

# Reinforcement Learning with Ensemble Model Predictive Safety Certification

Sven Gronauer  
 Technical University of Munich (TUM)  
 Munich, Germany  
 sven.gronauer@tum.de

Tom Haider  
 Fraunhofer IKS  
 Munich, Germany  
 tom.haider@iks.fraunhofer.de

Felippe Schmoeller da Roza  
 Fraunhofer IKS  
 Munich, Germany  
 felippe.schmoeller.da.roza@iks.fraunhofer.de

Klaus Diepold  
 Technical University of Munich (TUM)  
 Munich, Germany  
 kldi@tum.de

## ABSTRACT

Reinforcement learning algorithms need exploration to learn. However, unsupervised exploration prevents the deployment of such algorithms on safety-critical tasks and limits real-world deployment. In this paper, we propose a new algorithm called *Ensemble Model Predictive Safety Certification* that combines model-based deep reinforcement learning with tube-based model predictive control to correct the actions taken by a learning agent, keeping safety constraint violations at a minimum through planning. Our approach aims to reduce the amount of prior knowledge about the actual system by requiring only offline data generated by a safe controller. Our results show that we can achieve significantly fewer constraint violations than comparable reinforcement learning methods.

## KEYWORDS

Reinforcement Learning, Safe Reinforcement Learning, Safe Exploration, Predictive Safety Filter, Model-based Learning,

### ACM Reference Format:

Sven Gronauer, Tom Haider, Felippe Schmoeller da Roza, and Klaus Diepold. 2024. Reinforcement Learning with Ensemble Model Predictive Safety Certification. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 9 pages.

## 1 INTRODUCTION

Deep reinforcement learning (RL) is a powerful data-driven paradigm for learning control strategies and has recently achieved remarkable results in various domains [23, 32]. RL is beneficial in situations where the system dynamics are not known or only partially available, but data can be generated through interaction with the environment. However, gathering data in safety-critical tasks and real-world systems is not trivial since the agent is required to act safely at all times, *e.g.*, the actions taken by a robot are supposed to satisfy a series of state and control input constraints to avoid

harm to itself or its environment. On top of the difficulty of providing safety, other self-imposed challenges, such as the high sample complexity and the lack of interpretability, impede the adoption of deep RL in real-world applications that go beyond simulations and fail-safe academic environments [12].

The majority of safe RL algorithms are designed to merely incentivize safety instead of ensuring hard safety constraint satisfaction [6]. Policies are encouraged to maximize the task performance while constraint violations are in expectation less or equal to a predefined safety threshold, resulting in safety constraint satisfaction at the end but not throughout training. In contrast, methods derived from model predictive control (MPC) provide formal safety guarantees by making rather strong assumptions and, thus, are a popular choice for designing safe controllers. Yet, their applicability is often limited to low-dimensional systems, as evidenced in [5, 16]. Deep RL, however, has the potential of scaling to high-dimensional problems [1, 34].

In this paper, we introduce a novel algorithm called ensemble model predictive safety certification (X-MPSC) that extends model-based deep RL with tube-based MPC to certify potentially unsafe actions taken by a learning agent. The result is an algorithm that combines the best of both frameworks by leveraging an ensemble of probabilistic neural networks (NNs) to approximate the system dynamics trained on data sampled from the environment. To provide safe exploration, MPC is used to plan multiple tube-based trajectories that enforce all given safety constraints based on the NN ensemble. The actions of an RL-based agent are modified to safe ones if necessary. For initialization, our method only requires offline data collected by a low-performing but safe controller. The experimental results demonstrate that our algorithm can significantly reduce the number of constraint violations compared to alternative constrained RL algorithms. Constraint violations can even be reduced to zero when a coarse prior system dynamics model is incorporated into the learning loop.

## 2 RELATED WORK

The artificial intelligence community has different notions of what constitutes a safe system [4, 18]. For RL, safety can be achieved by preventing error states, *i.e.*, undesirable states from which the system’s original state cannot be recovered. Error states are tightly coupled to the concepts of reachability [24] and set invariance [3]. Another perspective is to define an RL system to be safe when it



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

is able to maximize a performance measure while fulfilling or encouraging safety constraints during both learning and deployment phases [13, 26, 29]. In this paper, we adopt the latter safety definition as an instance of the constrained Markov decision process (CMDP) framework [2].

In a recent survey, Brunke et al. [6] unify the perspectives from both control theory and RL and provide a systematic overview of safe learning-based control in the context of robotics. The authors classify the task of achieving safe learning-based control into three main categories: (1) the formal certification of safety, (2) RL approaches that encourage safety, and (3) the improvement of system performance by safely learning the uncertain dynamics. Because our proposed algorithm addresses safety certification but compares to safety-encouraging RL methods, we limit the following literature review to these two categories.

*Formal Certification.* Methods from this category utilize prior system dynamics knowledge to provide rigorous safety through hard constraint satisfaction. One way to achieve this is to use *safety filters* such as model predictive safety certification (MPSC) that adapt input actions as minimal as possible to fulfill safety constraints [37, 38]. Another kind of safety filter is a control barrier function (CBF), which is a mechanism to prevent the system from entering unsafe regions. A CBF maps the state space to a scalar value and is defined to change signs when the system enters an unsafe region of the state space [3]. Cheng et al. [7] showed that CBFs could be integrated into model-free RL to achieve safe exploration for continuous tasks, while Robey et al. [30] learned a CBF from expert demonstrations. Luo and Ma [21] used barrier certificates to certify the stability of a closed dynamical system. By iteratively learning a dynamics model and a barrier certificate alongside a policy, they can ensure that no safety violations occur during training. However, their experimental evaluation was conducted on very low-dimensional state spaces. Similar to our work, Koller et al. [16] used learning-based stochastic MPC for multi-step look ahead predictions to correct any potentially unsafe actions based on a single probabilistic Gaussian process (GP) model. Another work by Pfrommer et al. [27] proposed a chance-constrained MPC approach that uses a safety penalty term in the objective to guide policy gradient updates.

*Encouraging Safety.* The standard formulation in safe RL is to model the environment as a CMDP [29] to learn how to respect safety thresholds while learning the task, leading to soft constraint satisfaction in most cases. One instance is *Lagrangian relaxation methods*, which add a penalty term for constraint violations to the objective such that an unconstrained optimization problem is solved instead [8, 20, 34]. Another common approach is to perform a *constrained policy search* where typically the cost objective function is linearized around the current policy iterate [1, 39]. Lastly, *action projection methods* are applied to correct actions taken by an agent and turn the actions into safe ones, e.g., via Lyapunov functions [9], in closed form with linearized cost models [11] or by evaluating risk-aware Q-functions [33]. Thananjeyan et al. [35] simultaneously learned a task policy, focused solely on task performance, and a recovery policy, activated when constraint violation is likely, which guides the agent back to a safe state. By separating task performance

and constraint satisfaction into two separate policies, safety and reward maximization are balanced more efficiently.

Our algorithm extends previous methods based on MPC by relaxing the requirement of prior knowledge about the system dynamics or knowing the terminal set *a priori*. To the best of our knowledge, this is the first work that integrates an NN ensemble into the tube-based MPC framework. Luo and Ma [21] utilized an NN ensemble together with barrier certificates, while Lütjens et al. [22] deployed an ensemble of recurrent NNs for predictive uncertainty estimates realized by Monte Carlo dropout and bootstrapping.

