

Conformal Reachability for Safe Control in Unknown Environments

Xinhang Ma

Washington University in St. Louis
St. Louis, USA
m.owen@wustl.edu

Yiannis Kantaros

Washington University in St. Louis
St. Louis, USA
ioannisk@wustl.edu

Junlin Wu

Washington University in St. Louis
St. Louis, USA
junlin.wu@wustl.edu

Yevgeniy Vorobeychik

Washington University in St. Louis
St. Louis, USA
yvorobeychik@wustl.edu

ABSTRACT

Designing provably safe control is a core problem in trustworthy autonomy. However, most prior work in this regard assumes either that the system dynamics are known or deterministic, or that the state and action space are finite, significantly limiting application scope. We address this limitation by developing a probabilistic verification framework for unknown dynamical systems which combines conformal prediction with reachability analysis. In particular, we use conformal prediction to obtain valid uncertainty intervals for the unknown dynamics at each time step, with reachability then verifying whether safety is maintained within the conformal uncertainty bounds. Next, we develop an algorithmic approach for training control policies that optimize nominal reward while also maximizing the planning horizon with sound probabilistic safety guarantees. We evaluate the proposed approach in seven safe control settings spanning four domains—cartpole, lane following, drone control, and safe navigation—for both affine and nonlinear safety specifications. Our experiments show that the policies we learn achieve the strongest provable safety guarantees while still maintaining high average reward. Full paper is available at <https://arxiv.org/abs/2602.03799>.

KEYWORDS

Safe control; conformal reachability; verified safety

ACM Reference Format:

Xinhang Ma, Junlin Wu, Yiannis Kantaros, and Yevgeniy Vorobeychik. 2026. Conformal Reachability for Safe Control in Unknown Environments. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 10 pages. <https://doi.org/10.65109/SEHD8631>

1 INTRODUCTION

Trustworthy control of autonomous systems in dynamic environments is a fundamental problem in a broad range of domains including autonomous driving, drone navigation, and robotic manipulation, among others. While reinforcement learning techniques

have made significant strides in advancing autonomy capabilities in these domains [6, 19, 27, 47, 56], their trustworthiness is opportunistic, insofar as they tend to combine performance objectives and safety considerations into a single reward function, and achieve high reward accounting for safety only on average. In safety critical environments, such as self-driving, more rigorous safety guarantees are commonly expected.

An extensive literature has emerged to address this issue, addressing both the problem of verifying safety [3, 4, 22, 23, 42, 53, 54], as well as synthesizing safe control policies (using both formal methods and learning techniques) [17, 18, 32, 46, 51]. However, verification techniques typically assume known or deterministic dynamics, and many suffer from severe scalability limitations when dynamics are nonlinear, particularly for complex (e.g., neural network) control policies. Similarly, synthesis techniques, including those using safe reinforcement learning methods, either make very strong assumptions on the dynamical system (for example, assuming that it is known and deterministic) [46], assume a finite state or action space [10, 23], or achieve safety only approximately [17, 30, 49, 51, 52] (for example, by relying on learned estimates of the probability that safety is violated [17]), on average (as in the constrained Markov decision process (CMDP) framework) [2, 41, 50], or in the limit [17, 51]. Moreover, while a number of approaches have been proposed to deal with unknown dynamical systems, most still rely on restrictive assumptions to obtain provable safety guarantees, such as assuming that dynamics are deterministic [23, 44, 58], constraining approximation models to specific forms such as Gaussian process regression [23], or requiring Lipschitz continuity [58].

We present a novel approach for learning policies in unknown dynamical systems that approximately maximize rewards with rigorous probabilistic safety guarantees. Our approach combines model-based RL—we learn a Markovian approximation of the dynamics—with conformal safety analysis (CSA) to obtain safety guarantees, and model-free RL to optimize cumulative rewards. By jointly training the dynamics model and policy using a loss that incentivizes safety improvements over a target decision horizon, our approach facilitates a finer-granularity tradeoff between the nature of (probabilistic) safety guarantees and rewards than prior art. Moreover, the use of conformal prediction (CP) as a key building block allows us to make no assumptions on the nature of the underlying dynamics (which we allow, for example, to be stochastic) or the



This work is licensed under a Creative Commons Attribution International 4.0 License.

model architecture used for learning it. Finally, our method provides sound probabilistic safety guarantees by combining three technical elements: valid uncertainty bounds for dynamics prediction, finite-horizon reachability analysis, and probabilistic tail bounds with respect to the unknown initial state distribution. Our CSA approach builds on recent literature combining conformal prediction with reachability analysis [8, 20, 21, 29, 33], but unlike most prior efforts we make no smoothness assumptions on the underlying dynamics, allow verification for an arbitrary horizon post surrogate model training, and develop novel training methods for both surrogate dynamics and policy that obtain strong safety guarantees.

We evaluate the proposed approach experimentally on seven safe control experiment settings, comparing it to state of the art safe RL baselines. Our results show that our approach achieves state of the art safety guarantees while maintaining high reward.

In summary, our main contributions are as follows:

- (1) *Conformal Safety Analysis (CSA)* for safe reinforcement learning in unknown dynamical systems. In particular, our CSA framework combines conformal prediction with reachability analysis for learning-based controllers in systems with *unknown dynamics*, which we augment with distribution-free tail bounds for an unknown initial state distribution to obtain rigorous probabilistic safety guarantees.
- (2) A novel framework for learning provably safe policies that (approximately) maximize cumulative reward subject to safety constraints. Our approach integrates model-free RL and model-based reachability estimation with quantified uncertainties to enforce hard safety constraints.
- (3) Comprehensive experiments across seven environments, including both linear and nonlinear safety specifications, showing that our approach achieves high rewards while providing verifiable probabilistic safety guarantees.

2 RELATED WORK

Ensuring safety in dynamical systems has been classically modeled with Constrained Markov Decision Processes (CMDPs) [2], which seek to maximize expected rewards while constraining cumulative costs below a predefined threshold. Common approaches for solving CMDPs include Lagrangian and penalty methods, as well as constrained policy optimization [1, 24, 32, 39, 40, 51]. While these methods are effective at balancing reward and cost, they primarily enforce empirical safety and lack formal safety guarantees.

To address this limitation, extensive literature has also emerged for obtaining provably safe or stable control, particularly for policies that are represented as neural networks. Many of these rely on control barrier functions, which are themselves difficult to synthesize or verify at scale [9, 12, 45, 57]. Others proposed reinforcement learning-based training algorithms to obtain policies with provable safety and stability guarantees [46, 58], but these scale poorly with the number of states, and either assume known or deterministic dynamics. Data-driven methods such as [13–15, 35] use simulation-based reachability or scenario optimization, but are often conservative and require large sample sizes to achieve high confidence. Our approach, in contrast, uses conformal prediction to provide statistically calibrated, distribution-free uncertainty quantification.

This enables scalable and model-agnostic verifications under any model errors.

A number of recent approaches combine reachability analysis with conformal prediction, as we do [20, 29, 33], with several also addressing the issue of unknown dynamics [8, 21]. However, these either assume smooth dynamics or a fixed decision horizon K . In contrast, our approach trains Markovian surrogate models that can subsequently be applied for any K , and we make no assumptions on the underlying dynamics. Moreover, unlike prior work, we train verifiable combinations of surrogate dynamics and policies.

Another class of approaches for verifying safety makes use of Hamilton-Jacobi (HJ) reachability techniques. These compute exact reachable/avoidable sets via PDE solutions, but do not scale well and typically assume knowledge of system dynamics [16, 28, 37, 52]. Yu et al. [51] proposed an approach that guarantees safety when agent is in the feasible set, while Ganai et al. [17] extended it to general stochastic settings. However, our framework is fundamentally different: instead of solving the computationally expensive PDEs, we adopt a discrete-time finite-step forward reachability formulation that scales better with system complexity. Moreover, our use of conformal prediction enables us to apply our approach without making any assumptions about unknown dynamics.

