

Sample-efficient Safe Learning for Online Nonlinear Control with Control Barrier Functions

Wenhao Luo¹, Wen Sun², and Ashish Kapoor³

¹ Department of Computer Science, University of North Carolina at Charlotte,
Charlotte NC 28223, USA wenhao.luo@uncc.edu

² Computer Science Department, Cornell University, Ithaca, NY 14853 USA
ws455@cornell.edu

³ Microsoft Corporation, Redmond, Washington 98033 USA
akapoor@microsoft.com

Abstract. Reinforcement Learning (RL) and continuous nonlinear control have been successfully deployed in multiple domains of complicated sequential decision-making tasks. However, given the exploration nature of the learning process and the presence of model uncertainty, it is challenging to apply them to safety-critical control tasks due to the lack of safety guarantee. On the other hand, while combining control-theoretical approaches with learning algorithms has shown promise in safe RL applications, the sample efficiency of safe data collection process for control is not well addressed. In this paper, we propose a *provably* sample efficient episodic safe learning framework for online control tasks that leverages safe exploration and exploitation in an unknown, nonlinear dynamical system. In particular, the framework 1) extends control barrier functions (CBFs) in a stochastic setting to achieve provable high-probability safety under uncertainty during model learning and 2) integrates an optimism-based exploration strategy to efficiently guide the safe exploration process with learned dynamics for *near optimal* control performance. We provide formal analysis on the episodic regret bound against the optimal controller and probabilistic safety with theoretical guarantees. Simulation results are provided to demonstrate the effectiveness and efficiency of the proposed algorithm.

Keywords: Robot Safety · Safe reinforcement learning · Safe control.

1 Introduction

The control of safety-critical systems such robotic systems is a difficult challenge under uncertainty and lack of complete information in the real world applications. While Reinforcement Learning (RL) algorithms for long-term reward maximization have achieved significant results in many continuous control tasks [12, 29], they have not yet been widely applied to safety-critical control tasks as the rigorous safety requirements may be easily violated by intermediate policies during policy learning. Safe RL approaches [1, 5, 14, 25] with constraints satisfaction have been proposed to encode safety consideration in a modified optimality criterion

or in the constrained policy exploration process with external knowledge, e.g. an accurate probabilistic system model [25, 34].

On the other hand, model-based approaches utilizing Model Predictive Control (MPC) or Lyapunov-based methods have seen a number of success in demanding control tasks under different constraints with accurate system models [17, 39, 40] or approximated dynamics [38]. Meanwhile, enforcing safety in terms of *set forward invariance* has become an active research area using Lyapunov functions [6] and control barrier functions (CBFs) with perfectly known system models [3, 4, 39] or noisy models with known distributions [11, 15, 21, 22, 33]. However, these control-based approaches still require known system model uncertainty and could be overly conservative for system behaviors in presence of large uncertainty.

For this purpose, integrating data-driven learning-based approaches with model-based safe control has received significant attention to achieve model uncertainty reduction while ensuring provable safety [6, 8–10, 13, 19, 20, 26, 32, 37]. For example, standard CBFs [37] and robust CBFs [8] have been employed in conjunction with Gaussian Processes to ensure robot safety with uncertainty learned online. The process often involves safe policy exploration with data collection from a nominal dynamics model and iteratively reduce learned model uncertainty over time to expand certified safety region of the system’s state space [6, 10, 19, 32, 37]. However, such exhaustive data collection for safe learning could suffer from poor scalability and low efficiency for primary task. For example, instead of densely sampling over the space, it may be more beneficial to guide the safe exploration process towards task-prescribed policy optimization. Recent work [9] incorporates safe learning using Gaussian Process (GP) and CBF into a model-free RL framework (RL-CBF), so that the guided exploration process will not only learn model uncertainty impacting safe behaviors but also optimizing the policy performance. Similar idea was proposed in [24] to achieve proactive safe planning using CBF and optimal performance with an off-policy RL algorithm.

To the best of our knowledge, there is no work on regret guarantee reflecting sample efficiency of *safe learning* for nonlinear control with unknown model uncertainty. Note that safe learning for controlling unknown nonlinear dynamical systems is indeed a challenging problem, since safety and exploration are contrast to each other. In this paper, we propose a *provably correct* method that handles both sample efficient safe learning and online nonlinear control task in partially unknown system dynamics. In particular, we develop an Optimism-based Safe Learning for Control framework that integrates 1) stochastic discrete-time control barrier functions (CBF) to ensure safety with high probability under uncertainty, and 2) an optimism-based exploration strategy that enjoys a formally provable regret bound. To leverage between safe exploration and exploitation, the framework utilizes an optimism-based exploration strategy in face of uncertainty [18]. This encourages efficient dynamics exploration and simultaneously synthesizes with model-based nonlinear control algorithms to safely optimize the policy performance under the learned dynamics. Compared to existing works on safe learning or safe RL [6, 9, 10, 32, 37], our framework is able to simultaneously guarantee safety with high probability (remain in safe set during learning and

execution), performance (task rewards maximization), and sample efficiency (near-optimal regret bound). The effectiveness of the algorithm is validated through simulation results on a unicycle mobile robot and an inverted pendulum. Our **main contributions** are: 1) a provably sample efficient episodic online learning framework that integrates safe model-based nonlinear control approaches with optimism-based exploration strategy to simultaneously achieve safe learning and policy optimization for online nonlinear control tasks, and 2) rigorous theoretical analysis of guaranteed probabilistic safety under learned uncertainty and near-optimal online learning and policy performance with proved regret bound.

2 Preliminaries

2.1 Dynamical System and Stochastic Control

Consider the following partially *unknown* discrete-time control-affine system dynamics with state $x \in \mathcal{X} \subset \mathbb{R}^n$ and control input $u \in \mathcal{U} \subset \mathbb{R}^m$ for a discrete time index $h \in \mathbb{N}$

$$x_{h+1} = \hat{f}(x_h, u_h) + d(x_h, u_h) + \varepsilon_h, \quad \varepsilon_h \sim \mathcal{N}(0, \Sigma_\sigma) \quad (1)$$

where $\hat{f} : \mathcal{X} \times \mathcal{U} \mapsto \mathbb{R}^n$ is the known nominal discrete dynamics affine in the control input as $\hat{f}(x_h, u_h) = \hat{F}(x_h) + \hat{G}(x_h)u_h$. We assume $\hat{F} : \mathbb{R}^n \mapsto \mathbb{R}^n$, $\hat{G} : \mathbb{R}^n \mapsto \mathbb{R}^{n \times m}$ are locally Lipschitz continuous and the relative degrees of the nominal model and the actual system are the same, which are common assumptions as in [10, 32]. The system dynamics in this form is general and could describe a large family of nonlinear systems, e.g. 3-dof differential drive vehicles with unicycle dynamics [27, 36], 12-dof quadrotors with underactuated system [37], bipedal robots, automotive vehicle, and Segway robots [3, 32].