### 3 PRELIMINARIES

Throughout this work, we consider the dynamics of a system in discrete time described by

$$x_{t+1} = f(x_t, u_t, w_t), \quad (1)$$

with states  $x \in \mathbb{X}$ , actions  $u \in \mathbb{U}$ , and disturbances  $w \in \mathbb{W}$  at time step  $t$ . We assume that the disturbances are bounded and that the system dynamics are Lipschitz continuous. Further, we assume that the system is subject to polytopic constraints in the states and actions, i.e.,  $x \in \mathbb{X} = \{x \in \mathbb{R}^{n_x} \mid H_x x \leq d_x\}$  and  $u \in \mathbb{U} = \{u \in \mathbb{R}^{n_u} \mid H_u u \leq d_u\}$ .

#### 3.1 Deep Reinforcement Learning

The standard framework for RL problems is the Markov decision process (MDP) which is formalized by the tuple  $(\mathbb{X}, \mathbb{U}, f, r, \mathbb{X}_0)$ , where the system  $f$  underlies a random disturbance  $w_t$  with the transition probability distribution given by  $x_{t+1} \sim p(\cdot | x_t, u_t)$ . The set  $\mathbb{X}_0$  denotes the initial state distribution, and  $r : \mathbb{X} \times \mathbb{U} \rightarrow \mathbb{R}$  is the reward function. Note that the dynamics model from Equation (1) is equivalently described by the state probability function  $p(x_{t+1} | x_t, u_t)$ . The optimization objective of RL is given by

$$\text{maximize}_{\theta} J_{\text{RL}}(\pi_{\theta}) = \mathbb{E}_{\tau \sim \pi_{\theta}} \left[ \sum_{k=0}^{\infty} \gamma^k r(x_k, u_k) \right],$$

where the expected return along the trajectories  $\tau = (x_0, u_0, x_1, \dots)$  produced under the policy  $\pi_{\theta}$  is optimized. A policy  $\pi_{\theta} : \mathbb{X} \rightarrow P(\mathbb{U})$  describes the mapping from states to a distribution over actions, where the vector  $\theta$  parameterizes the NN representing the policy. The shortcut  $\tau \sim \pi_{\theta}$  describes trajectories generated under policy  $\pi_{\theta}$  given  $x_{t+1} \sim p(\cdot | x_t, u_t)$ ,  $u_t \sim \pi_{\theta}(\cdot | x_t)$ , and  $x_0 \sim \mathbb{X}_0$ . Finally,  $\gamma \in (0, 1)$  denotes the discount factor.

#### 3.2 Model Predictive Safety Certification

The nominal MPSC problem as introduced in [36] seeks a control input  $v_0$  that changes the learning input  $u_t$  as minimal as possible by solving the objective

$$\begin{aligned} & \text{minimize}_{v_0, \dots, v_{N-1}} && \|u_t - v_0\|_2^2 \\ & \text{subject to} && z_0 = x_t \\ & && z_{k+1} = f_{\text{prior}}(z_k, v_k) \quad \forall k = [0, N-1] \\ & && v_k \in \mathbb{U} \quad \forall k = [0, N-1] \\ & && z_k \in \mathbb{X} \quad \forall k = [0, N] \\ & && z_N \in \mathbb{X}_{\text{term}} \end{aligned} \quad (2)$$

over a finite horizon  $N$  such that the nominal state-action sequence  $(z_0, v_0, \dots, v_{N-1}, z_N)$  lies within the given state-action constraints.

Note that we use  $t$  as the time index for state measurements and actions that are applied to (1), whereas  $k$  indicates states and actions used for planning such that the predicted states  $z_k$  are  $k$  stages ahead of time step  $t$ . The terminal set  $\mathbb{X}_{\text{term}}$  acts as a constraint that must be reached from  $x_t$  within  $N$  stages. MPSC assumes access to a model  $f_{\text{prior}}$ , which is usually derived from first principles. MPSC was initially proposed for linear systems [37] and extended to systems with nonlinear dynamics in later works [36, 38].

### 3.3 Tube-based Model Predictive Control

If uncertainties and disturbances exist in the system under control, feedback control is superior to open loop control. While conventional MPC finds a nominal action sequence as a solution for the open loop control problem, robust MPC returns a sequence of feedback policies. In this paper, we focus on *tube-based* MPC [28] as implementation to approach robustness. Tube-based MPC utilizes a model  $f_{\text{prior}}$  to plan a nominal state trajectory  $(z_0, \dots, z_N)$  associated with the action sequence  $(v_0, \dots, v_{N-1})$  based on the latest measurement  $x_t$  from (1). In presence of uncertainty, it is assumed that the tube contains all possible realizations of the actual system, where each realization implements a series of disturbances. Since tubes can grow large under uncertainty, a closed loop feedback

$$u_{t+k} = v_k + K(x_{t+k} - z_k) \quad (3)$$

is used to track the state trajectory  $(x_t, x_{t+1}, \dots)$  of the actual system toward the nominal trajectory. The matrix  $K \in \mathbb{R}^{n_u \times n_x}$  implements the feedback and is chosen such that the error system  $e_k = x_{t+k} - z_k$  is stable.

### 3.4 Ellipsoidal Calculus

An ellipsoid  $\mathcal{E}(c, S) = \{x \mid (x - c)^T S^{-1} (x - c) \leq 1\}$  describes an affine transformation of the unit ball with center  $c \in \mathbb{R}^{n_x}$  and positive definite shape matrix  $S \in \mathbb{R}^{n_x \times n_x}$ . Ellipsoids are preserved under affine transformations since  $A\mathcal{E}(c, S) = \mathcal{E}(Ac, ASA^T)$ . Although the result of a set addition of two ellipsoids is, in general, not ellipsoidal, we can outer approximate the operation [17]. The over-approximated ellipsoid can be computed through

$$\mathcal{E}(c_1, S_1) \oplus \mathcal{E}(c_2, S_2) \subseteq \mathcal{E}(c_+, S_+) \quad (4)$$

with shape matrix  $S_+ = (1 + \alpha^{-1})S_1 + (1 + \alpha)S_2$  and center  $c_+ = c_1 + c_2$  given by  $\alpha = \sqrt{\text{Tr}(S_1)/\text{Tr}(S_2)}$ . Another useful expression is to check whether the ellipsoid  $\mathcal{E}(c, S)$  is contained in the polytope  $\mathcal{P} = \{x \mid Hx \leq d\}$ , where  $Hx \leq d$  describes a system of linear inequalities. The inscription is evaluated by

$$h_j^T c - d_j + \sqrt{h_j^T S h_j} \leq 0 \quad \forall j, \quad (5)$$

where  $h_j$  is the  $j$ -th row of  $H$  and  $d_j$  is the  $j$ -th vector component.

Ellipsoids have favorable geometrical properties compared to polytopes, *e.g.*, under uncertainty the computational complexity is linear over the predictive horizon. Also, the analytical expressions of (4) can be exploited to maintain differentiability along the predicted trajectory.

## 4 ENSEMBLE MODEL PREDICTIVE SAFETY CERTIFICATION

Our proposed algorithm integrates an ensemble of NNs and tube-based MPC into (2) to certify the actions of a model-based RL agent. Trained on trajectory data from the environment, the ensemble of dynamics models is leveraged for both RL policy optimization and planning with tube-based MPC. The ensemble is represented by probabilistic NNs and parametrizes state-action dependent ellipsoidal predictions, propagated over multiple time steps. An action is certified as safe when the planned tubes satisfy the given state-action constraints and capture the trajectory of the actual system.

In the remainder of this section, we first describe our approach for a single dynamics model in Sections 4.1–4.2 and then extend to ensembles in Section 4.3. We finally present the X-MPSC optimization problem and our algorithm in Sections 4.4–4.5.