3 PRELIMINARIES

3.1 Deep Reinforcement Learning

As our approach includes features of model-free reinforcement learning (RL), we provide a brief overview of the main concepts. Generally, such approaches train a critic V_{θ_c} (e.g., value or action-value function) and an actor (policy) π_{θ_a} , in some cases including also an entropy term \mathcal{H} that encourages exploration, with the general loss function of the form

$$\mathcal{J}(\theta_c, \theta_a) = \mathbb{E}_{s,a,r,s'} [\mathcal{L}_c(\theta_c) + \alpha_1 \mathcal{L}_a(\theta_a) + \alpha_2 \mathcal{H}[\pi_{\theta_a}]]$$

for observed sequences of states and actions s, a, s' and rewards r . In actor-critic methods in particular (such as PPO [36]), a common approach is to define an advantage function $A(s, a) = r + \gamma V_{\theta_c}(s') - V_{\theta_c}(s)$, where γ is the discount factor, and r and s' stem from taking action a , with $\mathcal{L}_c(\theta_c) = A(s, a)^2$ and $\mathcal{L}_a(\theta_a)$ is an algorithm-specific function of $A(s, a)$.

3.2 Conformal Prediction

Conformal prediction (CP) is a distribution-free and model-agnostic statistical uncertainty quantification technique. As a result, it has come to be particularly useful in supervised learning, where it can provide valid statistical guarantees for complex models such as deep neural networks. Of particular interest in our context is *regression learning*, which we use to learn a model of system dynamics in a state space that is embedded in a Euclidean space, and for which CP provides valid prediction intervals [43]. Here we provide an overview of conformal prediction in the form of split conformal prediction, which is most widely used. For a more comprehensive review and more examples, we refer the readers to [5].

Given a general input x and output y , first define the score function $s(x, y) \in \mathbb{R}$ where larger scores reflect larger uncertainties (larger errors between x and y). Let $\alpha \in (0, 1)$ be a user-specified miscoverage level that represents the desired probability of error in

the uncertainty estimates. We compute \hat{q} as the $\frac{\lceil (n+1)(1-\alpha) \rceil}{n}$ th quantile of the scores of a calibration set $\{s_i = s(x_i, y_i)\}$ for $1 \leq i \leq n$. Finally, for a new input x_{test} , we use \hat{q} to provide the uncertainty interval for the regression model \hat{f} as

$$\mathcal{C}(x_{\text{test}}) = [\hat{f}(x_{\text{test}}) - \hat{q}, \hat{f}(x_{\text{test}}) + \hat{q}] \quad (1)$$

Crucially, the interval $\mathcal{C}(x_{\text{test}})$ achieves a $\mathbf{P}(y_{\text{test}} \in \mathcal{C}(x_{\text{test}})) \geq 1 - \alpha$ coverage guarantee [43].

Because standard split conformal prediction assumes exchangeable nonconformity scores s_0, \dots, s_n , it warrants adaptation when applied in time-series prediction. A naive example is to use a union bound over the prediction horizon to connect per-step and full-trajectory bounds. However, this can result in bounds that are overly conservative.

To obtain tighter conformal trajectory bounds, we make use of the approach by Cleaveland et al. [11], which learns a trajectory-level nonconformity score that can directly be optimized for efficiency. Specifically, one can define

$$s = \max(w_1 s(x_1, y_1), \dots, w_T s(x_T, y_T))$$

where $w_1, \dots, w_T \geq 0$ are parameters. Cleaveland et al formulate the problem of finding the parameters w as the following optimization problem:

$$\begin{aligned} \min_{w \geq 0} \quad & \text{Quantile}(\{s^{(1)}, \dots, s^{(n)}\}, 1 - \alpha) \\ \text{s.t.} \quad & s^{(i)} = \max(w_1 s_1^{(i)}, \dots, w_T s_T^{(i)}) \quad \forall i; \quad \sum_{j=1}^T w_j = 1, \end{aligned}$$

where $s^{(i)}$ is the score for the i -th trajectory in a held-out dataset. The details of solving this optimization problem efficiently can be found in Section 4 of [11]. Importantly, conformal region computed using this nonconformity score with optimized w parameters maintains the desired coverage guarantee (Theorem 3 in [11]).

4 MODEL

We consider a safety-constrained Markov decision process (MDP) (SMDP), defined by a tuple $\mathcal{M} := \langle \mathcal{S}, \mathcal{A}, P, \rho, r, h, K, \gamma \rangle$, where $\mathcal{S} \subseteq \mathbb{R}^n$ is the state space, \mathcal{A} the action space, $P : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{P}(\mathcal{S})$ the transition distribution (i.e. the environment dynamics) which maps current state and action to a distribution \mathcal{P} over states, ρ the initial state distribution, $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ the reward function assigning a real value reward to each state-action pair, $h : \mathcal{S} \rightarrow \mathbb{R}_+^m$ is the cost function assigning a real vector to a state representing its associated m -dimensional safety value (e.g., corresponding to m safety constraints, such as m distinct obstacles to avoid), K the decision horizon, and $\gamma \in (0, 1)$ is the discount factor. Thus, $h(s) \leq 0$ pointwise implies that the state $s \in \mathcal{S}$ is *safe* (safety cost is non-negative) while $h(s) > 0$ means that it is *unsafe*. Let $\mathcal{S}_s \subseteq \mathcal{S} := \{s \in \mathcal{S} : h(s) \leq 0\}$ denote the set of safe states, while $\mathcal{S}_u := \{s \in \mathcal{S} : h(s) > 0\}$ is the set of unsafe states. Correspondingly, the SMDP imposes a safety constraint that the system dynamics remain safe until the safety horizon K . Formally, for any trajectory of states $\tau = \{s_0, \dots, s_K\}$, we wish to ensure that $s_t \in \mathcal{S}_s$ for all $1 \leq t \leq K$. We assume that the support of ρ is over \mathcal{S}_s (that is, all initial states are safe). If this constraint holds, we say that the system is *safe*. We remark that the safety constraints in SMDPs are distinct from the typical nature of constraints in *constrained MDPS (CMDPs)*,

which accrue (discounted) costs additively over time analogously to rewards in the objective. Many algorithmic frameworks for CMDPs leverage this structure (for example, using Lagrangian methods to transform it into an MDP). In contrast, the safety constraints in SMDPs are non-additive, and such algorithmic approaches cannot be leveraged without the use of additive proxies.

Let $\pi : \mathcal{S} \rightarrow \mathcal{A}$ be a deterministic Markov stationary policy, and Π the set of all such policies. As is conventional in MDPs, we define the state-conditional *policy value function* $V^\pi(s)$ as the discounted sum of rewards, i.e.,

$$V^\pi(s) = \mathbb{E} \left[\sum_{t=0}^K \gamma^t r(s_t, a_t) \mid s_0 = s, s_t \sim P(s_{t-1}, a_{t-1}), a_t = \pi(s_t) \right].$$

Informally, our goal is to find a policy π that maximizes $V^\pi(s)$ while ensuring that the system is safe. However, we suppose that we do not know the system dynamics P or reward r . While we can use conventional reinforcement learning approaches to learn a near-optimal policy in an MDP without safety considerations, guaranteeing perfect safety is impractical. Instead, we aim to achieve safety *with high probability*. Define

$$\Sigma^\pi = \Pr [s_t \in \mathcal{S}_s, \forall t \in \{0, 1, \dots, K\}]$$

where $s_0 \sim \rho, s_t \sim P(s_{t-1}, a_{t-1}), a_t = \pi(s_t)$, that is, Σ^π is the probability that the entire trajectory of horizon K generated by a policy π is safe, accounting for all relevant uncertainty (which we will make more precise in Section 5.2).