To describe a prior unknown model error representing uncertainty, $d : \mathcal{X} \times \mathcal{U} \mapsto \mathbb{R}^n$ denotes the unmodelled part of the system dynamics which is unknown. This unmodelled part $d(x, u)$ could represent state (action)-dependent external motion disturbances [37] or system model error due to parameter mismatch [9, 32]. ε_h is i.i.d noise sampled from a known Multivariate Gaussian distribution with the covariance matrix $\Sigma_\sigma = \text{diag}(\sigma_1^2, \dots, \sigma_n^2)$, i.e., $\sigma_1, \dots, \sigma_n$ are known to the learner. For notation simplicity, we denote the stochastic state transition as $P(\cdot|x, u)$. In particular, we assume that $d(x, u)$ is modelled by the nonlinear model $d(x, u) := W^* \phi(x, u)$ where $\phi : \mathcal{X} \times \mathcal{U} \mapsto \mathbb{R}^r$ is a known nonlinear feature mapping, e.g. Random Fourier Features (RFF) [28], and the linear mapping $W^* \in \mathbb{R}^{n \times r}$ is the unknown system parameters that need to be learned. Such model has been studied in [18] for unconstrained online control and in [23] for pure system identification. The model is flexible enough to describe the classic linear dynamical systems, nonlinear dynamical systems such as high order polynomials, and piece-wise linear systems [23]. The control task is described by a cost function. Given an immediate cost function $c : \mathcal{X} \times \mathcal{U} \mapsto \mathbb{R}^+$, the primary objective is

$$\min_{\pi \in \Pi} J^\pi(x_0; c, W^*) = \min_{\pi \in \Pi} \mathbb{E} \left[\sum_{h=0}^{H-1} c(x_h, u_h) | \pi, x_0, W^* \right] \quad (2)$$

where $\mathbf{x}_0 \in \mathcal{X}$ is a given starting state and Π is some set of pre-defined feasible controllers. Each controller (or a policy) is a mapping $\pi \in \Pi : \mathcal{X} \mapsto \mathcal{U}$. We denote $J^\pi(x; c, W)$ as the expected total cost of a policy π under cost function c , initial state x , and the dynamical system in Eq. 1 whose $d(x, u)$ is parameterized by W . To achieve optimal task performance, we need to learn the unmodelled $d(x, u)$ by taking samples to approximate to the true linear mapping W^* . On the other hand, simply optimizing the cost function is not enough for safety critical application. Below we consider a specific formulation that uses Control Barrier Functions (CBFs) [3] to enforce safety constraints.

2.2 Discrete-time Control Barrier Functions For Gaussian Dynamical Systems

We introduce CBFs in this section that are specialized to *discrete-time Gaussian stochastic dynamics* and provide safety analysis in a stochastic setting. Note that most existing continuous or discrete-time CBFs are defined only for deterministic system, e.g. [2, 3, 9, 10, 32, 37]. For stochastic system, we introduce a new CBF definition and prove the safety guarantee *with high probability* where the high probability is corresponding to the system Gaussian noise (Proposition 1). Consider a stochastic Gaussian discrete-time dynamics described in Eq. 1. A desired safety set $x \in \mathcal{S} \subset \mathcal{X}$ can be denoted by the following safety function $h^s : \mathbb{R}^n \mapsto \mathbb{R}$

$$\mathcal{S} = \{x \in \mathbb{R}^n : h^s(x) \geq 0\} \quad (3)$$

Formally, a safety condition is forward invariant if $x_{h=0} \in \mathcal{S}$ implies $x_h \in \mathcal{S}$ for all time step $h > 0$ with some designed controller $u \in \mathcal{U}$. Control barrier functions (CBF) [3] are often used to derive such designed controllers that enforce the forward invariance of a set of the system state space. Due to the stochasticity of the dynamics, we need to take the unbounded noise into consideration and thus we define a stochastic discrete-time CBF $h^s(\cdot)$ with the following condition.

Definition 1. [Discrete-time Control Barrier Function under Known Gaussian Dynamics] Assume $h^s(\cdot)$ is L -Lipschitz continuous when $x \in \mathcal{X}$ is bounded. Given $\delta_s \in (0, 1)$ and time horizon H , let \mathcal{S} be the 0-superlevel set of $h^s : \mathbb{R}^n \rightarrow \mathbb{R}$ which is a continuously differentiable function. We call $h^s(\cdot)$ a stochastic discrete-time control barrier function (CBF) for dynamical system Eq. 1 if there exists a $\eta \in (0, 1)$, such that for all time steps $h = 0, \dots, H - 1$, given any $x \in \mathcal{S}$:

$$\sup_{u \in \mathcal{U}} \left[h^s \left(\hat{f}(x, u) + d(x, u) \right) - L\bar{\sigma} \sqrt{2n \ln \left(\frac{Hn}{\delta_s} \right)} - h^s(x) \right] \geq -\eta h^s(x) \quad (4)$$

where $\bar{\sigma} = \max\{\sigma_1, \dots, \sigma_n\}$. Note that different from conventional CBF [3, 9] that is defined with respect to deterministic transition, the above definition takes the stochasticity into consideration. Also note that when $\bar{\sigma} \rightarrow 0$, i.e., the Gaussian dynamical system in Eq. 1 becomes a near deterministic system, then the above definition converges to the usual discrete-time CBF definition [9, 39]. We now show that under Definition 1, the resulting system state will stay in the safe set \mathcal{S} with probability at least $1 - \delta_s$.

Proposition 1 (Guaranteed Safety with High Probability). *Consider a CBF $h^s(\cdot)$ in Definition 1. Given $x_0 \in \mathcal{S}$, consider any policy $\pi : \mathcal{X} \rightarrow \mathcal{U}$ such that at any state x , this policy outputs an action $u = \pi(x)$ that satisfies the constraint Eq. 4. Then executing π to generate a trajectory starting at x_0 : $\tau = \{x_0, u_0, \dots, x_{H-1}, u_{H-1}\}$, with probability at least $1 - \delta_s$ we have $h^s(x_h) \geq 0$ for all $h \in [H]$, i.e., all states on the trajectory belong to the safe set \mathcal{S} .*

Proof. For notation simplicity, let us denote $f(x, u) := \hat{f}(x, u) + d(x, u)$. First consider as follows the probability of the event: $\exists h = 1, \dots, H, i = 1, \dots, n$ such that $\epsilon_h[i] \geq p$. Via union bound and the normal distribution's property, we have:

$$\mathbb{P}(\exists h \in [H], i \in [n], \text{ s.t., } \epsilon_h[i] \geq p) \leq Hn \exp(-p^2/(2\bar{\sigma}^2))$$

Let us set the failure probability to be δ_s , i.e., $Hn \exp(-p^2/(2\bar{\sigma}^2)) = \delta_s$. Solve for p and we indeed have $p := \bar{\sigma} \sqrt{2 \ln \left(\frac{Hn}{\delta_s} \right)}$. Below we conditioned the event that for all $h \in [H]$ and $i \in [n]$, we have $|\epsilon_h[i]| \leq p$. Note that this event happens with probability at least $1 - \delta_s$.

