

SABLAS: Learning Safe Control for Black-Box Dynamical Systems

Zengyi Qin , Dawei Sun , and Chuchu Fan 

Abstract—Control certificates based on barrier functions have been a powerful tool to generate probably safe control policies for dynamical systems. However, existing methods based on barrier certificates are normally for white-box systems with differentiable dynamics, which makes them inapplicable to many practical applications where the system is a black-box and cannot be accurately modeled. On the other side, model-free reinforcement learning (RL) methods for black-box systems suffer from lack of safety guarantees and low sampling efficiency. In this paper, we propose a novel method that can learn safe control policies and barrier certificates for black-box dynamical systems, without requiring for an accurate system model. Our method re-designs the loss function to back-propagate gradient to the control policy even when the black-box dynamical system is non-differentiable, and we show that the safety certificates hold on the black-box system. Empirical results in simulation show that our method can significantly improve the performance of the learned policies by achieving nearly 100% safety and goal reaching rates using much fewer training samples, compared to state-of-the-art black-box safe control methods. Our learned agents can also generalize to unseen scenarios while keeping the original performance. The source code can be found at <https://github.com/Zengyi-Qin/bcbf>.

Index Terms—Robust/adaptive control, robot safety.

I. INTRODUCTION

GUARANTEERING safety is an open challenge in designing control policies for many autonomous robotic systems, ranging from consumer electronics to self-driving cars and aircrafts. In recent years, the development of machine learning (ML) has created unprecedented opportunities to control modern autonomous systems with growing complexity. However, ML also poses great challenges for developing high-assurance autonomous systems that are provably dependable. While many learning-based approaches [1]–[4] have been proposed to train controllers to accomplish complex tasks with improved empirical performance, the lack of safety certificates for the learning-enabled components has been a fundamental hurdle that blocks the massive deployment of the learned solutions.

Manuscript received September 9, 2021; accepted January 5, 2022. Date of publication January 13, 2022; date of current version January 24, 2022. This letter was recommended for publication by Associate Editor Stephan Weiss and Editor Clement Gosselin upon evaluation of the reviewers' comments. This work was supported by the Defense Science and Technology Agency in Singapore. (Corresponding author: Chuchu Fan.)

Zengyi Qin and Chuchu Fan are with the Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: qinzy@mit.edu; chuchufan1990@gmail.com).

Dawei Sun is with the University of Illinois at Urbana-Champaign, Champaign, IL 61820 USA (e-mail: daweis2@illinois.edu).

This letter has supplementary downloadable material available at <https://doi.org/10.1109/LRA.2022.3142743>, provided by the authors.

Digital Object Identifier 10.1109/LRA.2022.3142743

For decades, mathematical control certificates has been used as proofs that the desired properties of the system are satisfied in closed-loop with certain control policies. For example, Control Lyapunov Function [5], [6] and Control Contraction Metrics [7]–[10] ensure the existence of a controller under which the system converges to an equilibrium or a desired trajectory. Control Barrier Function [11]–[17] ensures the existence of a controller that keeps the system inside a safe invariant set. Classical control synthesis methods based on Sum-of-Squares [12], [18]–[20] and linear programs [21], [22] usually only work for systems with low-dimensional state space and simple dynamics. This is because they choose polynomials as the candidate certificate functions, where low-order polynomials may not be a valid certificate for complex dynamical systems and high-order polynomials will increase the number of decision variables exponentially in those optimization problems. Recent data-driven methods [9], [17], [23], [24] have shown significant progress in overcoming the limitations of the traditional synthesis methods. They can jointly learn the control certificates and policies as neural networks (NN). However, data-driven methods that can generate control certificates still require an explicit (or white-box) differentiable model of the system dynamics, as the derivatives of the dynamics are required in the learning process. Finally, many works study the use of control certificates such as CBF to provide safety guarantees during and after training the policies [16], [25], but they more or less requires an accurate model of the system dynamics or a reasonable modeling of the system uncertainties to build the certificates.