### 4.1 Neural Network Parametrization

To approximate the system dynamics, we use probabilistic NNs

$$f_\phi(x_t, u_t) = \mathcal{N}(m_\phi(x, u), S_\phi(x, u)) \quad (6)$$

parametrized by Gaussian probability distribution functions with the mean  $m_\phi(x, u)$  and the diagonal covariance matrix  $S_\phi(x, u)$ . The predicted uncertainties are state- and action-dependent and are determined by the vector  $\phi$  that holds the flattened weights and biases of the NN. Similar to [10], we train a dynamics model by optimizing the maximum-likelihood

$$\underset{\phi}{\text{minimize}} \quad (m_\phi - x_{t+1})^T S_\phi^{-1} (m_\phi - x_{t+1}) + \log \det S_\phi, \quad (7)$$

over trajectory data sampled from (1). An apparent advantage of using probabilistic over deterministic NNs is that aleatoric uncertainty can be captured, *i.e.*, the stochasticity inherent to a system.

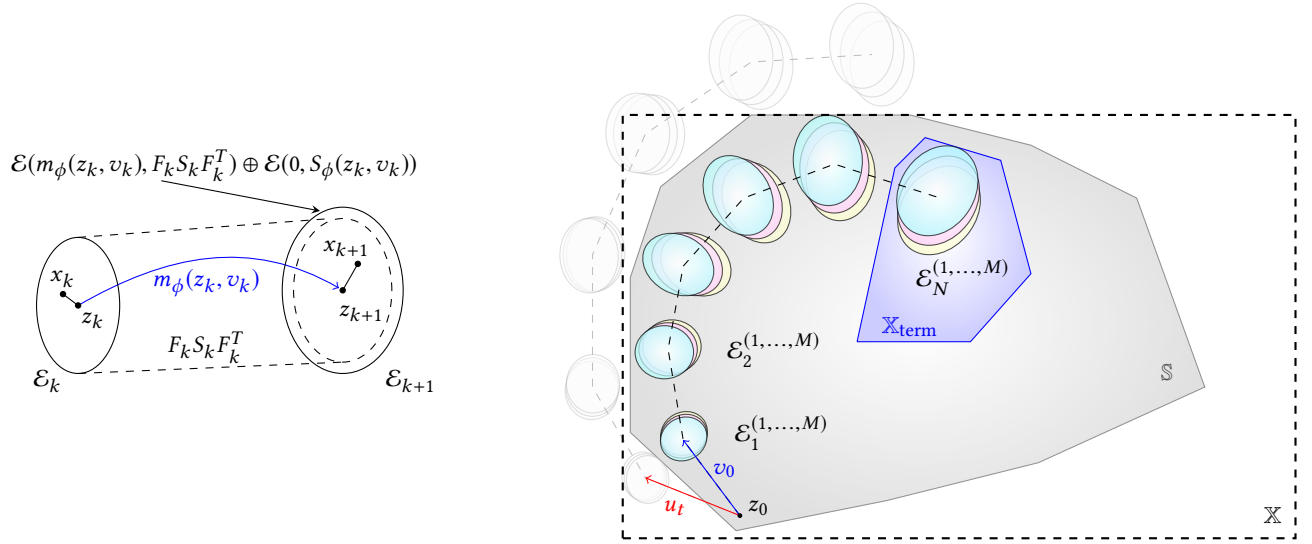
### 4.2 Single-Model Uncertainty Propagation

To predict the future evolution of (1), we use a probabilistic dynamics model  $f_\phi$  for multi-step look ahead rollouts. The nominal state trajectory described by  $(z_0, z_1, \dots, z_N)$  is produced by the action sequence  $(v_0, v_1, \dots, v_{N-1})$  with respect to the nonlinear model  $z_{k+1} = m_\phi(z_k, v_k)$  for all  $k \in [0, N-1]$ . In addition to the mean, the probabilistic model provides an ellipsoidal uncertainty estimate  $S_\phi$  that is propagated over multiple time steps. However, the resulting tube can grow large when an open loop action sequence rather than a closed loop feedback is used for planning [28]. Thus, we employ the affine control law from (3) to keep the actual system close to the nominal trajectory.

We now consider the uncertainty tube along the nominal trajectory  $(z_0, v_0, \dots, v_{N-1}, z_N)$ . For a one-step error prediction, we make use of the first-order Taylor-series expansion

$$x_{t+k+1} \approx m_\phi(z_k, v_k) + A_k(x_{t+k} - z_k) + B_k(u_{t+k} - v_k),$$

around a fixed point  $(z_k, v_k)$  given  $A_k = \nabla_x m_\phi(x, u)^T |_{x=z_k, u=v_k}$  and  $B_k = \nabla_u m_\phi(x, u)^T |_{x=z_k, u=v_k}$  as the Jacobians. The predicted error  $e_{k+1} = x_{t+k+1} - z_{k+1}$  between the nominal state and the



**Figure 1: X-MPSC uses multi-step planning with ellipsoidal uncertainty estimates. (Left) Ellipsoidal uncertainty propagation with a single model. (Right) Tube-based predictions generated by an ensemble of NN models. By utilizing multiple models, an unsafe action  $u_t$  (red) is corrected to  $v_0$  (blue) that keeps the system within the safety constraints over the horizon  $N$ .**

actual system state satisfies the error difference equation

$$\begin{aligned} e_{k+1} &\approx A_k(x_{t+k} - z_k) + B_k(u_{t+k} - v_k) \\ &= A_k(x_{t+k} - z_k) + B_k K(x_{t+k} - z_k) \\ &= (A_k + B_k K)(x_{t+k} - z_k) \\ &= F_k e_k \end{aligned}$$

that is accurate to the first order, given  $z_{k+1} = m_\phi(z_k, v_k)$ . The matrix  $F_k = A_k + B_k K$  describes the closed loop error system. To account for state-action dependent uncertainties  $S_\phi$ , we combine the nonlinear nominal trajectory with the linearized error dynamics. A one-step ellipsoidal forward propagation is computed by

$$\mathcal{E}_{k+1} = g_\phi(\mathcal{E}_k, v_k)$$

with the non-linear mapping

$$g_\phi(\mathcal{E}_k, v_k) = \mathcal{E}\left(m_\phi(z_k, v_k), F_k S_k F_k^T\right) \oplus \mathcal{E}\left(0, S_\phi(z_k, v_k)\right),$$

where the evolution of  $\mathcal{E}_k$ 's only depends on the action sequence  $(v_0, \dots, v_{N-1})$  and the NN parameters  $\phi$ . Note that  $\mathcal{E}(z_k, S_k)$  is abbreviated to  $\mathcal{E}_k$  to improve readability. An illustration of the one-step uncertainty propagation is depicted in Figure 1 (Left).

### 4.3 Ensemble Uncertainty Propagation

A frequent problem in model-based RL is that NN predictions exhibit inaccuracies that grow with the length of the predictive horizon, limiting the applicability to short rollouts [25]. Ensembles, however, demonstrated their effectiveness in preventing the exploitation of inaccuracies during planning [10]. Therefore, we adopt an NN ensemble  $\tilde{f}_\phi = \{f_{\phi_1}, \dots, f_{\phi_M}\}$  composed of  $M$  models. The ellipsoidal uncertainty propagation for each model  $i$  is described by

$$\mathcal{E}_{k+1}^{(i)} = g_{\phi_i}\left(\mathcal{E}_k^{(i)}, v_k\right),$$

resulting in  $M$  tubes used for planning.