Due to the Lipschitz assumption with a Lipschitz constant L and the assumption that $h^s(\cdot)$ is differentiable, we have that for the system $x_{h+1} = f(x_h, u_h) + \epsilon_h$ at any time $h \in [H]$,

$$h^s(x_{h+1}) = h^s(f(x_h, u_h)) + \nabla h^s(\xi)^\top \epsilon_h \geq h^s(f(x_h, u_h)) - L\bar{\sigma} \sqrt{2n \ln \left(\frac{Hn}{\delta_s} \right)}$$

Since the control policy satisfies the control barrier function constraint in Eq. 4, we must have:

$$h^s(x_{h+1}) - h^s(x_h) \geq \left[h^s(f(x_h, u_h)) - L\bar{\sigma} \sqrt{2n \ln \left(\frac{Hn}{\delta_s} \right)} - h^s(x_h) \right] \geq -\eta h^s(x_h)$$

Namely, we have:

$$h^s(x_{h+1}) \geq (1 - \eta)h^s(x_h) \geq (1 - \eta)^{h+1}h^s(x_0) \geq 0,$$

under the condition that $h^s(x_0) \geq 0$. The above argument holds for all $h \in [H]$ which thus concludes the proof. \square

In presence of uncertainty on the considered system dynamics, it is noted that an initially chosen control barrier function $h^s(\cdot)$ in Eq. 4 describing the desired safety set \mathcal{S} should exist and remain valid during the safe learning task. Thus similar to [32], we make the following assumption without loss of generality.

Assumption 1. *If a function $h^s(\cdot)$ is a valid stochastic CBF as defined in Eq. 4 for the nominal stochastic system dynamics $\hat{f}(x_h, u_h) + \epsilon_h$, then it is a valid stochastic CBF for the uncertain stochastic dynamical system in Eq. 1.*

Mathematically this implies that

$$\begin{aligned} \sup_{u \in \mathcal{U}} \left[h^s \left(\hat{f}(x, u) \right) - L\bar{\sigma} \sqrt{2n \ln \left(\frac{Hn}{\delta_s} \right)} - h^s(x) \right] &\geq -\eta h^s(x) \\ \implies \sup_{u \in \mathcal{U}} \left[h^s \left(\hat{f}(x, u) + d(x, u) \right) - L\bar{\sigma} \sqrt{2n \ln \left(\frac{Hn}{\delta_s} \right)} - h^s(x) \right] &\geq -\eta h^s(x) \end{aligned}$$

Similar to the deterministic CBF [32], our assumption indicates that a set within the state space that can be kept safe for the nominal model of the stochastic system can also be kept safe for the actual uncertain system. This typically amounts to the assumption that the relative degree of $h^s(\cdot)$ for the nominal model is the same as the relative degree of $h^s(\cdot)$ for the actual system, so that the chosen CBF $h^s(\cdot)$ remains a valid CBF during learning [32]. Note that in order to reduce conservativeness during learning for improved sample efficiency, a strategic safe exploration process is still needed to efficiently learn the unmodelled $d(\cdot)$ while ensuring safety as the predicted model of $d(\cdot)$ being updated overtime.

2.3 Learning Objective

If assuming known dynamics $d(\cdot)$ in Eq. 1, then the safe nonlinear control problem can be modeled as follows:

$$\min_{\pi \in \Pi} J^\pi(x_0; c) \quad (5)$$

$$s.t. \quad h^s \left(\hat{f}(x, \pi(x)) + d(x, \pi(x)) \right) - L\bar{\sigma} \sqrt{2n \ln \left(\frac{Hn}{\delta_s} \right)} - h^s(x) \geq -\eta h^s(x) \quad (6)$$

Assume the above constrained optimization problem is feasible, and denote π^* as the optimal policy. Then, via Proposition 1, we know that with probability at least $1 - \delta_s$, π^* generates a trajectory whose states are all in the safe set \mathcal{S} . However, as the unmodelled part $d(x, u)$ is initially unknown, we cannot directly solve the above constrained optimization program using standard RL or MPC approaches. Instead we need to learn d (more specifically, the unknown linear mapping $W^* \in \mathbb{R}^{n \times r}$) online using data-driven approach for probabilistic safety guarantee as well as cost minimization.

Following the standard online learning setup [16, 35], in our episodic finite horizon learning framework we start with some initialization \bar{W}_0 which is used to parameterize $d_0(x, u) := \bar{W}_0 \phi(x, u)$. At every episode t , the learner will propose a policy $\pi_t \in \Pi$ (probably based on the current guess $d_t(x, u)$ with \bar{W}_t), execute π_t in the real system to generate one trajectory $\{x_0^t, u_0^t, \dots, x_{H-1}^t, u_{H-1}^t\}$ for H time steps; the learner then incrementally updates model parameter to \bar{W}_{t+1} using observations from all of the past trajectories, and move to the next episode $t+1$ starting from the same initial state $x \leftarrow x_0$. The goal is to ensure that π_t is safe (with high probability) in terms of satisfying CBF constraint Eq. 4, and also optimize the cost function over episodes:

$$\text{Regret}_T := \sum_{t=0}^{T-1} \sum_{h=0}^{H-1} c(x_h^t, u_h^t) - \sum_{t=0}^{T-1} J^{\pi^*}(x_0; c) = o(T) \quad (7)$$

Namely, comparing to the best policy π^* (i.e., the optimal solution of the constrained optimization program in Eq. 5 if assuming perfect model information), the cumulative regret grows sublinearly with respect to the number of episodes T . This goal implies that when $T \rightarrow \infty$ in a long run, the average episodic cost incurred by the learner is the same as the episodic cost incurred by the best policy π^* generated with the ground-truth dynamics Eq. 1.