Many dynamical systems in the real-world are black-box and lack accurate models. The most popular model-free approach to handle such black-box systems is safe reinforcement learning (RL). Safe RL methods enforce safety and performance by maximizing the expectation of the cumulative reward and constraining the expectation of the cost to be less or equal to a given threshold. The biggest disadvantage of safe RL methods is the lack of systematic or theoretically grounded way of designing cost function and reward functions, which heavily rely on empirical trials and errors. The lack of explainable safety guarantees and low sampling efficiency also make safe RL methods difficult to exhibit satisfactory performance.

Instead of stressing about the trade-off between the strong guarantees from control certificates and the practicability of model-free methods, in this work, we propose SABLAS to achieve both. SABLAS is a general-purpose approach to learning safe control for black-box dynamical systems. SABLAS enjoys the guarantees provided by the safety certificate from CBF

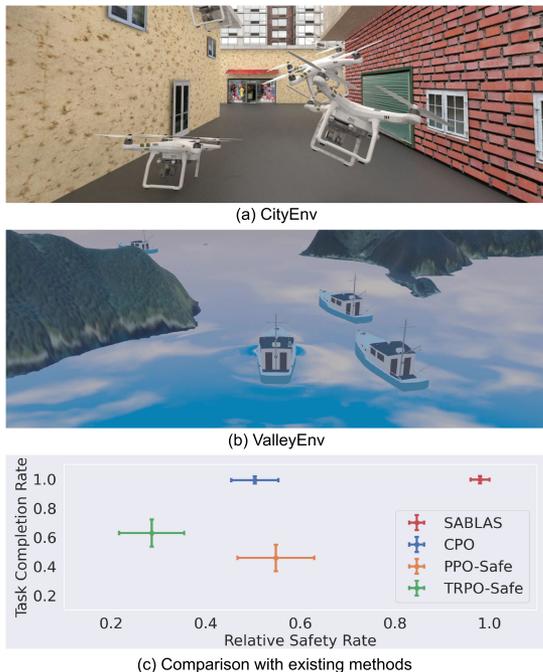


Fig. 1. (a) and (b) are two task environments we consider. In both CityEnv and ValleyEnv, the controller should control the ego vehicle to avoid collision with NPC vehicles. There are 1024 NPC drones in the city and 32 NPC ships in the valley. Each vehicle is assigned a random initial and goal location. (c) shows the average result for the two tasks, where the errorbars represent two times the standard deviation. SABLAS exhibits significant improvement in relative safety rate and task completion rate.

theory without requiring for an accurate model for the dynamics. Instead, SABLAS only needs a nominal dynamics function that can be obtained through regressions over simulation data. There is no need to model the error between the nominal dynamics and the real dynamics since SABLAS essentially re-design the loss function in a novel way to back-propagate gradient to the controller even when the black-box dynamical system is non-differentiable. The resulting CBF (and the corresponding safety certificate) holds directly on the original black-box system if the training process converges. The proposed algorithm is easy-to-implement and follows almost the same procedure of learning CBF for white-box systems with minimal modification. SABLAS fundamentally solves the problem that control certificates cannot be learned directly on black-box systems, and opens up the next chapter use of CBF theory on synthesizing safe controllers for black-box dynamical systems.

Experimental results demonstrates the superior advantages of SABLAS over leading learning-based safe control methods for black-box systems including CPO [26], PPO-Safe [3] and TRPO-Safe [27], [28]. We evaluate SABLAS on two challenging tasks in simulation: drone control in a city and ship control in a valley (as shown in Fig. 1). The dynamics of the drone and ship are assumed unknown. In both tasks, the controlled agent should avoid collision with uncontrolled agents and other obstacles, and reach its goal before the testing episode ends. We also examine the generalization capability of SABLAS on testing scenarios that are not seen in training. Fig. 1 shows that SABLAS can

reach a near 1.0 relative safety rate and task completion rate while using only 1/10 of the training data compared to existing safe RL methods, demonstrating a significant improvement. We also study the effect of model error (between the nominal model and actual dynamics) on the performance of the learned policy. It is shown that SABLAS is tolerant to large model errors while keeping a high safety rate. A detailed description of the results is presented in the experiment section. Video results can be found at supplementary materials.