### 4.4 Safety Certification

In this section, we introduce the optimization problem that is solved by X-MPSC. More formally, X-MPSC extends nominal MPSC from (2) with tube-based MPC and a probabilistic ensemble of NNs (the modified parts are highlighted in blue color). The goal of X-MPSC is to certify an action  $u_t$  proposed by an RL-based policy in each time step  $t$  by solving the optimization problem

$$\begin{aligned} &\text{minimize} && \|u_t - v_0\|_2^2 \\ &v_0, \dots, v_{N-1} \\ &\text{subject to} && \mathcal{E}_0^{(i)} = \mathcal{E}(x_t, 0) \quad \forall i \\ &&& \mathcal{E}_{k+1}^{(i)} = g_{\phi_i}(\mathcal{E}_k^{(i)}, v_k) \quad \forall k \in [0, N-1], i \\ &&& v_k \in \tilde{\mathbb{U}}(\mathcal{E}_k^{(i)}) \quad \forall k \in [0, N-1], i \\ &&& \mathcal{E}_k^{(i)} \subseteq \mathbb{X} \quad \forall k \in [0, N], i \\ &&& \mathcal{E}_N^{(i)} \subseteq \mathbb{X}_{\text{term}} \quad \forall i \end{aligned} \quad (8)$$

over the horizon  $N$ . The objective is solved in a receding horizon fashion, where a safe action  $v_0$  that deviates as minimally as possible from the learning input is obtained. For (8) to be feasible, a series of constraints must be satisfied. All propagated ellipsoids must be contained in the polytopic state space, which is checked through (5). Further, every member of the ensemble is subject to the constraints. If a single member violates a constraint, the optimization problem becomes infeasible. The final ellipsoid of each tube is required to be contained in the terminal set  $\mathbb{X}_{\text{term}}$ . Depending on the ellipsoidal uncertainty estimates, the set of feasible actions shrinks  $v_k \in \tilde{\mathbb{U}}(\mathcal{E}_k) = \bigcup \mathcal{E}(0, K S_k K^T)$ .

When (8) is feasible, we do not only obtain a safe action but also a sequence of feedback policies  $\Pi_t = (\pi_t, \pi_{t+1}, \dots, \pi_{t+N-1})$  that can steer the system back to  $\mathbb{X}_{\text{term}}$  within  $N$  steps. Here, each  $\pi_{t+k}(x_{t+k}) = v_k + K(x_{t+k} - z_k)$  is an affine controller that tracks the actual system toward the nominal trajectory. In case of infeasibility, the solution of the former solver iteration  $\Pi_{t-1}$  is reused. An

illustration of the optimization problem and the safety certification can be seen in Figure 1 (Right).

#### 4.5 Proposed Algorithm

The solution to the optimization problem in (8) provides a safe action in each step. However, to provide safe exploration for an RL agent throughout the training, we rely on certain assumptions.

**Assumption 4.1.** There exists a safe backup policy  $\pi_b \in \Pi_b$  such that when following  $\pi_b$  the states

$$x_t \in \mathbb{S} \Rightarrow x_{t+k} \in \mathbb{X} \quad \forall k > 0$$

are contained in  $\mathbb{X}$ . The set  $\mathbb{S}$  is called safe set.

The safe set  $\mathbb{S}$  is a control-forward invariant set that allows us to gather safely offline data when having access to a safe backup controller. We utilize offline data collected by such a safe backup controller for pre-training the ensemble of dynamics model  $\tilde{f}_\phi$  and the initial policy  $\pi_\theta$  before starting the RL training. In practice, we can satisfy this assumption through a simple stabilizing local controller, which has low task performance but keeps the system safe within a small region of the state space.

**Assumption 4.2.** The set of initial states is inscribed in the safe set, *i.e.*,  $\mathbb{X}_0 \subseteq \mathbb{S}$ . Also, the terminal set is a subset of the safe set, *i.e.*,  $\mathbb{X}_{\text{term}} \subseteq \mathbb{S}$ . Both sets are convex.

Since all initial states lie in the safe set, a safe backup policy is able to keep the system safe for all future time steps.

**Assumption 4.3.** The actual system (1) underlies bounded disturbances and is Lipschitz continuous.

The predictions of the dynamics models (6) can be also bounded by this assumption.

**Assumption 4.4.** The ensemble of dynamics models  $\tilde{f}_\phi$  is sufficiently accurate such that it captures the trajectories of the actual system (1).

A trajectory is captured by the ensemble when all states of the actual system are contained in any of the predicted ellipsoids, *i.e.*,

$$\forall k \in [0, N] \exists i \text{ s. t. } x_{t+k} \in \mathcal{E}_k^{(i)}.$$

Even though a state  $x_{t+k}$  is not contained in any of the predicted ellipsoids but lies in between the tube-based rollouts, *i.e.*,

$$x_{t+k} \in \text{conv} \left( \bigcup_i \mathcal{E}_k^{(i)} \right),$$

safety constraints can be still enforced. We consider the ensemble of dynamics to be sufficiently accurate in such a scenario. The assumption of an accurate dynamics model is very strong since the model can only be a good approximation in regions where data is available. Since the terminal set acts as a natural regularizer to the exploration in (8), *i.e.*, the agent must reach  $\mathbb{X}_{\text{term}}$  within  $N$  steps, exploration relies on the terminal set growing throughout training to acquire novel samples. However, the growth must happen at an appropriate speed so that new data is informative (*i.e.*, from regions where prediction uncertainty is high) and safe (the system can be steered back to the terminal set).

---

#### Algorithm 1 Safe Reinforcement Learning with X-MPSC

---

```

1: Input: Initial data  $\mathbb{D}_0$  collected by a safe policy  $\pi_b$  (and optionally a prior model  $f_{\text{prior}}$ )
2: Pre-train  $\pi_\theta$  on  $\mathbb{D}_0$  and set  $\mathbb{D} \leftarrow \mathbb{D}_0$ 
3: for epoch  $j = 1, \dots$  do
4:   Train actor-critic and ensemble model  $\tilde{f}_\phi$  via (7) on  $\mathbb{D}$ 
5:   Estimate  $\tilde{\mathbb{S}}_j$  based on (9) and set  $\mathbb{X}_{\text{term}} \leftarrow \tilde{\mathbb{S}}_{j-\delta}$ 
6:   for time step  $t = 1, \dots$  do
7:     Sample (possibly unsafe)  $u_t \sim \pi_\theta(x_t)$  from RL policy
8:     Obtain (feasible,  $\Pi$ ) by solving X-MPSC problem (8)
9:     Retrieve sequence of controllers
10:     $\Pi_t \leftarrow \begin{cases} \Pi = (\pi_t, \pi_{t+1}, \dots, \pi_{t+N-1}) & \text{if feasible} \\ \Pi_{t-1} = (\pi_t, \dots, \pi_{t+N-2}) & \text{otherwise} \end{cases}$ 
11:    Get safe action  $v_t = \pi_t(x_t)$  and apply  $v_t$  to system (1)
12:    Store  $\mathbb{D} \leftarrow \mathbb{D} \cup (x_t, v_t, x_{t+1}, \text{feasible})$ 
13:   end for
14: end for

```

---

We claim that when the Assumptions 4.1–4.4 are fulfilled, then Algorithm 1 becomes itself a safe backup policy due to its recursive feasibility and, hence, can safely certify the actions of an RL-based agent. In the next paragraph, we will give an intuition of this claim and show empirical evidence with the experiments conducted in Section 6. Our method is summarized in Algorithm 1.