3 Algorithm and Analysis

3.1 Calibrated Model and Approximate Safety Guarantee

Due to the unknown d , our initial guess d_0 using \bar{W}_0 estimated from pre-collected data could be arbitrarily bad, and hence it is impossible to achieve safe learning while simultaneously optimizing the regret without further assumptions. To ensure high-confidence safety (which could be overly conservative before exploration) starting with initially learned uncertainty from a pre-collected training data set, existing literature on safe learning often made the standard assumption that the prior calibrated model of d could yield initial safe policy to start the exploration process (e.g. [6, 9, 10, 19, 32, 37]), for example, by assuming the true d stays within the high confidence region estimated from Gaussian Processes using an initial data set. Hence without loss of generality, we make the typical assumption similar to existing literature of a calibrated initial model in the following.

Given N triples $(x_i, u_i, x'_i)_{i=1}^N$ with $x' \sim P(\cdot|x, u)$ as a set of pre-collected initial training data, we can compute \bar{W}_0 — the initialization parameters of $d(x, u)$ via ridge linear regression under known feature mapping $\phi : \mathcal{X} \times \mathcal{U} \mapsto \mathbb{R}^r$:

$$\bar{W}_0 = \arg \min_W \sum_{i=1}^N \left\| W\phi(x_i, u_i) - (x'_i - \hat{f}(x_i, u_i)) \right\|_2^2 + \lambda \|W\|_F^2 \quad (8)$$

where λ is a regularizer parameter and $\|W\|_F$ is the Frobenius norm of the model parameter matrix $W \in \mathbb{R}^{n \times r}$. Denote the initial empirical regularized covariance matrix as

$$V_0 = \sum_{i=1}^N \phi(x_i, u_i) \phi(x_i, u_i)^\top + \lambda I \quad (9)$$

Assumption 2. (*Calibrated model*) With \bar{W}_0, V_0 from the initial data $(x_i, u_i, x'_i)_{i=1}^N$ and $\epsilon, \delta \in (0, 1)$ where ϵ is a small error term, we can build the initial confidence ball describing the uncertain region of the linear mapping W^* with probability at least $1 - \delta$ as follows:

$$\text{Ball}_0 = \left\{ W : \left\| (W - \bar{W}_0) V_0^{-1/2} \right\|_2 \leq \beta, \quad \|W\|_2 \leq \|W^*\|_2 \right\} \quad (10)$$

where β is a hyper-parameter describing an appropriate confidence radius. Then for all $\tilde{W} \in \text{Ball}_0$:

$$\forall x, u \in \mathcal{X} \times \mathcal{U} : \left\| (\tilde{W} - W^*) \phi(x, u) \right\|_2 \leq \mathcal{O}(\epsilon).$$

See [Appendix Section B](#) for details. This entails for any $\widetilde{W} \in \mathbf{Ball}_0$, its prediction $\widetilde{d}(x, u) = \widetilde{W}\phi(x, u)$ for any x, u is reasonably good to the true prediction $d(x, u) = W^*\phi(x, u)$ from W^* (note that however we cannot guarantee \widetilde{W}_0 will be close to W^* in terms of ℓ_2 norm). As discussed in the following Theorem 1, this assumption ensures when we control the dynamical system using CBF with any model $\widetilde{W} \in \mathbf{Ball}_0$, we can ensure safety update to $\mathcal{O}(\epsilon)$ since $W^* \in \mathbf{Ball}_0$ with high probability. As more data are collected by each future training episode t , we show in Section 3.2 the decreased confidence region \mathbf{Ball}_t covers true W^* with high probability, yielding less conservative safe learning behavior over time and achieving provable sample efficiency during learning.

Theorem 1 (Policy for Approximate High-Probability Safety Guarantee with Learned Dynamics). *Under Assumption 2, consider any $\widetilde{W} \in \mathbf{Ball}_0$, and define any policy $\pi_s : \mathcal{X} \mapsto \mathcal{U}$ that satisfies the CBF constraint parameterized by \widetilde{W} , i.e.,*

$$\forall x \in \mathcal{X} : \pi_s(x) \in \mathcal{U}_x := \left\{ u : h^s \left(\hat{f}(x, u) + \widetilde{W}\phi(x, u) \right) - L\bar{\sigma} \sqrt{2n \ln \left(\frac{Hn}{\delta_s} \right)} \geq (1 - \eta)h^s(x) \right\} \quad (11)$$

Then with probability at least $1 - \delta_s$, starting at any safe initial state $h^s(x_0) \geq 0$, π_s generates a safe trajectory $\{x_0, u_0, \dots, x_{H-1}, u_{H-1}\}$, such that for all time steps $h \in [H]$, $h^s(x_h) \geq -\mathcal{O}(\frac{L\epsilon}{\eta})$, where L is the Lipschitz constant of $h^s(\cdot)$ under bounded $x \in \mathcal{X}$.

The detailed proof (See [Appendix Section C](#)) follows from Assumption 2 and Proposition 1. Next, we show how to search for a policy using an optimism-based algorithm that performs as good as the best benchmark π^* in the sense of minimizing regret defined in Eq. 7 and subject to Eq.11.

3.2 Optimism-based Safe Learning for Control

To achieve no-regret performance for efficient safe learning for control, we leverage the LC^3 algorithm [18] for strategic exploration. However, LC^3 is designed for unconstrained optimization and hence not suitable for safety-critical applications. Here we adopt our CBF constraint Eq. 11 to select safe exploration policy. Similar to LC^3 , we need to leverage the principle of optimism in the face of uncertainty to achieve small regret. We propose the following framework of Optimism-based Safe Learning for Control (Algorithm 1).

At the beginning of each episode t , given all previous trajectories, $\tau^i = \{x_0^i, u_0^i, \dots, x_{H-1}^i, u_{H-1}^i, x_H^i\}$ from episode $i = 0$ to $t - 1$, we perform ridge linear regression to find \widetilde{W}_t in Line 6, i.e.,

$$\begin{aligned} \widetilde{W}_t = \arg \min_W \sum_{i=1}^N & \left(W\phi(x_i, u_i) - (x'_i - \hat{f}(x_i, u_i)) \right)^2 \\ & + \sum_{i=0}^{t-1} \sum_{h=0}^{H-1} \left(W\phi(x_h^i, u_h^i) - (x_{h+1}^i - \hat{f}(x_h^i, u_h^i)) \right)^2 + \lambda \|W\|_F^2 \end{aligned} \quad (12)$$

Algorithm 1 Optimism-based Safe Learning for Control

Input: CBF h^s , cost function c , initial data $(x_i, u_i, x'_i)_{i=1}^N$, initial confidence region Ball_0 with \bar{W}_0, Σ_0 , number of training episodes T , horizon H , regularizer λ , initial state x_0