Our goal is to solve the following constrained optimization problem:

$$\max_{\pi \in \Pi} \mathbb{E}_{s \sim \rho} [V^\pi(s)] \quad \text{s.t.} : \Sigma^\pi \geq 1 - \zeta \quad (2)$$

for a target safety tolerance level ζ . Crucially, while we can accept approximately optimal policies π in terms of the objective, we wish to obtain *verified* safety at the desired tolerance, since these often capture considerations that have high stakes. For example, rewards would naturally capture efficiency, such as getting to the destination quickly, whereas safety constraints focus on avoiding traffic accidents that may potentially lead to fatalities.

Solving Problem (2) entails two considerations: 1) obtaining sound probabilistic safety guarantees for a given deterministic policy π as expressed in the safety constraint (which we call *verification*) and 2) synthesizing (or learning) a policy that is approximately optimal while also satisfying the safety constraint for the target decision horizon K and tolerance ζ . Moreover, we must accomplish both goals *without knowing the true dynamics P and reward r* . We tackle these problems next.

5 SOLUTION APPROACH

In this section, we introduce our approach, *ReCORS (Reachability with Conformal uncertainty sets for Reinforcement learning with Safety guarantees)*, which involves three main components:

- (1) Approximating the true system dynamics P with a neural network \hat{f}_β that is optimized for accuracy in safety-critical state dimensions.
- (2) A novel integration of conformal prediction with reachability analysis performed using the learned dynamics model \hat{f}_β . The resulting *conformal safety analysis (CSA)* enables verification of whether the system is likely to remain in safe

set \mathcal{S}_s over horizon K with a pre-defined confidence level $(1 - \alpha)$.

- (3) A fully differentiable pipeline allowing end-to-end gradient-based learning of a parametric control policy π_β : we use a neural network η_ω to approximate the conformal prediction intervals, allowing the safety objective (derived from \hat{f}_β and η_ω) to be directly integrated as a differentiable objective term in any gradient-based policy optimization algorithm.

Notably, our approach is a novel blend of model-based and model-free reinforcement learning. In the model-based component, we learn the model of the dynamics (but not the rewards), which is pivotal for training with probabilistic safety guarantees and associated verification. Our learning pipeline then leverages a combination of model-free RL (e.g., PPO) to maximize rewards with the model-based component focusing on optimizing for provable safety. We detail the components of our approach next.

5.1 Learning System Dynamics with Safety Specifications

A key component of our approach is training a parametric discrete-time dynamics model $\hat{f}_\beta(\hat{s}, a)$ with parameters β , which *deterministically* maps state \hat{s} and action a to predicted next state \hat{s}' . Thus, our *learned* approximate dynamical system model has the form

$$\hat{s}_{t+1} = \hat{f}_\beta(\hat{s}_t, a_t), \quad (3)$$

where \hat{s}_t and a_t are the (previously predicted) state and action taken at time step t , respectively, and \hat{s}_{t+1} is the predicted next state. Notably, even though the underlying dynamics may be stochastic, while \hat{f} is deterministic, our use of conformal prediction will account for the concomitant prediction uncertainty (essentially combining both the prediction error and any underlying uncertainty of the dynamics itself); we discuss this further below.

To train the dynamical system model, we collect a set of ground-truth trajectories $\{\tilde{\tau}_l\}$ where $\tilde{\tau}_l = (s_{l0}, a_{l0}, s_{l1}, \dots, s_{lK-1}, a_{lK-1}, s_{lK})$ by sampling initial states s_0 i.i.d. from ρ . This induces a dataset of N tuples $\mathcal{D} = \{(s_i, a_i, s'_i)\}_{i=1}^N$ that we can use for training. Notably, we train the dynamics \hat{f} in two phases. First, we pre-train \hat{f} by generating the trajectory data by sampling actions a_t uniformly at random. Subsequently, we keep fine-tuning the dynamics model as part of the proposed RL approach (see below), where in each training epoch we resample the trajectory data using the latest policy π_β obtained during the RL training.

The loss function we use for training \hat{f}_β is a variant of ℓ_2 loss. Since a major concern in our context is safety, we modify the standard ℓ_2 loss to account for safety considerations by adding weights that amplify the importance of safety-critical dimensions of the state space. Specifically, let $\phi_j(h(s)) > 0$ be a positive scalar that measures the safety criticality of dimension j with respect to the safety function $h(s)$. In particular, suppose without loss of generality that the i th safety constraint has the form $h_i(s) = \bar{h}_i(s) + b_i$ for some scalar b_i . We propose the following measure of safety criticality of dimension j :

$$\phi_j(h(s)) = \prod_{i=1}^m \left(\frac{|\partial \bar{h}_i(s) / \partial s_j|}{|b_i|} + 1 \right),$$

where $|\cdot|$ is an absolute value. For example, if $h(s) = As - b$, $\partial \bar{h}_i(s) / \partial s_j = A_{ij}$. Intuitively, ϕ_j introduces a higher weight for dimensions which are more likely to induce safety violations.

Next, define the safety weight of dimension j as $w_j = \phi_j(h(s))$. Our training loss then becomes

$$\mathcal{L}_{\text{dyn}}(\beta) = \frac{1}{N} \sum_{i=1}^N \|\hat{f}_\beta(s_i, a_i) - s'_i\|_W, \quad (4)$$

where W is a diagonal matrix with w_{js} along the diagonal, and $\|x\|_W = x^T W x$.

The expected effect of adding such safety weights as part of the loss function used to train the dynamics is two-fold. On the one hand, it will incentivize training to focus on making better predictions along the dimensions that are more safety-critical. On the other hand, this may increase overall (unweighted) prediction error, which would in turn lead to reduced reward. This tradeoff is deliberate: while we aim to maximize total reward, we prioritize *verified safety*. Notably, our approach is heuristic, to be sure, but as our ablations in Section 6.3 demonstrate, it has the intended effect of improving verified safety.

5.2 Conformal Safety Analysis

Consider the approximate dynamical system model in Equation (3) with a known initial state s_0 . We can represent this as a collection of K predictions, each corresponding to a model that predicts the t th state in the sequence for $t \in [1, K]$ given an initial state s_0 and a policy π_θ , that is,

$$\hat{s}_t = \hat{f}_\beta^{(t)}(s_0; \pi_\theta) \quad (5)$$

for $t > 0$, where $\hat{s}_1 = \hat{f}_\beta^{(1)}(s_0; \pi_\theta) = \hat{f}_\beta(s_0, \pi_\theta(s_0))$, $\hat{s}_2 = \hat{f}_\beta^{(2)}(s_0; \pi_\theta) = \hat{f}_\beta(\hat{s}_1, \pi_\theta(\hat{s}_1))$, and so on, with $a_t = \pi_\theta(\hat{s}_t)$. In other words, given a policy π_θ , the predicted dynamical system model \hat{f}_β effectively gives rise to K separate state predictions, each with s_0 as input. Moreover, if s_0 are sampled i.i.d. from the initial state distribution ρ , we can apply conformal prediction bounds for each step t to obtain

$$\mathbb{P}[\|s_t - \hat{f}_\beta^{(t)}(s_0, \pi_\theta)\| \leq \eta_t] \geq 1 - \alpha_t, \quad (6)$$

where s_t is a true state generated by the actual system dynamics after t steps starting at s_0 , α_t the error tolerance rate, and η_t the conformity score obtained using a calibration dataset to ensure the validity of (6).