To summarize the strength of SABLAS: 1) SABLAS can jointly find a safety certificate (i.e. CBF) and the corresponding control policy on black-box dynamics; 2) Unlike RL-based methods that need tedious trial-and-error on designing the rewards, SABLAS provides a systematic way of learning certified control policy, without parameters (other than the standard hyper-parameters in NN training) that need fine-tuning. 3) Empirical results that SABLAS can achieve a nearly perfect performance in terms of guaranteeing safety and goal-reaching, using much less samples than state-of-the-art safe RL methods.

II. RELATED WORK

There is a rich literature on controlling black-box systems and safe RL. Due to the space limit, we only discuss a few directly related and commonly used techniques on black-box system control. We will also skip the literature review for the large body of works on model-based safe control and trajectory planning as the research problems we are solving are very different.

A. Controller Synthesis for Black-Box Systems

Proportional–integral–derivative (PID) controller is widely used in controlling black-box systems. The advantage of PID controller is that it does not rely on a model of the system and only requires the measurement of the state. A drawback of PID controller is that it does not guarantee safety or stability, and the system may overshoot or oscillate about the control setpoint. If the underlying black-box system is linear time-invariant, existing work [29] has presented a polynomial-time control algorithm without relying on any knowledge of the environment. For non-linear black-box systems, the dynamics model can be approximated using system identification and controlled using model-based approaches [30], and PID can be used to handle the error part. The concept of practical relative degree [31] is also proposed to enhance the control performance on systems with heavy uncertainties. Recent advance in reinforcement learning [3], [4], [27], [32] also gives us insight into treating the system as a pure black-box and estimating the gradient for the black-box functions in order to optimize the control variables. However, in safety-critical systems, these black-box control methods still lack formal safety guarantee or certificate.

Simulation-guided controller synthesis methods can also generate control certificates for black-box systems, and sometimes those certificates can indicate how policies should be constructed [20]–[22]. However, most of these techniques use polynomial templates for the certificates, which limits their use on high-dimensional and complex systems. Another line of work studies the use of [33], [34] data-driven reachability analysis,

jointly with receding-horizon control to constructed optimal control policies. These methods rely on side information about the black-box systems (e.g. lipschitz constants of the dynamics, monotonicity of the states, decoupling in the states' dynamics) to do the reachability, which is not needed in our method.

B. Safe Reinforcement Learning

Safe RL [2], [26], [28], [35], [36] extends RL by adding constraints on the expectation of certain cost functions, which encode safety requirements or resource limits. CPO [26] derived a policy improvement step that increases the reward while satisfying the safety constraint. DCRL [2] imposes constraint on state density functions rather than cost value functions, and shows that density constraint has better expressiveness over cost value function-based constraints. RCPO [28] weights the cost using Lagrangian multipliers and add it to the reward. FISAR [36] uses a meta-optimizer to achieve forward invariance in the safe set. A disadvantage of safe RL is that it does not provide safety guarantee, or their safety guarantee cannot be realized in practice. The problem of sampling efficiency and sparsity of cost also increase the difficulty to synthesize safe controller through RL.

C. Safety Certificate and Control Barrier Functions

Mathematical certificates can serve as proofs that the desired property of the system is satisfied under the corresponding planning [37], [38] and control components. Such certificate is able to guide the controller synthesis for dynamical systems in order to ensure safety. For example, Control Lyapunov Function [5], [6], [24] ensures the existence of a controller so that the system converges to desired behavior. Control Barrier Function [11]–[17], [39] ensures the existence of a controller that keep the system inside a safe invariant set. However, the existing controller synthesis with safety certificate heavily relies on a white-box model of the system dynamics. For black-box systems or system with large model uncertainty, these methods are not applicable. While recent work [40] proposes to learn the model uncertainty effect on the CBF conditions, it still assumes that a handcrafted CBF is given, which is not always available for complex non-linear dynamical systems. Our approach represents a substantial improvement over the existing CBF-based safe controller synthesis strategy. The proposed SABLAS framework simultaneously enjoys the safety certificate from the CBF theory and the effectiveness on black-box dynamical systems.