Output: a sequence of policies for $t = 0, \dots, T$

- 1: **for** $t = 0, \dots, T$ **do**
- 2: $x_0^t \leftarrow x_0$
- 3: Sample $\widetilde{W}_t \sim \mathcal{N}(\bar{W}_t, \Sigma_t^{-1})$ in Ball_t # Thompson Sampling for Exploration
- 4: $\pi_s^t \leftarrow \arg \min_{\pi \in \Pi_{\widetilde{W}}} J^\pi(x_0^t; c, \widetilde{W}_t)$ with $\Pi_{\widetilde{W}}$ specified in Eq. 17 # Safe MPC Planning
- 5: Execute π_s^t to sample a trajectory $\tau^t := \{x_h^t, u_h^t, c_h^t, x_{h+1}^t\}_{h=0}^{H-1}$ # Execution and Data Collection
- 6: $\bar{W}_{t+1}, \Sigma_{t+1} \leftarrow \text{Update Ball}_{t+1}$ # Model Update
- Return a sequence of policies for $t = 0, \dots, T$

and we have the shape of the estimate region as

$$\Sigma_t = V_0 + \sum_{i=0}^{t-1} \sum_{h=0}^{H-1} \phi(x_h^i, u_h^i) \phi(x_h^i, u_h^i)^\top \quad (13)$$

where we use the data from trajectories $\{\tau^i\}_{i=0}^{t-1}$ and the initial data $(x_i, u_i, x'_i)_{i=1}^N$ for computing Ball_0 . Then the confidence region of W is defined as:

$$\text{Ball}_t = \text{Ball}_0 \cap \left\{ W : \|(W - \bar{W}_t) \Sigma_t^{1/2}\|_2 \leq \beta_t \right\} \quad (14)$$

with the confidence radius parameter β_t defined as:

$$\beta_t := \sqrt{\lambda} C_1 + \bar{\sigma} \sqrt{8n \ln(5) + 8r \ln(1 + (TH + N)/\lambda) + 8 \ln(1/\delta)} \quad (15)$$

where C_1 is a design hyper-parameter (see [Appendix Section B](#)). The uncertainty region Ball_t decreases as more data are collected over episodes. LC³ [18] shows that with probability $1 - \delta$, for all t , $W^* \in \{W : \|(W - \bar{W}_t) \Sigma_t^{1/2}\|_2 \leq \beta_t\}$. Then under our Assumption 2, we have the following proposition proving the existence of such intersection with true W^* living in the confidence interval Ball_t (Eq. 14).

Proposition 2. *Given the uncertainty regions $W^* \in \{W : \|(W - \bar{W}_t) \Sigma_t^{1/2}\|_2 \leq \beta_t\}$ (Proof of Lemma B.5 in [18]) and Ball_0 (Eq. 10) with the probability of $\Pr(W^* \in \{W : \|(W - \bar{W}_t) \Sigma_t^{1/2}\|_2 \leq \beta_t\}) \geq 1 - \delta$ and $\Pr(W^* \in \text{Ball}_0) \geq 1 - \delta$, then for all episode t we have*

$$\Pr\left(W^* \in \text{Ball}_t := \text{Ball}_0 \cap \left\{ W : \|(W - \bar{W}_t) \Sigma_t^{1/2}\|_2 \leq \beta_t \right\}\right) \geq 1 - 2\delta \quad (16)$$

where $\Pr(\cdot)$ denotes the probability of an event. Detailed proof is provided in [Appendix Section D](#). This ensures as more data are collected, given any model

$\widetilde{W} \in \text{Ball}_t \Rightarrow \widetilde{W} \in \text{Ball}_0$, we can constrain our policy class Π based on CBF constraint under \widetilde{W} (Eq. 11) in Theorem 1, i.e., we denote $\Pi_{\widetilde{W}}$ as follows:

$$\Pi_{\widetilde{W}} = \left\{ \pi_s \in \Pi : \forall x \in \mathcal{X}, \pi_s(x) \in \left\{ u : h^s \left(\hat{f}(x, u) + \widetilde{W}\phi(x, u) \right) - L\bar{\sigma} \sqrt{2n \ln \left(\frac{Hn}{\delta_s} \right)} \geq (1 - \eta)h^s(x) \right\} \right\} \quad (17)$$

As discussed in [2, 9, 39], a nonlinear discrete-time barrier function $h^s(\cdot)$ generally makes the constrained optimization process (Line 4 in our case subject to Eq. 17) a nonlinear programming problem (NLP) that is possibly non-convex, while with an affine barrier function $h^s(\cdot)$ it could become tractable as a convex optimization problem [9]. Such an optimization problem could be solved by using standard MPC approach. To obtain an approximate solution for improved computation efficiency, with \hat{f} affine in control and assumed local Lipschitz continuity of the dynamics and feature mapping function ϕ , Eq. 17 could also be approximated by linear control constraint w.r.t. u as follows to define the admissible control space

$$\begin{aligned} L_{\hat{F}}^{\Delta} h^s(x) + L_{\hat{G}}^{\Delta} h^s(x)u - L\bar{\sigma} \sqrt{2n \ln \left(\frac{Hn}{\delta_s} \right)} \\ \geq -\eta h^s(x) + |\Delta h^s(x) \widetilde{W} \phi(x, u^*)| + |\Delta h^s(x) \widetilde{W} L_{x,\phi}(u^+ - u^-)| \end{aligned} \quad (18)$$

where $L_{\hat{F}}^{\Delta} h^s(x)$ and $L_{\hat{G}}^{\Delta} h^s(x)$ are discrete-time Lie-derivatives of $h^s(x)$ obtained through Taylor's theorem along $\hat{F}(x)$ and $\hat{G}(x)$ respectively. $L_{x,\phi}$ is the local Lipschitz constant vector for the known feature mapping ϕ w.r.t. u at x . $\Delta h^s(x)$ is the discrete derivative of h^s and u^*, u^+, u^- are the nominal, max and min value of u respectively. [Appendix Section E](#) proves that any u satisfying Eq. 18 ensures $u \in \pi_s(x)$ in Eq. 17, thus constructing the safe policy class $\Pi_{\widetilde{W}}$. Note that if the ground truth dynamics d is only state-dependent as assumed in [6, 9, 37], then Eq. 18 is also linear in control where $L_{x,\phi} = \mathbf{0}$ with feature mapping $\phi(x)$. Now we select model and safe policy optimistically at each episode t :

$$(W_t, \pi^t) := \arg \min_{\widetilde{W} \in \text{Ball}_t} \arg \min_{\pi \in \Pi_{\widetilde{W}}} J^{\pi}(x_0^t; c, \widetilde{W}). \quad (19)$$