At the beginning of each episode, we obtain  $x_0 \in \mathbb{X}_0$ . By Assumptions 4.1 and 4.2, there always exists a solution  $\Pi_0 = \{\pi_b, \dots, \pi_b\}$  at  $t = 0$  since  $x_0 \in \mathbb{X}_0 \subseteq \mathbb{S} \Rightarrow x_t \in \mathbb{X} \quad \forall t > 0$  when following a safe backup controller  $\pi_b \in \Pi_b$ . We can now show by induction that the system can be kept safe  $\forall t \geq 0$ . Let the previous step  $t - 1$  have the feasible solution  $\Pi_{t-1} = (\pi_{t-1}, \pi_t, \pi_{t+1}, \dots)$  that holds the system safe for all future time steps  $t \geq 0$ . Then, there exists also a solution at step  $t$  because either X-MPSC will find it by solving (8) or the solution  $\Pi_{t-1}$  from the former step gives a sequence of policies  $(\pi_t, \pi_{t+1}, \dots, \pi_{t+N-2})$  as a fallback solution that leads to a safe state  $x_N \in \mathbb{X}_{\text{term}} \subseteq \mathbb{S}$  within  $N - 1$  steps, from where a safe backup controller  $\pi_b$  can keep the system safe.

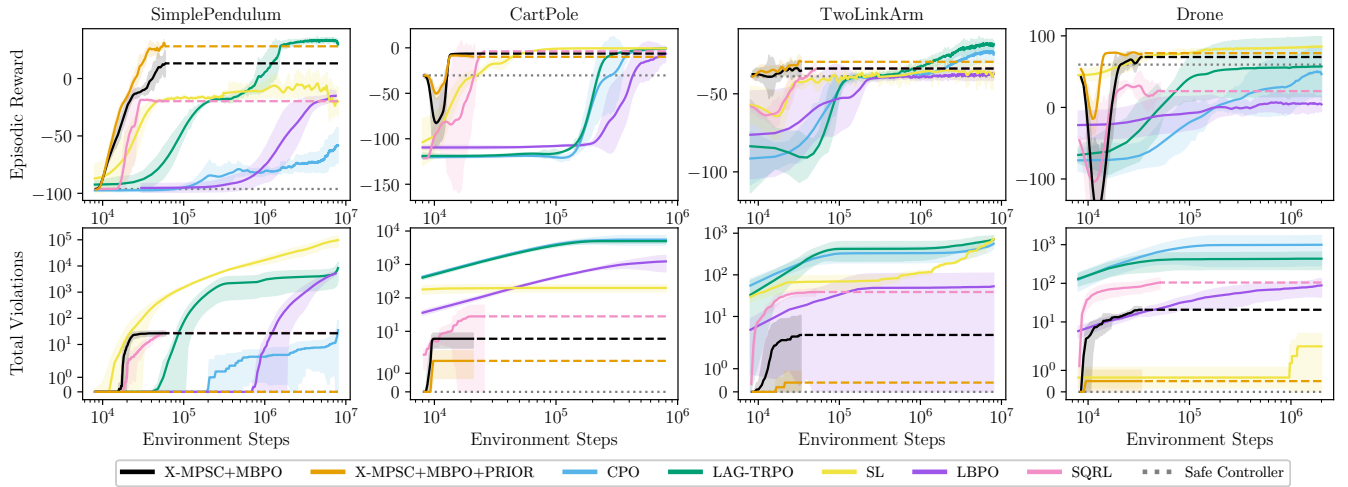
## 5 PRACTICAL IMPLEMENTATION

In the previous section, we did not specify how the terminal set  $\mathbb{X}_{\text{term}}$  can be obtained nor how the safe set  $\mathbb{S}$  can be determined. In this section, we elaborate on the estimation of both sets and introduce prior system models that we used to improve the accuracy of the NN dynamics models.

*Estimation of the Safe Set.* The safe set  $\mathbb{S}$  is difficult to compute in practice. However, from the data  $\mathbb{D}$  collected so far over the training, we can build the outer approximation

$$\tilde{\mathbb{S}} = \text{conv} \{x_t \mid (x_t, \cdot, \cdot, \text{feasible} = \text{true}) \in \mathbb{D}\} \quad (9)$$

as convex hull over all states where X-MPSC found a solution to (8), which is polytopic and convex. Here we use the fact that all states  $x_t$  kept the system within  $\mathbb{X}$  under the model  $\tilde{f}_\phi$ . Note that  $\tilde{\mathbb{S}}$  is just an approximation of  $\mathbb{S}$  and changes as soon as new data samples are collected in the training process.



**Figure 2: Experimental results. Thick lines show the average over five independent seeds and the shaded area denotes the standard deviation. (Top) The cumulative reward of one episode reported over the total environment steps. (Bottom) Total constraint violations over the whole training.**

*Terminal Set.* Since the terminal set naturally limits exploration, the terminal set is supposed to grow throughout the training to learn about the areas of the state space that have not been visited yet. Due to Assumption 4.2, we can set  $\mathbb{X}_{\text{term}} \leftarrow \mathbb{X}_0$  as the first choice for the terminal set. The initial state distribution can be estimated by building the convex hull of all initial states from  $\mathbb{D}_0$ . As new samples are collected, the safe set  $\tilde{\mathbb{S}}_j$  is re-estimated at each epoch  $j$ . In order to prevent model exploitation, we use a safe set estimation of a former epoch, i.e.,  $\mathbb{X}_{\text{term}} \leftarrow \tilde{\mathbb{S}}_{j-\delta}$  such that the estimated safe set is delayed by  $\delta$  epochs.

*Handling of Infeasibility.* Due to divergence in the model ensemble predictions or large uncertainty ellipsoids, the X-MPSC problem might not always be feasible, i.e., the solver cannot find an action sequence that keeps all tubes within the constraints. In such cases, the solution  $\Pi_{t-1}$  found in the former solver iteration returns a policy sequence that steers the system back to  $\mathbb{X}_{\text{term}}$  in  $N-1$  steps. Safety can still be guaranteed due to recursive feasibility. In practice, however, a single failure event can lead to a series of infeasibility events. When the solver fails to find a solution to (8) in  $N$  consecutive steps, we transform the hard constraints to soft constraints with high penalty terms. The number of infeasibility events depends on the selected hyper-parameters and varied between 0.0% (for the best seeds) and 2.0% (worst case) of the time steps. However, in the TwoLinkArm task, we could also observe a worst-case failure rate of approximately 39%.

*Prior Model.* To improve the validity of Assumption 4.4, we tested the use of a prior model. Thus, we extended each ensemble member

$$f_{\phi_i}(x_t, u_t) = \mathcal{N}(m_{\phi_i}(x, u), S_{\phi_i}(x, u)) + f_{\text{prior}}$$

with an additive component  $f_{\text{prior}}$  that is derived from first principles. We set the system parameters of the prior model with an error of 20% (offset) compared to the actual system’s parameters of (1).

## 6 EXPERIMENTS

Our experimental evaluation is intended to give empirical evidence that X-MPSC can certify the actions taken by an RL agent and that the Assumptions 4.1–4.4 apply to typical RL problem settings. The software implementation is published on GitHub and can be found at: <https://github.com/SvenGronauer/x-mpsc>.