However, obtaining conformal bounds for each t is insufficient, since our goal is to obtain such bounds over K -step trajectories, and predictions can be arbitrarily correlated across time steps. Specifically, let $S_t(s_0; \pi_\theta) = \{s \mid \|s_t - \hat{f}_\beta^{(t)}(s_0, \pi_\theta)\| \leq \eta_t\}$ denote the set of states that can be reached after exactly t steps given the conformity score η_t if we start at s_0 and follow the policy π_θ . Let $S(s_0; \pi_\theta) = \cup_{t \in [1, K]} S_t(s_0; \pi_\theta)$. We can then observe that if $S(s_0; \pi_\theta) \cap \mathcal{S}_u = \emptyset$, the policy π_θ is safe, modulo the predictive uncertainty of \hat{f} . Our goal is to obtain this claim rigorously for a given policy π_θ , with a target probability $1 - \zeta$. One way to do this is to simply apply a union bound to the step-wise conformal bounds α_t over all steps t . However, this is often too loose. As an alternative, we additionally consider a more efficient time-series conformal bound proposed by Cleaveland et al. [11]. Armed with a

time-series-level conformal bound, we obtain the following result, which borrows the probabilistic reasoning with respect to the initial state distribution from semi-probabilistic verification [31] and gives us a sound probabilistic safety bound, accounting for uncertainty about (a) initial state, (b) imperfect predictions by \hat{f}_β , and (c) samples of finite datasets of s_0 used to obtain empirical safety.

THEOREM 1. *Let $\{s_0^i\}_{i=1}^N$ be i.i.d. samples from ρ , independent of the calibration data used to construct the conformal bounds. Suppose the conformal procedure provides a trajectory-level coverage guarantee that for every fixed initial state s_0 , $\Pr(S(s_0; \pi_\theta) \cap \mathcal{S}_u = \emptyset) \geq 1 - \alpha$. Let $V = \{s_0^i : S(s_0^i; \pi_\theta) \cap \mathcal{S}_u = \emptyset\}$. Then, for any $\delta \in (0, 1)$,*

$$\Sigma^{\pi_\theta} \geq \left(\frac{|V|}{N} - \sqrt{\frac{1}{2N} \log \frac{2}{\delta}} \right) (1 - \alpha)(1 - \delta).$$

PROOF SKETCH. For any fixed s_0 , the conformal coverage guarantee implies that $\Pr(S(s_0; \pi_\theta) \cap \mathcal{S}_u = \emptyset) \geq 1 - \alpha$. Now let $x = \Pr_{s \sim \rho}[S(s; \pi_\theta) \cap \mathcal{S}_u = \emptyset]$ and $\hat{x} = |V|/N$. Applying Hoeffding’s inequality, we obtain

$$\Pr\left(x \geq \hat{x} - \sqrt{\frac{1}{2N} \log \frac{2}{\delta}}\right) \geq 1 - \delta,$$

where the probability is over the i.i.d. draws of N initial states s_0 . Combining this with the conformal trajectory guarantee $(1 - \alpha)$ yields the result. \square

We refer to this general approach as *conformal safety analysis* (CSA). A key practical benefit of CSA is that the conformal procedure produces per-time-step reachable sets $S_t(s_0; \pi_\theta)$ so the safety condition can be checked separately for each t . Consequently, our verification problem reduces to checking

$$S_t(s_0; \pi_\theta) \cap \mathcal{S}_u = \emptyset \quad \text{for each } t,$$

Modern neural network verification tools, such as (α, β) -CROWN [38, 48], can handle nonlinear models (e.g., ReLU or tanh networks) and thus general nonlinear safety specifications. Moreover, many practically relevant cases involve relatively simple constraints: for example, safety specification often decomposes across coordinates (e.g., axis-aligned box constraints), in which case it suffices to check safety independently for each dimension.

5.3 Training with Differentiable Conformal and Safety Loss

CSA provides us with a safety verification approach. Our key concern, however, is to *train* policies and dynamic models to enable us to achieve safety. This entails two challenges: first, how do we design a loss function that promotes smaller conformal uncertainty sets and higher confidence (which, in combination, allow us to achieve tighter reachability bounds), and second, how do we design a loss function that specifically promotes safety? We tackle these in turn next.

Conformal Loss. While traditional conformal prediction provides statistical guarantees, it is not differentiable and thus incompatible with end-to-end gradient-based training. To address this issue, we make use of an approach proposed by [7] that enables us to obtain a differentiable approximation of conformal sets using a neural network, which takes state and action as an input, and predicts

the corresponding conformity threshold η . We denote the neural network by $\eta_\omega(s, a)$, with $\eta_{\omega,j}(s, a)$ the threshold predicted for state variable j . This neural network is trained to balance two objectives: 1) minimizing the prediction set size (“efficiency”), and 2) ensuring that the empirical coverage level is above a target threshold.

To formalize, we define the efficiency loss (which promotes smaller conformal sets) as

$$\mathcal{L}_{\text{eff}}(\omega) = \frac{1}{N} \sum_{(s,a) \in \mathcal{D}_e} \left(\prod_{j=1}^n \eta_{\omega,j}(s, a) \right)$$

where n is the state dimension and \mathcal{D}_e is a dataset of tuples (s, a, s') obtained from trajectories $\tilde{\tau}$ sampled in epoch e of training using the latest policy π_θ , where we fix the datasets size $N = |\mathcal{D}_e|$.

Next, the coverage loss penalizes violations of the target coverage, and is defined as

$$\mathcal{L}_{\text{cov}}(\beta, \omega) = \max(0, 1 - \alpha - C_{\text{emp}}),$$

where C_{emp} is the empirical coverage defined as

$$C_{\text{emp}}(\beta, \omega) = \frac{1}{N} \sum_{(s,a,s') \in \mathcal{D}_e} \mathbb{1}[\|s' - \hat{f}_\beta(s, a)\| \leq \eta_\omega(s, a)].$$

However, note that empirical coverage is not differentiable with respect to the parameters (β, ω) . To address this issue, we propose the following differentiable proxy:

$$\hat{C}_{\text{emp}}(\beta, \omega) = \frac{1}{N} \sum_{(s,a,s') \in \mathcal{D}_e} \sigma \left(1 - \max_j \frac{|\hat{f}_\beta(s, a) - s'_j|}{\max\{\eta_{\omega,j}(s, a), c\}} \right)$$

where c is a small constant ensuring that the loss is well-behaved and $\sigma(\cdot)$ is a sigmoid that smoothly approximates the indicator. To unpack, observe that for the max norm, $\max_j \frac{|\hat{f}_\beta(s, a) - s'_j|}{\eta_{\omega,j}(s, a)} \leq 1$ iff $\|s' - \hat{f}_\beta(s, a)\| \leq \eta_\omega(s, a)$, i.e., it is equivalent to s' lying inside the conformal set.

Finally, we combine these into conformal loss as

$$\mathcal{L}_{\text{conf}}(\beta, \omega) = \max_{\lambda \geq 0} (\mathcal{L}_{\text{eff}}(\omega) + \lambda \mathcal{L}_{\text{cov}}(\beta, \omega)),$$

where λ is a Lagrange multiplier periodically updated during training using projected gradient ascent.

Safety Loss. Next, we consider the safety aspect of the problem explicitly. In particular, we construct a safety loss $\mathcal{L}_{\text{safety}}(\theta, \omega)$ that takes the learned dynamics as given, and consists of two components: one which maintains safety over the entire sequence of time steps, and the second which explicitly encourages a policy to improve the safety score between initial and K th step, with K increasing over training iterations somewhat echoing curriculum learning methods for learning safe control [46]. Note that the safety loss is explicitly also a function of the conformal interval parameters η_ω , since these have a direct impact on verified safety. Additionally, safety loss is computed on a pre-collected initial state dataset $\mathcal{D}_0 = \{s_0^i\}_{i=1}^{N_0}$, $s_0^i \sim \rho$. By evaluating safety consistently on this fixed set of initial states rather than randomly sampled states each epoch, we ensure the safety loss signal remains stable and meaningful throughout training, avoiding noise from sampling variance that could obscure true safety improvements.