In face of uncertainty, we solve this problem by first utilizing Thompson Sampling [30] for a \widetilde{W}_t in Ball_t (Eq. 14) (Line 3) and then computing the safe optimal policy using MPC planning oracle such as model predictive path integral control (MPPI) [38] under \widetilde{W}_t , subject to safety control constraint Eq. 18 (Line 4). For computation efficiency in practice, the safe planning problem in Line 4-5 at each episode t could be cast into the following step-wise optimization process: at each time step $h \in [H]$ in episode t , the optimal control sequence to Line 4 without safety consideration could be derived as $\bar{u}_{h:h+H-1|h}^t = [\bar{u}_{h|h}^t, \bar{u}_{h+1|h}^t, \dots, \bar{u}_{h+H-1|h}^t]$ starting at x_h^t by using MPPI planning algorithm [38] under dynamics with sampled mapping \widetilde{W}_t . Then we denote the next control input for execution as nominal control $u^* \leftarrow \bar{u}_{h|h}^t$ and map it to the safe policy class by solving the following quadratic programming (QP) with linear control constraint Eq. 18:

$$u_h^t = \arg \min_u \|u - u^*\|^2 \quad \text{s.t.} \quad \text{Eq. 18} \quad \text{and} \quad u \in [u^-, u^+]$$

Note that the QP may be infeasible when no control solution can be found with the chosen barrier function constraints where the violation of safety is inevitable, e.g. starting at the boundary of safe set and too late to enforce safety. Readers are referred to [31] for discussion on how to create a barrier function that guarantees feasibility of the QP through nominal evading maneuvers.

Then after executing $\pi_s^t(x_h^t) \rightarrow u_h^t$ at x_h^t , the process is repeated at $h + 1$ on state x_{h+1}^t , computes unconstrained nominal control sequence $\bar{u}_{h+1:h+H|h+1}^t$, obtains safe control $\pi_s^t(x_{h+1}^t) \rightarrow u_{h+1}^t$ mapped from $\bar{u}_{h+1|h+1}^t$ as above, and a trajectory $\tau^t := \{x_h^t, u_h^t, c_h^t, x_{h+1}^t\}_{h=0}^{H-1}$ is iteratively sampled.

Given Eq. 19 and conditioned on the high probability event that $W^* \in \text{Ball}_t$, and $\pi^* \in \Pi_{W^*}$ by definition of π^* , we can show optimism in the sense that:

$$J^{\pi^t}(x_0^t; c, W_t) \leq J^{\pi^*}(x_0; c, W^*).$$

Hence the optimism allows us to prove the following main statement regarding the bounded regret similar to [18].

Theorem 2. [Regret Bound] Set $\lambda = \bar{\sigma}^2/C_1^2$. Our algorithm learns a sequence of policies π^0, \dots, π^{T-1} in T episodes, such that in expectation, we have:

$$\mathbb{E}[\text{Regret}_T] \leq \tilde{\mathcal{O}}\left(H\sqrt{Hr(r+n+H)T}\right).$$

Also with probability at least $1 - O(\delta_s)$, we have that for all $t \in [T]$, $h^s \in [H]$, $h(x_h^t) \geq -\mathcal{O}(L\epsilon/\eta)$.

Detailed proof is provided in [Appendix Section F](#).

4 Results

We evaluate our Optimism-based Safe Learning framework on two simulation platforms: simulated unicycle mobile robot and inverted pendulum. The cumulative rewards (negative cost) are used for evaluations, i.e. the higher the better.

Mobile Robot Navigation

To compare our method using exploration behavior (Algorithm 1) against exploitation (Nom-MPPI-CBF-Exploitation) in terms of sample efficiency, consider a mobile robot navigation task simulated in Matlab (Figure 1) where the unicycle robot moves in a wind field that represents the unmodelled part of the robot dynamics. In particular, the algorithm Nom-MPPI-CBF-Exploitation in comparison denotes our method with exploitation behavior only, i.e. replace Line 3 in Algorithm 1 by $\tilde{W}_t \leftarrow \bar{W}_t$. Algorithms GT-MPPI and Nom-MPPI denote the MPPI method [38] with ground-truth dynamics model and nominal dynamics (i.e. $d(\cdot) = 0$) respectively. Here we assume the unicycle dynamics is directly available as the nominal model for the learner but suffers from unknown wind field defined by $d^*(x) := [\cos(x_1 - 4)(x_2 - 3), \sin(x_1 - 4)(x_2 - 3)]^\top \in \mathbb{R}^2$ with $x = [x_1, x_2]^\top$ as the position of the robot. In particular, there is a rectangle area $[-2, 3] \times [-2.6, -0.2]$ in the environment where the wind $d^*(x)$ has uniform directions (East pointing) with larger magnitude. We use standard quadratic

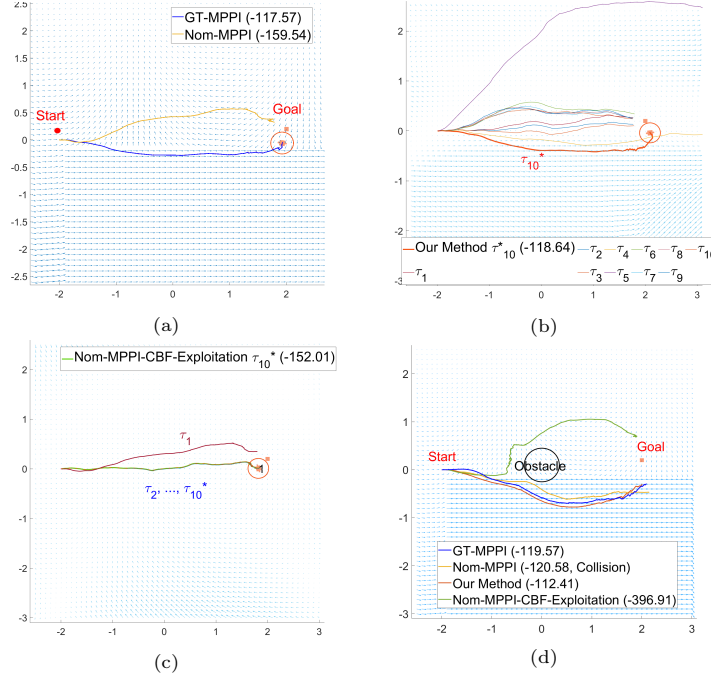


Fig. 1: Mobile robot navigation trajectories in an unknown wind field. (a) ground-truth wind field and trajectories from GT-MPPI (rewards: -117.57) and Nom-MPPI (rewards: -159.54) without obstacles, (b) predicted wind field and trajectories from our method (Algorithm 1) during training and testing after 10 episodes (rewards: -118.64) without obstacles, (c) predicted wind field and trajectories from our method using exploitation behavior during training and testing after 10 episodes (rewards: -152.01) without obstacles, (d) ground-truth wind field and trajectories from four different methods during testing with one static obstacle.

normalized cost $c = (x - x_{goal})^\top Q(x - x_{goal}) + u^\top Ru$ where Q, R are positive-definite to reflect the cost to go and to learn the optimal policy driving the robot towards the goal.