### 6.1 Environments

We tested X-MPSC on four tasks that differ in complexity and dynamics. (1) *Simple Pendulum* ( $\mathbb{X} \subset \mathbb{R}^2, \mathbb{U} \subset \mathbb{R}$ ) describes a swing-up task with restricted angle and input constraints. (2) In *Cart Pole* ( $\mathbb{X} \subset \mathbb{R}^4, \mathbb{U} \subset \mathbb{R}$ ), the agent is supposed to balance the pole in an upright position without violating cart position and pole angle constraints. (3) *Two-Link-Arm* ( $\mathbb{X} \subset \mathbb{R}^8, \mathbb{U} \subset \mathbb{R}^2$ ) is a two-joint manipulator where a target point should be reached with end-effector position limits. (4) The *Drone* ( $\mathbb{X} \subset \mathbb{R}^{12}, \mathbb{U} \subset \mathbb{R}^4$ ) environment describes the task to take off the ground and fly to position  $[0, 0, 1]^T$  with a restricted body angle. All environments terminated early when state constraint violations occurred. We denote the episodic reward as task performance while we refer to the total number of constraint violations as safety performance.

### 6.2 Algorithm Setup

In general, X-MPSC can be combined with any RL algorithm. For our experiments, we chose model-based policy optimization (MBPO) proposed by Janner et al. [15] for two reasons. First, we can use the NN ensemble for both planning with X-MPSC and for generating short model-based rollouts to train the policy. Second, MBPO has been shown to be more sample efficient than other RL algorithms, which reduces the wall clock time since the computational bottleneck is finding a solution to (8). For training, we used an ensemble size of  $M = 5$ , where each NN was implemented as a multi-layer

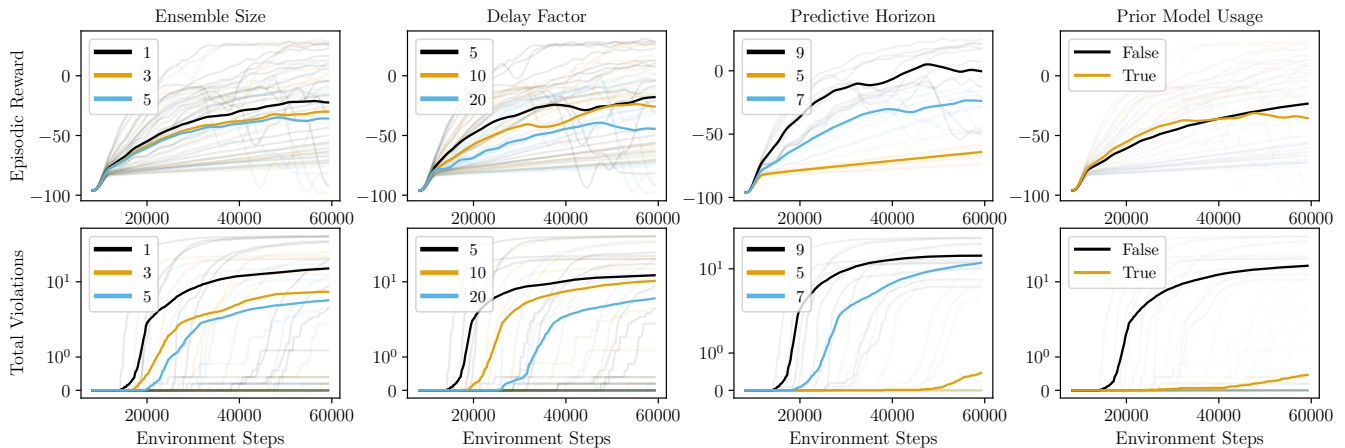


Figure 3: Impact of X-MPSC hyper-parameters on safety and performance in Simple Pendulum.

perceptron (MLP) with two hidden layers. The actor-critics were also MLPs with two hidden layers.

The optimization problem in (8) was solved with a primal-dual interior point method, for which we used *CasADi* and the *IPOPT* software library. Two relevant hyper-parameters of X-MPSC are the delay factor  $\delta$  and the horizon  $N$ , which depend on the characteristics of the environment. We adjusted both hyper-parameters for each environment individually. The feedback matrix  $K$  was hand-tuned, and all matrix entries were set to 0.5 for all environments. We expect that tuning this hyper-parameter individually for each environment will reduce the number of constraint violations but did not test this. An overview of all selected hyper-parameters can be found in the Appendix at: <https://arxiv.org/abs/2402.04182>.

We collected for each task  $|\mathbb{D}_0| = 8000$  initial samples, which equals less than three minutes of real-world experience on systems with a 50 Hz sampling rate. The safe backup controller used for initial data collection is specific to the environment dynamics and was implemented by a low-performing but stabilizing linear quadratic regulator (LQR) or proportional–integral–derivative (PID) controller. We compare our results with several baselines, namely constrained policy optimization (CPO) [1], safety Q-functions for RL (SQRL) [33], safety layer (SL) [11], Lyapunov barrier policy optimization (LBPO) [31], and Lagrangian trust-region policy optimization (Lag-TRPO). For SL and SQRL, we collected initial data samples that deliberately contained safety violations to pre-train their safety-aware functions.

### 6.3 Results

Figure 2 depicts the episodic reward as well as the total number of constraint violations over the course of training for each algorithm and environment. For each algorithm, we identified the best hyper-parameter choice via a coarse grid search and report only the best configuration. Each experiment was averaged over five independent random seeds. Additionally, we report the performance of the safe controller that was used to collect the initial data for X-MPSC. Note that MBPO and SQRL converge faster due to their off-policy nature.

We observe that, even without using a prior model, X-MPSC can significantly reduce constraint violations compared to the other algorithms while achieving only a slightly worse final reward than Lag-TRPO, the strongest baseline in terms of task performance. The on-policy algorithms CPO, SL, and LBPO show less consistent results across the experiments, with SL demonstrating comparable performance to X-MPSC only in the Drone task. The off-policy SQRL displays a good performance–safety ratio by learning policies with fewer violations than the other algorithms in most cases, excluding X-MPSC. However, when an inaccurate prior model is added to the NN ensemble, the total constraint violations with the X-MPSC can be reduced by approximately an order of magnitude without performance losses in terms of episodic reward.

### 6.4 Impact of Hyper-Parameters

We studied the impact of selected X-MPSC hyper-parameters on performance and constraint violations in the Simple Pendulum task. Since hyper-parameter settings can cross-correlate and influence each other, we deliberately do not fix all hyper-parameters except the one of interest to avoid cherry-picking good hyper-parameter settings and distorting the general impact of the selected hyper-parameter. Instead, we are interested in the average effect of the hyper-parameter choice and, thus, measure the impact over multiple random seeds and various hyper-parameter settings. To this end, we tested 54 different hyper-parameter configurations, where each configuration was averaged over five independent trials.

Figure 3 shows each configuration in pale color, while the average of the selected hyper-parameter is shown as a solid line. The plots indicate that a larger ensemble size results in fewer total constraint violations. Further, the usage of a coarse prior model can reduce the number of constraint violations on average by one order of magnitude. The delay factor regulates the speed of the terminal set updates, with a higher factor diminishing the number of violations at the expense of fewer cumulative rewards. Finally, a longer predictive horizon  $N$  can accelerate the learning progress but produces more costs.

## 7 DISCUSSION

The results show that our proposed X-MPSC algorithm is able to provide safe exploration when certain assumptions are fulfilled. We discuss our experimental results based on those assumptions first and then give reasons for the success or failure of X-MPSC. After that, we discuss the most critical limitations of our method.