Define a *safety score function* with respect to a set of states S , $g(S) = \max_{s \in S} h(s)$, quantifying the maximal safety violation of

any $s \in S$. The first safety loss component captures the maximal safety loss over a target safety horizon K_e in training epoch e :

$$\mathcal{L}_{\text{safety}}^{\max}(\theta, \omega) = \max_{\substack{t \in [1, K_e]; \\ s_0 \in \mathcal{D}_0}} g(\hat{f}_\beta^{(t)}(s_0; \pi_\theta) \pm \eta_\omega(s_t, a_t)).$$

The second part of the safety loss aims to incentivize learning to *improve* (or reduce the degradation) of safety over the longer-term evolution of the dynamical system (without consideration of transient terms), which we formalize as

$$\mathcal{L}_{\text{safety}}^{\text{improve}}(\theta, \omega) = \frac{1}{N_0 K_e} \sum_{s_0 \in \mathcal{D}_0} [g(\hat{f}_\beta^{(1)}(s_0; \pi_\theta) \pm \eta_\omega(s_0, a_0)) - g(\hat{f}_\beta^{(K_e)}(s_0; \pi_\theta) \pm \eta_\omega(s_{K_e}, a_{K_e}))]$$

The full safety loss is then a weighted combination of the two terms:

$$\mathcal{L}_{\text{safety}}(\theta, \omega) = w_{\max} \mathcal{L}_{\text{safety}}^{\max}(\theta, \omega) + w_{\text{improve}} \mathcal{L}_{\text{safety}}^{\text{improve}}(\theta, \omega),$$

with the weights w_{\max} and w_{improve} specified as hyperparameters.

5.4 Putting Everything Together: The ReCORS Algorithm

The complete ReCORS algorithm proceeds as follows. First, it pre-trains the dynamics model using a dataset generated by selecting random actions, collects the initial state dataset for safety loss computation, and initializes the policy (actor), critic (as in conventional RL), and conformal uncertainty neural networks. In each training epoch, the algorithm then (i) collects a new dataset of trajectories with the current policy, (ii) fine-tunes the dynamics model using this dataset with \mathcal{L}_{dyn} , and (iii) updates the remaining components. As in standard actor–critic methods, the critic parameters θ_c are trained separately using the conventional RL loss $\mathcal{L}_{\text{RL}}(\theta_c, \theta)$. The policy parameters θ and conformal network parameters ω are then optimized jointly using a combined loss

$$\mathcal{L}(\theta_c, \theta, \beta, \omega) = w_1 \mathcal{L}_{\text{RL}}(\theta_c, \theta) + w_2 \mathcal{L}_{\text{conf}}(\beta, \omega) + w_3 \mathcal{L}_{\text{safety}}(\theta, \omega),$$

where w_1, w_2, w_3 are hyperparameters, θ_c are the parameters of the critic, θ parameters of the actor (policy), β parameters of the dynamics model, and ω parameters of the conformal uncertainty estimation network.

An additional feature of the ReCORS algorithm is a form of *curriculum learning*: we initially begin with a smaller safety horizon K_0 , and incrementally increase the horizon whenever provable safety is achieved or after every E epochs. Algorithm 1 formalizes this discussion.

6 EXPERIMENTS

6.1 Experiment Setup

We evaluate ReCORS in six environments with linear safety constraints and one environment with nonlinear constraints. First, we designed a 2D Quadrotor environment with a nonlinear safe distance constraint for demonstrating ReCORS in the most general setting. And for the linear cases, the first four, Cartpole, Lane Following, and 2D and 3D Quadrotor control, have hard state space safety constraints with binary cost signals, where the episode terminates with any violation. The other two environments, CarGoal and HalfCheetah, are from Safety-Gymnasium [25], which have more

Algorithm 1 ReCORS

Input pretrained dynamics model β , pre-collected initial state dataset \mathcal{D}_0 , initial safety horizon K_0 , epoch counter E , total number of epochs T

- 1: Initialize θ_c, θ, ω
 - 2: **for** epoch $e = 0, 1, 2, \dots, T$ **do**
 - 3: Collect trajectories and induced dataset \mathcal{D}_e using the current policy π_{θ_e} , or random actions if $e = 0$
 - 4: Dynamics update: $\beta_{e+1} \leftarrow \arg \min_{\beta} \mathcal{L}_{\text{dyn}}(\beta_e)$.
 - 5: Critic update: $\theta_{c,e+1} \leftarrow \arg \min_{\theta_c} \mathcal{L}_{\text{RL}}(\theta_{c,e})$
 - 6: Actor and uncertainty update:

$$(\theta_{e+1}, \omega_{e+1}) \leftarrow \arg \min_{\theta, \omega} (w_1 \mathcal{L}_{\text{RL}}(\theta_e) + w_2 \mathcal{L}_{\text{conf}}(\beta_{e+1}, \omega_e) + w_3 \mathcal{L}_{\text{safety}}(\theta_e, \omega_e))$$
 - 7: Evaluate coverage:

$$C_{\text{emp}} = \frac{1}{N} \sum_{(s, a, s') \in \mathcal{D}_e} \mathbf{1}[s' \in \hat{f}_{\beta_{e+1}}(s, a) \pm \eta_{\omega_{e+1}}(s, a)]$$
 - 8: **if** $(\mathcal{L}_{\text{safety}}^{\max} < 0$ and $C_{\text{emp}} \geq 1 - \alpha$) or every E epochs **then**
 - 9: $K_{e+1} = K_e + 1$
 - 10: **else**
 - 11: $K_{e+1} = K_e$
 - 12: **end if**
 - 13: **end for**
-

complex underlying system dynamics, richer cost structures, and are higher dimensional (e.g., 72 dimensions for CarGoal). Detailed experiment descriptions are provided in the Appendix.

We compare to five SOTA baselines: 1) Lagrangian-based Proximal Policy Optimization (PPOLag) [34]; 2) Projection-based Constrained Policy Optimization (PCPO) [50]; 3) Constraint-Rectified Policy Optimization (CRPO) [49]; 4) Penalized Proximal Policy Optimization (P3O) [55]; and 5) Reachability Estimation for Safe Policy Optimization (RESPO) [17]. For consistency and fair comparison, all but RESPO were implemented based on [26] and all used similar hyperparameters. Details of specific settings and hyperparameters are provided in the Appendix.

We use both the time series conformal prediction method proposed by Cleaveland et al. [11] and the union bound applied to the step-wise split conformal prediction. For both methods, we set the trajectory level miscoverage level α to 0.1. Our calibration set contains 1000 trajectories, where, for the times series method, we use 100 for optimization of the w parameters and the rest to compute the actual prediction regions. We then perform forward reachability analysis on 2000 new trajectories (verification set) using the calibrated thresholds. After obtaining the empirical proportions of trajectories that are safe at each time step, we apply Theorem 1 (with $\delta = 0.05$) to compute the lower bound safety probabilities with respect to the entire initial state space.