For obstacle-free scenario shown in Figure 1(a)-(c), GT-MPPI [38] with known ground-truth wind model is able to plan the trajectory that takes advantage of the wind field to enjoy the lowest cost (highest reward). It is observed that our method with optimism-based exploration behavior (Algorithm 1) in Figure 1(b) is able to quickly find a near-optimal trajectory after data collection during training in 10 episodes. The predicted wind field correctly reflects the significant different wind distribution in the rectangle area due to the exploration process. This outperforms the exploitation behavior from Nom-MPPI-CBF-Exploitation in Figure 1(c) that quickly converges to a local minima without much exploration in the unknown wind field, which thus fails to find the different wind field below in the rectangle area. Figure 1(d) shows a safe learning scenario with one static




Fig. 2: Performance curves of (a) cumulative rewards, (b) maximum theta angle, and (c) minimum theta angle in Inverted Pendulum environment testing under the same initial condition.

obstacle and the corresponding trajectories by the four different algorithms during testing. Our method is able to achieve similar performance compared to GT-MPPI and significantly outperforms Nom-MPPI-CBF-Exploitation whose trajectory (green) takes a much longer detour. This empirically validated the learning and control performance of our method.

Inverted Pendulum

To further compare the performance with quantitative results, we use the standard inverted pendulum platform modified from the OpenAI gym environment [7] with additive disturbance of $0.05 \cos(\theta_t - 3)$ on state update. The pendulum has ground truth mass $m = 1$ and length $l = 1$, and is controlled by the limited torque input $u \in [-15, 15]$. The standard cost function $c = \theta^2 + 0.1\dot{\theta}^2 + 0.001$ is used to learn the optimal policy keeping the pendulum upright (i.e. $\theta = 0$). Similar to [9], we randomly set the safe region to be $\theta \in [-1/8\pi, +5/4\pi]$ radians. We define the true system dynamics as $\theta_{t+1} = \theta_t + \dot{\theta}_{t+1}\Delta t + 0.05 \cos(\theta_t - 3)$ and $\dot{\theta}_{t+1} = \dot{\theta}_t + \frac{3g}{2l} \sin \theta_t \Delta t + \frac{3}{m'l^2} u \Delta t$. To describe the partially known system dynamics, we assume a nominal model as $\theta_{t+1} = \theta_t + \dot{\theta}_{t+1}\Delta t$ and $\dot{\theta}_{t+1} = \dot{\theta}_t + \frac{3g}{2l'} \sin \theta_t \Delta t + \frac{3}{m'l'^2} u \Delta t$ with incorrect model parameters $m' = 1.8, l' = 1.8$ available to the learner (hence 80% error in model parameters).

Using the same and different initial conditions respectively, Figure 2 and Figure 3 compare the cumulative rewards, maximum and minimum theta angle achieved during testing after each training episode by using (1) MPPI [38] with ground-truth dynamics model (GT-MPPI), (2) MPPI with nominal dynamics model and CBF (Nom-MPPI-CBF), (3) our method of optimism-based safe learning (Algorithm 1), (4) our method with exploitation only, i.e. replace Line 3 in Algorithm 1 by $\tilde{W}_t \leftarrow \bar{W}_t$ (Nom-MPPI-CBF-Exploitation), and (5) unconstrained Lower Confidence-based Continuous Control algorithm (LC3) [18]. The last three learning-based algorithms are trained for 50 episodes with 20 testing trials after each training episode averaged from four random seeds. It is observed that our method quickly increased cumulative reward in early stage while satisfying the safety constraints as learning process evolves, and our method using exploration behavior (our method) is able to increase reward even faster than our method




Fig. 3: Performance curves of (a) cumulative rewards, (b) maximum theta angle, and (c) minimum theta angle in Inverted Pendulum environment with four different initial conditions.

using exploitation behavior (Nom-MPPI-CBF-Exploitation), empirically implying sample efficiency. In contrast, GT-MPPI and LC3 severely violate angle limitation due to lack of safety consideration, and safe MPPI using CBF with nominal model (Nom-MPPI-CBF) still violates safety constraints with lower cumulative rewards due to the inaccurate nominal model with large error.

5 Conclusion

In this paper, we address the problem of episodic safe learning for online nonlinear control tasks. Compared to existing safe learning and control approaches that exhaustively expanding safety region or optimizing policy performance without efficiency guarantee, we propose an optimism-based online safe learning framework that simultaneously achieve sample efficient learning for safe behaviors and nonlinear control optimization with bounded regret guarantee. We believe our presented work is an important first step to bridge provably efficient learning based methods and model based safety-critical control with formal guarantees. Future work include extending our sample efficient learning on more complex high-dimensional dynamical systems.

References

1. Achiam, J., Held, D., Tamar, A., Abbeel, P.: Constrained policy optimization. In: Proceedings of the 34th International Conference on Machine Learning-Volume 70. pp. 22–31 (2017)
2. Agrawal, A., Sreenath, K.: Discrete control barrier functions for safety-critical control of discrete systems with application to bipedal robot navigation. In: Robotics: Science and Systems. vol. 13. Cambridge, MA, USA (2017)
3. Ames, A.D., Coogan, S., Egerstedt, M., Notomista, G., Sreenath, K., Tabuada, P.: Control barrier functions: Theory and applications. In: 18th European Control Conference (ECC). pp. 3420–3431. IEEE (2019)
4. Ames, A.D., Xu, X., Grizzle, J.W., Tabuada, P.: Control barrier function based quadratic programs for safety critical systems. IEEE Transactions on Automatic Control **62**(8), 3861–3876 (2017)

5. Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., Mané, D.: Concrete problems in ai safety. arXiv preprint arXiv:1606.06565 (2016)
6. Berkenkamp, F., Turchetta, M., Schoellig, A., Krause, A.: Safe model-based reinforcement learning with stability guarantees. In: Advances in neural information processing systems. pp. 908–918 (2017)
7. Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., Zaremba, W.: Openai gym. arXiv preprint arXiv:1606.01540 (2016)
8. Cheng, R., Khojasteh, M.J., Ames, A.D., Burdick, J.W.: Safe multi-agent interaction through robust control barrier functions with learned uncertainties. In: 59th IEEE Conference on Decision and Control (CDC). pp. 777–783. IEEE (2020)
9. Cheng, R., Orosz, G., Murray, R.M., Burdick, J.W.: 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, pp. 3387–3395 (2019)
10. Choi, J., Castañeda, F., Tomlin, C., Sreenath, K.: Reinforcement Learning for Safety-Critical Control under Model Uncertainty, using Control Lyapunov Functions and Control Barrier Functions. In: Proceedings of Robotics: Science and Systems. Corvallis, Oregon, USA (July 2020)
11. Clark, A.: Control barrier functions for complete and incomplete information stochastic systems. In: 2019 American Control Conference (ACC). pp. 2928–2935. IEEE (2019)
12. Duan, Y., Chen, X., Houthoofd, R., Schulman, J., Abbeel, P.: Benchmarking deep reinforcement learning for continuous control. In: International Conference on Machine Learning. pp. 1329–1338 (2016)
13. Fisac, J.F., Akametalu, A.K., Zeilinger, M.N., Kaynama, S., Gillula, J., Tomlin, C.J.: A general safety framework for learning-based control in uncertain robotic systems. IEEE Transactions on Automatic Control **64**(7), 2737–2752 (2018)
14. Garcia, J., Fernández, F.: A comprehensive survey on safe reinforcement learning. Journal of Machine Learning Research **16**(1), 1437–1480 (2015)
15. Gurriet, T., Singletary, A., Reher, J., Ciarletta, L., Feron, E., Ames, A.: Towards a framework for realizable safety critical control through active set invariance. In: 2018 ACM/IEEE 9th International Conference on Cyber-Physical Systems (ICCPs). pp. 98–106. IEEE (2018)
16. Hazan, E.: Introduction to online convex optimization. Foundations and Trends in Optimization **2**(3-4), 157–325 (2016)
17. Hogan, F.R., Rodriguez, A.: Reactive planar non-prehensile manipulation with hybrid model predictive control. The International Journal of Robotics Research **39**(7), 755–773 (2020)
18. Kakade, S., Krishnamurthy, A., Lowrey, K., Ohnishi, M., Sun, W.: Information theoretic regret bounds for online nonlinear control. Advances in Neural Information Processing Systems **33** (2020)
19. Khojasteh, M.J., Dhiman, V., Franceschetti, M., Atanasov, N.: Probabilistic safety constraints for learned high relative degree system dynamics. In: Learning for Dynamics and Control. pp. 781–792 (2020)
20. Koller, T., Berkenkamp, F., Turchetta, M., Krause, A.: Learning-based model predictive control for safe exploration. In: IEEE conference on decision and control (CDC). pp. 6059–6066. IEEE (2018)
21. Luo, W., Sun, W., Kapoor, A.: Multi-robot collision avoidance under uncertainty with probabilistic safety barrier certificates. Advances in Neural Information Processing Systems **33**, 372–383 (2020)

22. Lyu, Y., Luo, W., Dolan, J.M.: Probabilistic safety-assured adaptive merging control for autonomous vehicles. In: IEEE International Conference on Robotics and Automation (ICRA). pp. 10764–10770. IEEE (2021)
23. Mania, H., Jordan, M.I., Recht, B.: Active learning for nonlinear system identification with guarantees. arXiv preprint arXiv:2006.10277 (2020)
24. Marvi, Z., Kiumarsi, B.: Safe reinforcement learning: A control barrier function optimization approach. *International Journal of Robust and Nonlinear Control* **31**(6), 1923–1940 (2021)
25. Moldovan, T.M., Abbeel, P.: Safe exploration in markov decision processes. arXiv preprint arXiv:1205.4810 (2012)
26. Ohnishi, M., Wang, L., Notomista, G., Egerstedt, M.: Barrier-certified adaptive reinforcement learning with applications to brushbot navigation. *IEEE Transactions on robotics* **35**(5), 1186–1205 (2019)
27. Pickem, D., Glotfelter, P., Wang, L., Mote, M., Ames, A., Feron, E., Egerstedt, M.: The robotarium: A remotely accessible swarm robotics research testbed. In: IEEE International Conference on Robotics and Automation (ICRA). pp. 1699–1706. IEEE (2017)
28. Rahimi, A., Recht, B.: Random features for large-scale kernel machines. *Advances in neural information processing systems* **20**, 1177–1184 (2007)
29. Recht, B.: A tour of reinforcement learning: The view from continuous control. *Annual Review of Control, Robotics, and Autonomous Systems* **2**, 253–279 (2019)
30. Russo, D., Van Roy, B.: An information-theoretic analysis of thompson sampling. *The Journal of Machine Learning Research* **17**(1), 2442–2471 (2016)
31. Squires, E., Pierpaoli, P., Egerstedt, M.: Constructive barrier certificates with applications to fixed-wing aircraft collision avoidance. In: IEEE Conference on Control Technology and Applications (CCTA). pp. 1656–1661. IEEE (2018)
32. Taylor, A., Singletary, A., Yue, Y., Ames, A.: Learning for safety-critical control with control barrier functions. In: *Learning for Dynamics and Control*. pp. 708–717. PMLR (2020)
33. Taylor, A.J., Ames, A.D.: Adaptive safety with control barrier functions. In: 2020 American Control Conference (ACC). pp. 1399–1405. IEEE (2020)
34. Turchetta, M., Berkenkamp, F., Krause, A.: Safe exploration in finite markov decision processes with gaussian processes. *Advances in Neural Information Processing Systems* **29**, 4312–4320 (2016)
35. Wagener, N., Cheng, C., Sacks, J., Boots, B.: An online learning approach to model predictive control. In: *Proceedings of Robotics: Science and Systems (RSS)* (2019)
36. Wang, L., Ames, A.D., Egerstedt, M.: Safety barrier certificates for collisions-free multirobot systems. *IEEE Transactions on Robotics* **33**(3), 661–674 (2017)
37. Wang, L., Theodorou, E.A., Egerstedt, M.: Safe learning of quadrotor dynamics using barrier certificates. In: IEEE International Conference on Robotics and Automation (ICRA). pp. 2460–2465. IEEE (2018)
38. Williams, G., Wagener, N., Goldfain, B., Drews, P., Rehg, J.M., Boots, B., Theodorou, E.A.: Information theoretic mpc for model-based reinforcement learning. In: IEEE International Conference on Robotics and Automation (ICRA). pp. 1714–1721. IEEE (2017)
39. Zeng, J., Zhang, B., Sreenath, K.: Safety-critical model predictive control with discrete-time control barrier function pp. 3882–3889 (2021)
40. Zhu, H., Alonso-Mora, J.: Chance-constrained collision avoidance for mavs in dynamic environments. *IEEE Robotics and Automation Letters* **4**(2), 776–783 (2019)