016. Combating Reinforcement Learning’s Sisyphian Curse with Intrinsic Fear. In *arXiv*.
- [20] Changliu Liu, Tomer Arnon, Christopher Lazarus, Christopher Strong, Clark Barrett, and Mykel J. Kochenderfer. 2021. Algorithms for Verifying Deep Neural Networks. *Foundations and Trends® in Optimization* (2021).
- [21] Yongshuai Liu, Jiaxin Ding, and Xin Liu. 2020. IPO: Interior-point Policy Optimization under Constraints. In *AAAI Conference on Artificial Intelligence*.
- [22] Yongshuai Liu, Avishai Halev, and Xin Liu. 2021. Policy Learning with Constraints in Model-free Reinforcement Learning: A Survey. In *International Joint Conference on Artificial Intelligence (IJCAI)*. 4508–4515.
- [23] Oded Maler and D. Niković. 2004. Monitoring Temporal Properties of Continuous Signals. In *FORMATS/FRITTT*.
- [24] Enrico Marchesini and Christopher Amato. 2022. Safety-Informed Mutations for Evolutionary Deep Reinforcement Learning. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. 1966–1970.
- [25] Enrico Marchesini and Christopher Amato. 2023. Improving Deep Policy Gradients with Value Function Search. In *International Conference on Learning Representations (ICLR)*.
- [26] Enrico Marchesini, Davide Corsi, and Alessandro Farinelli. 2021. Benchmarking Safe Deep Reinforcement Learning in Aquatic Navigation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 5590–5595. <https://doi.org/10.1109/IROS51168.2021.9635925>
- [27] E. Marchesini, D. Corsi, and A. Farinelli. 2022. Exploring Safer Behaviors for Deep Reinforcement Learning. In *AAAI Conference on Artificial Intelligence Conference on Artificial Intelligence*.
- [28] Enrico Marchesini and Alessandro Farinelli. 2020. Discrete deep reinforcement learning for mapless navigation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 10688–10694.
- [29] Enrico Marchesini and Alessandro Farinelli. 2021. Centralizing State-Values in Dueling Networks for Multi-Robot Reinforcement Learning Mapless Navigation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 4583–4588. <https://doi.org/10.1109/IROS51168.2021.9636349>
- [30] Enrico Marchesini and Alessandro Farinelli. 2022. Enhancing Deep Reinforcement Learning Approaches for Multi-Robot Navigation via Single-Robot Evolutionary Policy Search. In *International Conference on Robotics and Automation (ICRA)*. 5525–5531. <https://doi.org/10.1109/ICRA46639.2022.9812341>
- [31] Luca Marzari, Davide Corsi, Ferdinando Cicalese, and Alessandro Farinelli. 2023. The #DNN-Verification problem: Counting Unsafe Inputs for Deep Neural Networks. *ArXiv abs/2301.07068* (2023).
- [32] Luca Marzari, Davide Corsi, Enrico Marchesini, and Alessandro Farinelli. 2022. Curriculum learning for safe mapless navigation. In *Proceedings of the 37th ACM/SIGAPP Symposium on Applied Computing*. 766–769.
- [33] Luca Marzari, Enrico Marchesini, and Alessandro Farinelli. 2023. Online Safety Property Collection and Refinement for Safe Deep Reinforcement Learning in Mapless Navigation. In *International Conference on Robotics and Automation (ICRA)*.
- [34] Andrew Y. Ng, Daishi Harada, and Stuart Russell. 1999. Policy invariance under reward transformations: Theory and application to reward shaping. In *ICML*.
- [35] J. Nocedal and S. Wright. 2006. *Numerical Optimization* (2 ed.). Springer.
- [36] Ameya Pore, Davide Corsi, Enrico Marchesini, Diego Dall’Alba, Alicia Casals, Alessandro Farinelli, and Paolo Fiorini. 2021. Safe Reinforcement Learning using Formal Verification for Tissue Retraction in Autonomous Robotic-Assisted Surgery. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 4025–4031. <https://doi.org/10.1109/IROS51168.2021.9636175>
- [37] Alex Ray, Joshua Achiam, and Dario Amodei. 2019. Benchmarking Safe Exploration in Deep Reinforcement Learning. In *arXiv*.
- [38] Julien Roy, Roger Girgis, Joshua Romoff, Pierre-Luc Bacon, and Christopher Pal. 2022. Direct Behavior Specification via Constrained Reinforcement Learning. In *AAAI Conference on Artificial Intelligence*.
- [39] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv* (2017).
- [40] David Silver, Satinder Singh, Doina Precup, and Richard S. Sutton. 2021. Reward is enough. *Artificial Intelligence* 299 (2021), 103535.
- [41] Adam Stooke, Joshua Achiam, and Pieter Abbeel. 2020. Responsive Safety in Reinforcement Learning by PID Lagrangian Methods. In *International Conference on Machine Learning (ICML)*.
- [42] L. Tai, G. Paolo, and M. Liu. 2017. Virtual-to-real deep reinforcement learning: Continuous control of mobile robots for mapless navigation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- [43] Chen Tessler, Daniel J. Mankowitz, and Shie Mannor. 2019. Reward Constrained Policy Optimization. In *International Conference on Learning Representations (ICLR)*.
- [44] Cumhur Erkan Tuncali, James Kapinski, Hisahiro Ito, and Jyotirmoy V. Deshmukh. 2018. Reasoning about Safety of Learning-Enabled Components in Autonomous Cyber-physical Systems. In *Design Automation Conference (DAC)*.
- [45] Shiqi Wang, Kexin Pei, Justin Whitehouse, Junfeng Yang, and Suman Jana. 2018. Efficient Formal Safety Analysis of Neural Networks. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- [46] Shiqi Wang, Kexin Pei, Justin Whitehouse, Junfeng Yang, and Suman Jana. 2018. Formal security analysis of neural networks using symbolic intervals. In *USENIX Security Symposium*.
- [47] Ziyu Wang, Tom Schaul, Matteo Hessel, Hado Hasselt, Marc Lanctot, and Nando Freitas. 2016. Dueling network architectures for deep reinforcement learning. In *International conference on machine learning (ICML)*. 1995–2003.
- [48] Kaidi Xu, Huan Zhang, Shiqi Wang, Yihan Wang, Suman Jana, Xue Lin, and Cho-Jui Hsieh. 2021. Fast and Complete: Enabling Complete Neural Network Verification with Rapid and Massively Parallel Incomplete Verifiers. In *International Conference on Learning Representations (ICLR)*.
- [49] Tsung-Yen Yang, Justinian Rosca, Karthik Narasimhan, and Peter J. Ramadge. 2020. Projection-Based Constrained Policy Optimization. In *International Conference on Learning Representations (ICLR)*.
- [50] Mojtaba Zarei, Yu Wang, and Miroslav Pajic. 2020. Statistical Verification of Learning-Based Cyber-Physical Systems. In *International Conference on Hybrid Systems: Computation and Control*.
- [51] J. Zhang, J. T. Springenberg, J. Boedecker, and W. Burgard. 2017. Deep reinforcement learning with successor features for navigation across similar environments. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.

- [52] Xinglong Zhang, Yaoqian Peng, Biao Luo, Wei Pan, Xin Xu, and Haibin Xie. 2021. Model-Based Safe Reinforcement Learning with Time-Varying State and Control Constraints: An Application to Intelligent Vehicles. In *arXiv*.
- [53] Yiming Zhang, Quan Vuong, and Keith W. Ross. 2020. First Order Constrained Optimization in Policy Space. In *Conference on Neural Information Processing Systems (NeurIPS)*.