### 7.1 Discussion of Results

*Safety.* As shown in Figure 2, our method offers a better safety operation in terms of constraint satisfaction, with the violations kept at zero in the Simple Pendulum and nearly zero in the Drone environment when prior knowledge is used. Note that the Cart Pole task is particularly challenging since the starting pole upward position is an unstable equilibrium point, and expanding the safe set will rapidly reach states where the controller is not yet able to stabilize it. X-MPSC was able to learn a safe policy faster, but collecting the data necessary to approximate the dynamics model around unstable regions is a challenge that remains unsolved.

*Performance.* A safety-performance trade-off is reflected in the results. Slowly expanding the safe set and using ensembles based on tube-based MPC, which ensures that all tubes are contained in the safety constraints, result in a conservative system. X-MPSC’s performance without prior model is relatively close to the best baseline algorithms in the Cart Pole and Drone environments but presents a more significant difference to Lag-TRPO in the other two environments. When adding a prior dynamics model, we do not only see a significant decrease in constraint violations, but also slight improvements in terms of task reward in the Drone and Two-Link-Arm environments. In the Simple Pendulum environment this increase was even more substantial. We argue that the usage of a prior model leads to more accurate nominal trajectories, which facilitates learning in terms of reward performance as well as improved safety satisfaction capabilities.

### 7.2 Limitations of Our Method

In the remainder of this section, we list the limitations of the X-MPSC algorithm, which are ordered from weak to strong. We see these limitations as a starting point to be addressed in future work.

*Conservatism.* With a larger ensemble size, the safety certification can lead to more conservative behavior since the constraints imposed by every single dynamics model must be satisfied. As soon as a single model deviates from the rest of the ensemble, the agent is enforced to satisfy wrong constraints, and, hence, the set of feasible actions is reduced, which can result in diminished reward performance. With an increasing number of models, the safety certification becomes safer but also more conservative.

*Safe Backup Controller.* Our algorithm only requires offline data to be collected within the safe set to pre-train the ensemble and estimate the initial safe and terminal sets. In contrast to our work, related methods require offline data to contain mixed safe and unsafe trajectories (e.g., SL [11], SQRL [33] and Recovery RL [35]), which can be a limiting factor for the deployment on real-world robots. In practice, the safe backup controller can be implemented as a local control law that has low task performance but keeps the

system close to the initial state and within the safe set. The time required to develop a safe backup controller largely depends on the task and the robot dynamics. For CartPole and Simple Pendulum, we used an LQR, while we used P-controllers for Drone and Two-Link-Arm. Because a safe backup controller does not aim for good task performance but only for stabilizing within a small region of the state space, the design can be achieved with a few trial-and-error attempts. Conversely, the design of a safe backup controller with high reward performance can take considerably more time and requires accurate prior knowledge about the robot system. X-MPSC can be used, however, with relatively little prior knowledge about the system and can thus offer an approach to safely learn about the system while being able to improve task performance.

*Accurate Dynamics Model.* The access to an accurate model is a strong assumption since the model is only a good approximation in regions where data samples are available. Through the terminal set constraint, we naturally limit the exploration of the state space since the terminal set must be reached within  $N$  steps and, hence, force the agent to stay close to regions where data exists. Incorporating prior models helped significantly reduce the number of constraint violations, although we only used inaccurate models with parameters deviating by 20% from the ground-truth.

*Computation Time.* The bottleneck in the training loop is finding a solution to (8). The X-MPSC problem has cubic computational complexity, i.e.,  $\mathcal{O}(N^3(n_x + n_u + n_c)^3)$  with  $N$  being the predictive horizon and  $n_x, n_u, n_c$  being the dimensions of the state, action spaces and number of constraints, respectively. Note that the computational complexity grows linearly with the ensemble size  $M$  and quadratically with the number of neurons in each layer. Thus, we use at most 20 neurons in the hidden layers of the NNs. Furthermore, the estimation of the safe set involves the computation of the convex hull, which also grows with the number of state constraints. Thus, we limited the number of constrained state space variables to two, e.g., position and velocity.

## 8 CONCLUSION

We proposed X-MPSC, a novel algorithm that integrates an ensemble of NNs and tube-based MPC into nominal MPSC to correct the actions taken by an RL agent. To provide safe exploration, we utilized an ensemble of probabilistic NNs trained on sampled environment data to plan multiple tube-based trajectories that satisfy *a priori* defined safety constraints. The experimental results demonstrate that our method can achieve significantly fewer constraint violations than comparable RL methods, requiring only offline data with safe trajectories. When an inaccurate prior dynamics model is added to the NN ensemble, the constraint violations can be reduced by an order of magnitude without forfeiting reward performance.

Although the results indicate an improvement over other state-of-the-art algorithms, the scalability to higher dimensional state-action spaces and larger NN models remains a research frontier. Future work may improve upon the computational speed of solving the X-MPSC problem such that it can be implemented on a real robot or applied to high-dimensional tasks, as presented in the *Safety Gym* [29] or *Bullet Safety Gym* [14].