6.2 Results

6.2.1 Verified Safety Evaluation. We begin by comparing ReCORS to the baselines in terms of verified lower bound on the safety probability (y -axis) as a function of the decision horizon K (x -axis). The results with affine safety constraints are provided in Figure 1,

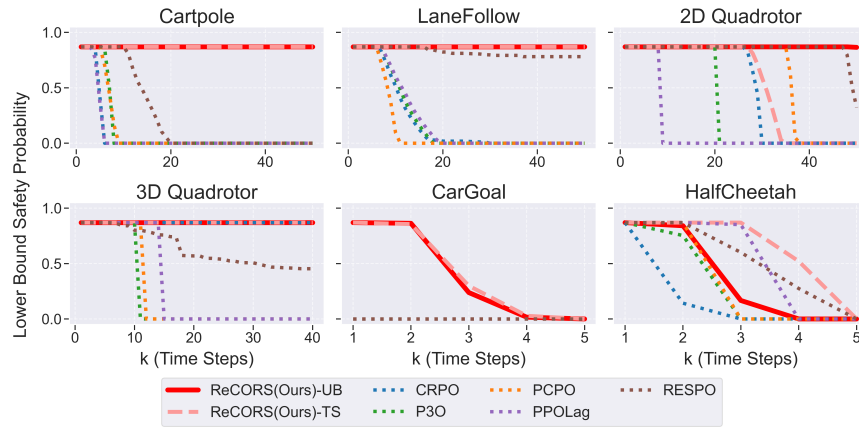


Figure 1: Probabilistic safety verification results. Higher is better. ReCORS-UB (red solid line) is applying conformal prediction per-step and union bounding the results. ReCORS-TS (pink dotted line) uses the time series conformal prediction method from [11]. All the baseline plotted lines use the conformal prediction method that results in better verified safety results.

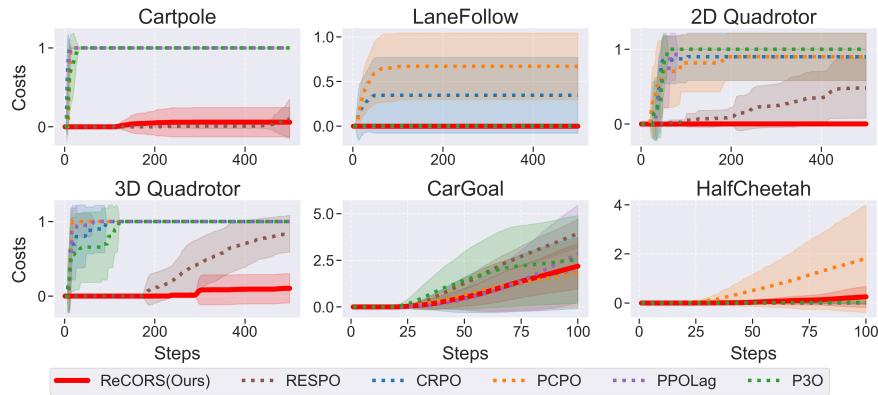


Figure 2: Empirical safety evaluation, comparing safety cost ReCORS (red line) with baselines. Lower is better.




Figure 3: Verified safety (left) and empirical safety (right) for 2D Quadrotor with nonlinear safe distance constraint.

and exhibit notable improvement over the baselines in every environment. For example, in CartPole, none of the baselines provide useful safety guarantees beyond $K = 20$, while ReCORS guarantees safety with high probability through $K \geq 50$. In 2D and 3D Quadrotor, RESPO and CRPO can be verified safe for relatively long horizons just as ReCORS, but this is a superficial consequence of the drone learning to essentially stay in place for both these baselines (obtaining essentially minimal reward in both cases; see Figure 4), while ReCORS achieves significantly higher reward than these *without compromising safety guarantees*. Curiously, while in

most cases the tighter bound either coincides with or improves on the union bound, the latter yields the best performance in the 2D quadrotor case; indeed, in most cases the two bounds essentially coincide. The reason stems from the fact that our trained systems often have bounds that are relatively similar over time, so that the optimal time-series bound parameters are often essentially uniform, corresponding to the union bound.

Figure 3 (left) compares verified safety of ReCORS against baselines in the 2D Quadrotor environment with a nonlinear safe-distance constraint. ReCORS significantly outperforms all baselines except RESPO, but this is again a superficial consequence of RESPO learning resulting in the drone staying in place. Consequently, ReCORS reward is significantly higher than RESPO (see Figure 5).

6.2.2 Empirical Safety Evaluation. Next, we consider to what extent policies are *empirically safe* over far longer horizons than we can verify. We quantify this in terms of *costs of safety violations*, which is just the fraction of time steps for which safety constraints are violated. Figure 2 presents the results for affine safety, while non-linear safety results are in Figure 3 (right). The results are roughly consistent with verified safety, showing that ReCORS results in exceedingly few safety violations over long decision horizons, and

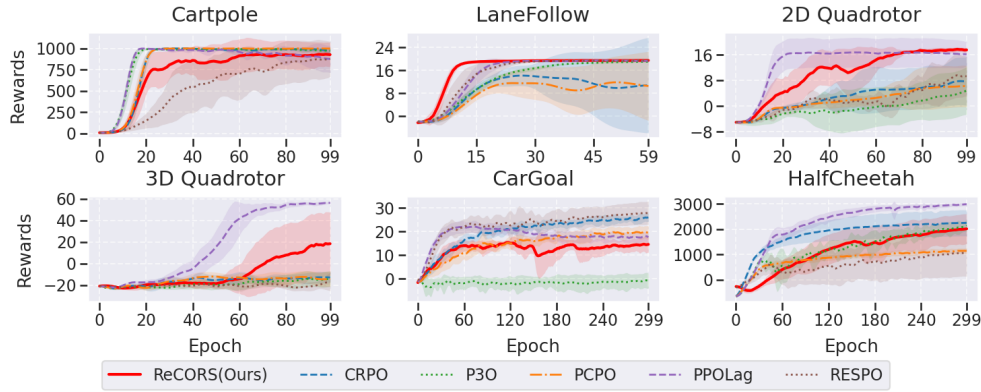


Figure 4: Reward comparisons of ReCORS (red line) with baselines. Higher is better.




Figure 5: Reward comparisons in 2D Quadrotor environment with nonlinear safe distance constraint.

is in all cases either far better than the baselines (e.g., 3D quadrotor), or comparable to the best baseline. In the latter instances, the baselines which are competitive fair far worse in terms of reward.

6.2.3 Empirical Performance Evaluation. Finally, we compare performance of ReCORS and the baselines in Figure 4, which shows the mean and standard deviation of empirical rewards across 10 random seeds as a function of training epochs. Since we expect a natural tradeoff between safety—particularly, verified safety—and reward, it is no surprise that the relative performance of ReCORS varies, and in several domains other methods, such as PPOLag and CRPO, exhibit better performance (reward). However, in all cases in which other methods achieve better reward than ReCORS, they exhibit very poor safety properties; indeed, ReCORS is the only method that can achieve both strong performance and verified safety. To illustrate, consider the HalfCheetah environment, where CRPO and PPOLag achieve higher reward. Here, CRPO achieves the worst safety, and while PPOLag is better, ReCORS achieves significantly better verified safety. In several domains, such as 2D Quadrotor, ReCORS achieves both, superior safety and reward than all baselines.

6.3 Ablations

Finally, we provide ablation results on two design choices: 1) including a heuristic safety weight in training the dynamics model, and 2) jointly optimizing the dynamics and uncertainty model in the conformal coverage loss $\mathcal{L}_{cov}(\beta, \omega)$. Results are shown in Figure 6. We find that we have the best performance omitting dependence on β in the coverage loss (ReCORS is better than D_conv_loss). Additionally, using safety-aware weights for learning \hat{f}_β provides better

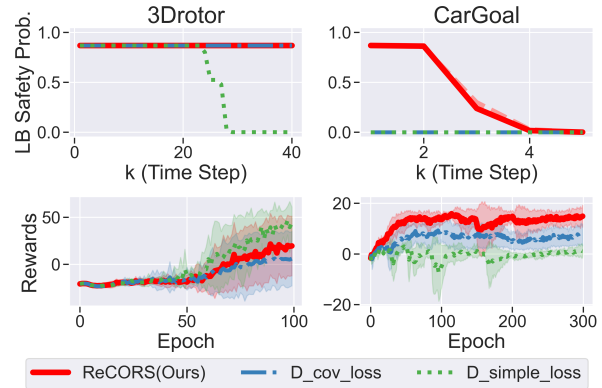


Figure 6: Ablation results for 3D Quadrotor (left) and CarGoal (right). D_cov_loss jointly optimizes β and ω in the conformal coverage loss, while ReCORS only uses the latter (ω). D_simple_loss uses uniform weights in learning \hat{f}_β .

verifies safety while maintaining good performances (ReCORS is safer than D_simple_loss).