## REFERENCES

- [1] Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. 2017. Constrained Policy Optimization. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, Doina Precup and Yee Whye Teh (Eds.), 22–31. <http://proceedings.mlr.press/v70/achiam17a.html>
- [2] Eitan Altman. 1999. *Constrained Markov decision processes*. CRC Press.
- [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 (ECC)*. 3420–3431.
- [4] Dario Amodei, Chris Olah, Jacob Steinhardt, Paul F. Christiano, John Schulman, and Dan Mané. 2016. Concrete Problems in AI Safety. Preprint. <https://doi.org/10.48550/arXiv.1606.06565>
- [5] Felix Berkenkamp, Matteo Turchetta, Angela Schoellig, and Andreas Krause. 2017. Safe Model-based Reinforcement Learning with Stability Guarantees. In *Advances in Neural Information Processing Systems 30 (NeurIPS)*. 908–918.
- [6] Lukas Brunke, Melissa Greeff, Adam W. Hall, Zhaocong Yuan, Siqi Zhou, Jacopo Panerati, and Angela P. Schoellig. 2022. Safe Learning in Robotics: From Learning-Based Control to Safe Reinforcement Learning. *Annual Review of Control, Robotics, and Autonomous Systems* 5, 1 (2022), 411–444. <https://doi.org/10.1146/annurev-control-042920-020211> arXiv:<https://doi.org/10.1146/annurev-control-042920-020211>
- [7] 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 33rd AAAI Conference on Artificial Intelligence (AAAI, Vol. 33)*. 3387–3395. <https://doi.org/10.1609/aaai.v33i01.33013387>
- [8] Yinlam Chow, Mohammad Ghavamzadeh, Lucas Janson, and Marco Pavone. 2017. Risk-constrained reinforcement learning with percentile risk criteria. *The Journal of Machine Learning Research* 18, 1 (2017), 6070–6120.
- [9] Yinlam Chow, Ofir Nachum, Aleksandra Faust, Edgar Duenez-Guzman, and Mohammad Ghavamzadeh. 2019. Lyapunov-based safe policy optimization for continuous control. *arXiv preprint arXiv:1901.10031* (2019).
- [10] Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. 2018. Deep Reinforcement Learning in a Handful of Trials Using Probabilistic Dynamics Models. In *Advances in Neural Information Processing Systems 31 (NeurIPS)*. 4759–4770.
- [11] Gal Dalal, Krishnamurthy Dvijotham, Matej Vecerik, Todd Hester, Cosmin Paduraru, and Yuval Tassa. 2018. Safe Exploration in Continuous Action Spaces. Preprint. <https://doi.org/10.48550/arXiv.1801.08757> arXiv:1801.08757
- [12] Gabriel Dulac-Arnold, Nir Levine, Daniel J. Mankowitz, Jerry Li, Cosmin Paduraru, Sven Gowal, and Todd Hester. 2021. Challenges of real-world reinforcement learning: definitions, benchmarks and analysis. *Machine Learning* 110, 9 (2021), 2419–2468. <https://doi.org/10.1007/s10994-021-05961-4>
- [13] Javier García, Fern, and o Fernández. 2015. A Comprehensive Survey on Safe Reinforcement Learning. *Journal of Machine Learning Research* 16, 42 (2015), 1437–1480.
- [14] Sven Gronauer. 2022. *Bullet-Safety-Gym: A Framework for Constrained Reinforcement Learning*. Technical Report. mediaTUM. <https://doi.org/10.14459/2022md1639974>
- [15] Michael Janner, Justin Fu, Marvin Zhang, and Sergey Levine. 2019. When to Trust Your Model: Model-Based Policy Optimization. In *Advances in Neural Information Processing Systems 32 (NeurIPS)*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (Eds.), 12519–12530.
- [16] Torsten Koller, Felix Berkenkamp, Matteo Turchetta, and Andreas Krause. 2018. Learning-Based Model Predictive Control for Safe Exploration. In *IEEE Conference on Decision and Control (CDC)*. 6059–6066. <https://doi.org/10.1109/CDC.2018.8619572>
- [17] Alex A. Kurzhanskiy and Pravin Varaiya. 2006. Ellipsoidal Toolbox (ET). In *IEEE Conference on Decision and Control (CDC, Vol. 45)*. 1498–1503. <https://doi.org/10.1109/CDC.2006.377036>
- [18] Jan Leike, Miljan Martic, Victoria Krakovna, Pedro A. Ortega, Tom Everitt, Andrew Lefrancq, Laurent Orseau, and Shane Legg. 2017. AI Safety Gridworlds. Preprint. <https://doi.org/10.48550/arXiv.1711.09883> arXiv:1711.09883
- [19] Weiwei Li and Emanuel Todorov. 2004. Iterative Linear Quadratic Regulator Design for Nonlinear Biological Movement Systems. In *International Conference on Informatics in Control, Automation and Robotics*.
- [20] Qingkai Liang, Fanyu Que, and Eytan Modiano. 2018. Accelerated primal-dual policy optimization for safe reinforcement learning. *arXiv:1802.06480* (2018).
- [21] Yuping Luo and Tengyu Ma. 2021. Learning Barrier Certificates: Towards Safe Reinforcement Learning with Zero Training-time Violations. In *Advances in Neural Information Processing Systems 34 (NeurIPS)*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (Eds.), 25621–25632. [https://proceedings.neurips.cc/paper\\_files/paper/2021/file/d71fa38b648d86602d14ac610f2e6194-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2021/file/d71fa38b648d86602d14ac610f2e6194-Paper.pdf)
- [22] Björn Lütjens, Michael Everett, and Jonathan P. How. 2019. Safe Reinforcement Learning With Model Uncertainty Estimates. In *International Conference on Robotics and Automation (ICRA)*. 8662–8668. <https://doi.org/10.1109/ICRA.2019.8793611>
- [23] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemaire, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharmash Kumar, Daan Wierstra, Shane Legg, and Demis Hassabis. 2015. Human-level control through deep reinforcement learning. *Nature* 518 (2015), 529 EP. <https://doi.org/10.1038/nature14236>
- [24] Teodor Mihai Moldovan and Pieter Abbeel. 2012. Safe Exploration in Markov Decision Processes. In *Proceedings of the 29th International Conference on Machine Learning (ICML)*. 1451–1458.
- [25] Anusha Nagabandi, Gregory Kahn, Ronald S. Fearing, and Sergey Levine. 2018. Neural Network Dynamics for Model-Based Deep Reinforcement Learning with Model-Free Fine-Tuning. In *IEEE International Conference on Robotics and Automation (ICRA)*. 7559–7566. <https://doi.org/10.1109/ICRA.2018.8463189>
- [26] Martin Pecka and Tomas Svoboda. 2014. Safe Exploration Techniques for Reinforcement Learning – An Overview. In *Modelling and Simulation for Autonomous Systems (MESAS)*, Jan Hodičky (Ed.). Springer International Publishing, Cham, 357–375.
- [27] Samuel Pfrommer, Tanmay Gautam, Alec Zhou, and Somayeh Sojoudi. 2022. Safe Reinforcement Learning with Chance-constrained Model Predictive Control. In *Proceedings of The 4th Annual Learning for Dynamics and Control Conference*, Vol. 168. PMLR, 291–303.
- [28] James Blake Rawlings, David Q Mayne, and Moritz Diehl. 2017. *Model Predictive Control: Theory, Computation, and Design* (2nd edition ed.). Nob Hill Publishing, LLC.
- [29] Alex Ray, Joshua Achiam, and Dario Amodei. 2019. Benchmarking Safe Exploration in Deep Reinforcement Learning. Preprint. <https://cdn.openai.com/safexp-short.pdf>
- [30] Alexander Robey, Haimin Hu, Lars Lindemann, Hanwen Zhang, Dimos V Dimarogonas, Stephen Tu, and Nikolai Matni. 2020. Learning control barrier functions from expert demonstrations. In *2020 59th IEEE Conference on Decision and Control (CDC)*. IEEE, 3717–3724.
- [31] Harshit Sikchi, Wenxuan Zhou, and David Held. 2021. Lyapunov Barrier Policy Optimization. *CoRR* abs/2103.09230 (2021). arXiv:2103.09230
- [32] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature* 529, 7587 (2016), 484–489. <https://doi.org/10.1038/nature16961>
- [33] Krishnan Srinivasan, Benjamin Eysenbach, Sehoon Ha, Jie Tan, and Chelsea Finn. 2020. Learning to be safe: Deep rl with a safety critic. Preprint. <https://doi.org/10.48550/arXiv.2010.14603>
- [34] Chen Tessler, Daniel J. Mankowitz, and Shie Mannor. 2019. Reward Constrained Policy Optimization. In *7th International Conference on Learning Representations (ICLR)*.
- [35] Brijen Thananjeyan, Ashwin Balakrishna, Suraj Nair, Michael Luo, Krishnan Srinivasan, Minh Hwang, Joseph E. Gonzalez, Julian Ibarz, Chelsea Finn, and Ken Goldberg. 2021. Recovery RL: Safe Reinforcement Learning With Learned Recovery Zones. *IEEE Robotics and Automation Letters* 6, 3 (2021), 4915–4922. <https://doi.org/10.1109/LRA.2021.3070252>
- [36] Kim P. Wabersich, Lukas Hewing, Andrea Carron, and Melanie N. Zeilinger. 2022. Probabilistic Model Predictive Safety Certification for Learning-Based Control. *IEEE Trans. Automat. Control* 67, 1 (2022), 176–188. <https://doi.org/10.1109/TAC.2021.3049335>
- [37] K. P. Wabersich and M. N. Zeilinger. 2018. Linear Model Predictive Safety Certification for Learning-Based Control. In *IEEE Conference on Decision and Control (CDC)*. 7130–7135. <https://doi.org/10.1109/CDC.2018.8619829>
- [38] Kim Peter Wabersich and Melanie N. Zeilinger. 2021. A predictive safety filter for learning-based control of constrained nonlinear dynamical systems. *Automatica* 129 (2021), 109597. <https://doi.org/10.1016/j.automatica.2021.109597>
- [39] Tsung-Yen Yang, Justinian Rosca, Karthik Narasimhan, and Peter J. Ramadge. 2020. Projection-Based Constrained Policy Optimization. In *International Conference on Learning Representations*.