7 CONCLUSION

We present a novel approach for learning effective policies in general dynamical systems while abiding by probabilistic safety constraints. Our approach is designed to deal with unknown stochastic dynamics, and avoids restrictive assumptions on either true dynamics or how they are learned. This is accomplished by combining conformal predictions, finite-horizon reachability proofs, and finite-sample tail bounds to obtain verified safety, with model-free RL training methods facilitating strong performance. While we achieve state-of-the-art results on several standard benchmarks, many open problems remain, paramount among these being an assumption of perfect observability of the system state.

ACKNOWLEDGEMENTS

This work was partially supported by the NSF (IIS-2214141, CCF-2403758), ARO (W911NF-25-1-0059), ONR (N000142412663), Foresight Institute, and Amazon.

REFERENCES

- [1] Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. 2017. Constrained policy optimization. In *International conference on machine learning*. 22–31.
- [2] Eitan Altman. 2021. *Constrained Markov Decision Processes*. Routledge.
- [3] Aaron D Ames, Samuel Coogan, Magnus Egerstedt, Gennaro Notomista, Koushil Sreenath, and Paulo Tabuada. 2019. Control barrier functions: Theory and applications. In *European Control Conference*. 3420–3431.
- [4] Aaron D Ames, Xiangru Xu, Jessy W Grizzle, and Paulo Tabuada. 2016. Control barrier function based quadratic programs for safety critical systems. *IEEE Trans. Automat. Control* 62, 8 (2016), 3861–3876.
- [5] Anastasios N Angelopoulos and Stephen Bates. 2021. A gentle introduction to conformal prediction and distribution-free uncertainty quantification. *arXiv preprint arXiv:2107.07511* (2021).
- [6] Ahmad Taher Azar, Anis Koubaa, Nada Ali Mohamed, Habiba A Ibrahim, Zahra Fathy Ibrahim, Muhammad Kazim, Adel Ammar, Bilel Benjdira, Alaa M Khamis, Ibrahim A Hameed, et al. 2021. Drone deep reinforcement learning: A review. *Electronics* 10, 9 (2021), 999.
- [7] Yu Bai, Song Mei, Huan Wang, Yingbo Zhou, and Caiming Xiong. 2022. Efficient and differentiable conformal prediction with general function classes. *arXiv preprint arXiv:2202.11091* (2022).
- [8] Kaustav Chakraborty, Zeyuan Feng, Sushant Veer, Apoorva Sharma, Wenhao Ding, Sever Topan, Boris Ivanovic, Marco Pavone, and Somil Bansal. 2025. Safety Evaluation of Motion Plans Using Trajectory Predictors as Forward Reachable Set Estimators. *arXiv preprint arXiv:2507.22389* (2025).
- [9] Richard Cheng, Gábor Orosz, Richard M Murray, and Joel W Burdick. 2019. End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 3387–3395.
- [10] Yinlam Chow, Ofir Nachum, Edgar Duenez-Guzman, and Mohammad Ghavamzadeh. 2018. A Lyapunov-based approach to safe reinforcement learning. *Advances in neural information processing systems* 31 (2018).
- [11] Matthew Cleaveland, Insup Lee, George J Pappas, and Lars Lindemann. 2024. Conformal prediction regions for time series using linear complementarity programming. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 20984–20992.
- [12] Charles Dawson, Zengyi Qin, Sicun Gao, and Chuchu Fan. 2022. Safe nonlinear control using robust neural Lyapunov-barrier functions. In *Conference on Robot Learning*. PMLR, 1724–1735.
- [13] Chuchu Fan. 2019. *Formal methods for safe autonomy: Data-driven verification, synthesis, and applications*. Ph.D. Dissertation. University of Illinois at Urbana-Champaign.
- [14] Chuchu Fan and Sayan Mitra. 2019. Data-driven safety verification of complex cyber-physical systems. In *Design Automation of Cyber-Physical Systems*. Springer, 107–142.
- [15] Chuchu Fan, Bolun Qi, Sayan Mitra, and Mahesh Viswanathan. 2017. DryVR: Data-driven verification and compositional reasoning for automotive systems. In *International Conference on Computer Aided Verification*. Springer, 441–461.
- [16] Jaime F Fisac, Neil F Lugovoy, Vicenc Rubies-Royo, Shromona Ghosh, and Claire J Tomlin. 2019. Bridging hamilton-jacobi safety analysis and reinforcement learning. In *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 8550–8556.
- [17] Milan Ganai, Zheng Gong, Chenning Yu, Sylvia Herbert, and Sicun Gao. 2023. Iterative reachability estimation for safe reinforcement learning. *Neural Information Processing Systems* 36 (2023), 69764–69797.
- [18] Shangding Gu, Long Yang, Yali Du, Guang Chen, Florian Walter, Jun Wang, and Alois Knoll. 2022. A review of safe reinforcement learning: Methods, theory and applications. *arXiv preprint arXiv:2205.10330* (2022).
- [19] Dong Han, Beni Mulyana, Vladimir Stankovic, and Samuel Cheng. 2023. A survey on deep reinforcement learning algorithms for robotic manipulation. *Sensors* 23, 7 (2023), 3762.
- [20] Navid Hashemi, Lars Lindemann, and Jyotirmoy V Deshmukh. 2024. Statistical reachability analysis of stochastic cyber-physical systems under distribution shift. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 43, 11 (2024), 4250–4261.
- [21] Navid Hashemi, Xin Qin, Lars Lindemann, and Jyotirmoy V Deshmukh. 2023. Data-driven reachability analysis of stochastic dynamical systems with conformal inference. In *IEEE Conference on Decision and Control*. 3102–3109.
- [22] Radoslav Ivanov, James Weimer, Rajeev Alur, George J Pappas, and Insup Lee. 2019. Verisig: verifying safety properties of hybrid systems with neural network controllers. In *ACM International Conference on Hybrid Systems: Computation and Control*. 169–178.
- [23] John Jackson, Luca Laurenti, Eric Frew, and Morteza Lahijanian. 2020. Safety verification of unknown dynamical systems via gaussian process regression. In *IEEE Conference on Decision and Control*. 860–866.
- [24] Ashish K Jayant and Shalabh Bhatnagar. 2022. Model-Based Safe Deep Reinforcement Learning via a Constrained Proximal Policy Optimization Algorithm. *Advances in Neural Information Processing Systems* 35 (2022), 24432–24445.
- [25] Jiaming Ji, Borong Zhang, Jiayi Zhou, Xuehai Pan, Weidong Huang, Ruiyang Sun, Yiran Geng, Yifan Zhong, Josef Dai, and Yaodong Yang. 2023. Safety Gymnasium: A Unified Safe Reinforcement Learning Benchmark. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*. <https://openreview.net/forum?id=WZmlxluGR>
- [26] Jiaming Ji, Jiayi Zhou, Borong Zhang, Juntao Dai, Xuehai Pan, Ruiyang Sun, Weidong Huang, Yiran Geng, Mickel Liu, and Yaodong Yang. 2024. Omnisafe: An infrastructure for accelerating safe reinforcement learning research. *Journal of Machine Learning Research* 25, 285 (2024), 1–6.
- [27] B Ravi Kiran, Ibrahim Sobh, Victor Talpaert, Patrick Mannion, Ahmad A Al Sallab, Senthil Yogamani, and Patrick Pérez. 2021. Deep reinforcement learning for autonomous driving: A survey. *IEEE transactions on intelligent transportation systems* 23, 6 (2021), 4909–4926.
- [28] Niklas Kochdumper, Hanna Krasowski, Xiao Wang, Stanley Bak, and Matthias Althoff. 2023. Provably safe reinforcement learning via action projection using reachability analysis and polynomial zonotopes. *IEEE Open Journal of Control Systems* 2 (2023), 79–92.
- [29] Lars Lindemann, Yiqi Zhao, Xinyi Yu, George J Pappas, and Jyotirmoy V Deshmukh. 2024. Formal verification and control with conformal prediction. *arXiv preprint arXiv:2409.00536* (2024).
- [30] Haitong Ma, Yang Guan, Shegno Eben Li, Xiangteng Zhang, Sifa Zheng, and Jianyu Chen. 2021. Feasible actor-critic: Constrained reinforcement learning for ensuring statewise safety. *arXiv preprint arXiv:2105.10682* (2021).
- [31] Xinhang Ma, Junlin Wu, Hussein Sibai, Yiannis Kantaros, and Yevgeniy Vorobeychik. 2025. Learning Vision-Based Neural Network Controllers with Semi-Probabilistic Safety Guarantees. *arXiv preprint arXiv:2503.00191* (2025).
- [32] Ye Cheng Jason Ma, Andrew Shen, Osbert Bastani, and Jayaraman Dinesh. 2022. Conservative and Adaptive Penalty for Model-Based Safe Reinforcement Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 5404–5412.
- [33] Anish Muthali, Haotian Shen, Sampada Deglurkar, Michael H Lim, Rebecca Roelofs, Aleksandra Faust, and Claire Tomlin. 2023. Multi-agent reachability calibration with conformal prediction. In *2023 62nd IEEE Conference on Decision and Control (CDC)*. 6596–6603.
- [34] Alex Ray, Joshua Achiam, and Dario Amodei. 2019. Benchmarking safe exploration in deep reinforcement learning. *arXiv preprint arXiv:1910.01708* 7, 1 (2019), 2.
- [35] Ali Salamati and Majid Zamani. 2022. Data-driven safety verification of stochastic systems via barrier certificates: A wait-and-judge approach. In *Learning for Dynamics and Control Conference*. PMLR, 441–452.
- [36] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- [37] Mahmoud Selim, Amr Alanwar, Shreyas Kousik, Grace Gao, Marco Pavone, and Karl H Johansson. 2022. Safe reinforcement learning using black-box reachability analysis. *IEEE Robotics and Automation Letters* 7, 4 (2022), 10665–10672.
- [38] Zhouxing Shi, Qirui Jin, Zico Kolter, Suman Jana, Cho-Jui Hsieh, and Huan Zhang. 2025. Neural network verification with branch-and-bound for general nonlinearities. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 315–335.
- [39] Oswin So and Chuchu Fan. 2023. Solving Stabilize-Avoid Optimal Control via Epigraph Form and Deep Reinforcement Learning. In *Proceedings of Robotics: Science and Systems*.
- [40] Adam Stooke, Joshua Achiam, and Pieter Abbeel. 2020. Responsive Safety in Reinforcement Learning by Pid Lagrangian Methods. In *International Conference on Machine Learning*. PMLR, 9133–9143.
- [41] Chen Tessler, Daniel J Mankowitz, and Shie Mannor. 2018. Reward constrained policy optimization. *arXiv preprint arXiv:1805.11074* (2018).
- [42] Hoang-Dung Tran, Feiyang Cai, Manzan Lopez Diego, Patrick Musau, Taylor T Johnson, and Xenofon Koutsoukos. 2019. Safety verification of cyber-physical systems with reinforcement learning control. *ACM Transactions on Embedded Computing Systems* 18, 5s (2019), 1–22.
- [43] Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. 2005. *Algorithmic learning in a random world*. Vol. 29. Springer.
- [44] Hao Wang, Javier Borquez, and Somil Bansal. 2024. Providing safety assurances for systems with unknown dynamics. *IEEE Control Systems Letters* (2024).
- [45] Yixuan Wang, Simon Sinong Zhan, Ruochen Jiao, Zhilu Wang, Wanxin Jin, Zhuoran Yang, Zhaoran Wang, Chao Huang, and Qi Zhu. 2023. Enforcing hard constraints with soft barriers: Safe reinforcement learning in unknown stochastic environments. In *International Conference on Machine Learning*. PMLR, 36593–36604.
- [46] Junlin Wu, Huan Zhang, and Yevgeniy Vorobeychik. 2024. Verified Safe Reinforcement Learning for Neural Network Dynamic Models. In *Neural Information Processing Systems*.
- [47] Botian Xu, Feng Gao, Chao Yu, Ruize Zhang, Yi Wu, and Yu Wang. 2024. Omnidrones: An efficient and flexible platform for reinforcement learning in drone control. *IEEE Robotics and Automation Letters* 9, 3 (2024), 2838–2844.
- [48] Kaidi Xu, Huan Zhang, Shiqi Wang, Yihan Wang, Suman Jana, Xue Lin, and Cho-Jui Hsieh. 2020. Fast and complete: Enabling complete neural network

- verification with rapid and massively parallel incomplete verifiers. *arXiv preprint arXiv:2011.13824* (2020).
- [49] Tengyu Xu, Yingbin Liang, and Guanghui Lan. 2021. CRPO: A new approach for safe reinforcement learning with convergence guarantee. In *International Conference on Machine Learning*. 11480–11491.
- [50] Tsung-Yen Yang, Justinian Rosca, Karthik Narasimhan, and Peter J Ramadge. 2020. Projection-based constrained policy optimization. *arXiv preprint arXiv:2010.03152* (2020).
- [51] Dongjie Yu, Haitong Ma, Shengbo Li, and Jianyu Chen. 2022. Reachability constrained reinforcement learning. In *International Conference on Machine Learning*. 25636–25655.
- [52] Dongjie Yu, Wenjun Zou, Yujie Yang, Haitong Ma, Shengbo Eben Li, Yuming Yin, Jianyu Chen, and Jingliang Duan. 2023. Safe model-based reinforcement learning with an uncertainty-aware reachability certificate. *IEEE Transactions on Automation Science and Engineering* 21, 3 (2023), 4129–4142.
- [53] Hongchao Zhang, Zhizhen Qin, Sicun Gao, and Andrew Clark. 2024. SEEV: Synthesis with Efficient Exact Verification for ReLU Neural Barrier Functions. In *Neural Information Processing Systems*. 101367–101392.
- [54] Hongchao Zhang, Junlin Wu, Yevgeniy Vorobeychik, and Andrew Clark. 2023. Exact verification of relu neural control barrier functions. In *Neural Information Processing Systems*. 5685–5705.
- [55] Linrui Zhang, Li Shen, Long Yang, Shixiang Chen, Bo Yuan, Xueqian Wang, and Dacheng Tao. 2022. Penalized proximal policy optimization for safe reinforcement learning. *arXiv preprint arXiv:2205.11814* (2022).
- [56] Tengting Zhang and Hongwei Mo. 2021. Reinforcement learning for robot research: A comprehensive review and open issues. *International Journal of Advanced Robotic Systems* 18, 3 (2021), 1–21.
- [57] Weiye Zhao, Tairan He, and Changliu Liu. 2023. Probabilistic safeguard for reinforcement learning using safety index guided gaussian process models. In *Learning for Dynamics and Control Conference*. PMLR, 783–796.
- [58] Ruikun Zhou, Thanin Quartz, Hans De Sterck, and Jun Liu. 2022. Neural Lyapunov control of unknown nonlinear systems with stability guarantees. *Advances in Neural Information Processing Systems* 35 (2022), 29113–29125.