III. PRELIMINARIES

A. Safety of Black-Box Dynamical Systems

Definition 1 (Black-box Dynamical System): A black-box dynamical system is represented by tuple $\langle \mathcal{S}, \mathcal{U}, f \rangle$, where $\mathcal{S} \in \mathbb{R}^n$ is the state space and $\mathcal{U} \in \mathbb{R}^m$ is the control input space. $f : \mathcal{S} \times \mathcal{U} \mapsto \mathcal{S}$ is the system dynamics $\dot{s} = f(s, u)$, which is unknown due to the black-box assumption.

Let $\mathcal{S}_0, \mathcal{S}_g, \mathcal{S}_s$ and \mathcal{S}_d denote the set of initial states, goal states, safe states and dangerous states respectively. The problem we aim to solve is formalized in Definition 2.

Definition 2 (Safe Control of Black-box Systems): Given a black-box dynamical system modeled as in Definition 1, the safe control problem aims to find a controller $\pi : \mathcal{S} \mapsto \mathcal{U}$ such that under control input $u = \pi(s)$ and the unknown dynamics $\dot{s} = f(s, u)$, the following is satisfied:

$$\begin{aligned} &\exists t > 0, s(t) \in \mathcal{S}_g \\ &\text{s.t. } \forall t > 0, s(t) \notin \mathcal{S}_d \text{ and } s(0) \in \mathcal{S}_0 \end{aligned} \quad (1)$$

The above definition requires that starting from the initial set \mathcal{S}_0 , the system should never enter the dangerous set \mathcal{S}_d under controller π .

B. Control Barrier Function as Safety Certificate

A common approach for guaranteeing safety of dynamical systems is via control barrier functions (CBF) [11], which ensures that the state always stay in the safe invariant set. A control barrier function $h : \mathcal{S} \mapsto \mathbb{R}$ satisfies:

$$\begin{aligned} h(s) &\geq 0, \quad \forall s \in \mathcal{S}_0 \\ h(s) &< 0, \quad \forall s \in \mathcal{S}_d \\ \dot{h} + \alpha(h) &= \frac{\partial h}{\partial s} f(s, u) + \alpha(h) \geq 0, \quad \forall s \in \mathcal{S}_p \end{aligned} \quad (2)$$

where $\mathcal{S}_p = \{s \mid h(s) \geq 0\}$, $u = \pi(s)$ and $\alpha(\cdot)$ is a class- \mathcal{K} function that is strictly increasing and $\alpha(0) = 0$. It is proven [11] that if there exists a CBF h for a given controller π , the system controlled by π starting from $s(0) \in \mathcal{S}_0$ will never enter the dangerous set \mathcal{S}_d . Besides the formal safety proof in [11], there is an informal but straightforward way to understand the safety guarantee. Whenever $h(s)$ decreases to 0, we have $\dot{h} + \alpha(h) = \dot{h} \geq 0$, which means $h(s)$ no longer decreases and $h(s) < 0$ will not occur. Thus s will not enter the dangerous set \mathcal{S}_d where $h(s) < 0$.

C. Co-Learning Controller and CBF for White-Box Systems

For white-box dynamical systems where $f(s, u)$ is known, we can jointly synthesize the controller π and its safety certificate h that satisfy CBF conditions (2) via a learning-based approach. We model π and h as neural networks with parameters Θ and Ω . Given a dataset \mathcal{D} of the state samples in \mathcal{S} , the CBF conditions (2) can be translated into empirical loss functions:

$$\begin{aligned} \mathcal{L}_0(\Omega) &= \frac{1}{|\mathcal{S}_0 \cap \mathcal{D}|} \sum_{s \in \mathcal{S}_0 \cap \mathcal{D}} \max(0, -h(s; \Omega)) \\ \mathcal{L}_d(\Omega) &= \frac{1}{|\mathcal{S}_d \cap \mathcal{D}|} \sum_{s \in \mathcal{S}_d \cap \mathcal{D}} \max(0, h(s; \Omega)) \\ \mathcal{L}_p(\Omega, \Theta) &= \frac{1}{|\mathcal{S}_p \cap \mathcal{D}|} \sum_{s \in \mathcal{S}_p \cap \mathcal{D}} \\ &\max \left(0, -\frac{\partial h}{\partial s} f(s, \pi(s; \Theta)) - \alpha(h(s, \Omega)) \right). \end{aligned} \quad (3)$$

Each of the loss functions in (3) corresponds to a CBF condition in (2). In addition to safety, we also consider goal-reaching by penalize the difference between the safe controller π and the

goal-reaching nominal controller π_{nom} in (4). The synthesis of π_{nom} is well-studied and is not the contribution of our work.

$$\mathcal{L}_g(\Theta) = \frac{1}{|\mathcal{S}_p \cap \mathcal{D}|} \sum_{s \in \mathcal{S}_p \cap \mathcal{D}} \|\pi(s; \Theta) - \pi_{nom}(s)\|_2^2. \quad (4)$$

The total loss function is $\mathcal{L}_0(\Omega) + \mathcal{L}_d(\Omega) + \mathcal{L}_p(\Omega, \Theta) + \lambda \mathcal{L}_g(\Theta)$, where λ is a constant that balances the goal-reaching and safety objective. The total loss function is minimized via stochastic gradient descent to find the parameters Ω and Θ . The dataset \mathcal{D} is not fixed during training and will be periodically updated with new samples by running the current controller. When the loss $\mathcal{L}_0 + \mathcal{L}_d + \mathcal{L}_p$ converges to 0, the resulting controller $\pi(s; \Theta)$ and CBF $h(s; \Omega)$ will satisfy (2) on unseen testing samples with a generalization bound as proven in [17]. Therefore, a safe controller and the corresponding CBF is found for the white-box system.

However, for black-box dynamical systems where $f(s, u)$ is unknown, the co-learning method for safe controller and CBF described above is no longer applicable. In Section IV, we propose an important and easy-to-implement modification to (3) such that we can leverage similar co-learning framework to jointly synthesize the safe controller and its corresponding CBF as safety certificate.

IV. LEARNING CBF ON BLACK-BOX SYSTEMS

In this section, we first elaborate on why it is difficult to learn the safe controller and its CBF for black-box dynamical systems. Then we will propose an important and easy-to-implement re-formulation of the optimization objective, which makes learning safe controller in black-box systems as easy as in white-box systems.

A. Challenges in Black-Box Dynamical Systems

Among the three loss functions \mathcal{L}_0 , \mathcal{L}_d and \mathcal{L}_p in (3), \mathcal{L}_p is the only one that can propagate gradient to the controller parameter Θ . The main challenge of training safe controller with its CBF for black-box dynamical systems is that the gradient can no longer be back-propagated to Θ when f is unknown. Therefore, a safe controller cannot be trained by minimizing the loss functions in (3).

Given state samples $s(t)$ and $s(t + \Delta t)$ from the black-box system where Δt is a sufficiently small time interval, we can approximate \dot{h} and compute \mathcal{L}_p as:

$$\begin{aligned} \dot{h}_1(s) &= \frac{h(s(t + \Delta t)) - h(s(t))}{\Delta t} \\ \mathcal{L}_{p1} &= \frac{1}{|\mathcal{S}_p \cap \mathcal{D}|} \sum_{s \in \mathcal{S}_p \cap \mathcal{D}} \max\left(0, -\dot{h}_1(s) - \alpha(h(s))\right). \end{aligned} \quad (5)$$

\mathcal{L}_{p1} does give the value of \mathcal{L}_p , but its backward gradient flow to the controller is cut off by the black-box system that is non-differentiable. (5) can only be used to train CBF h but not the safe controller π . Even worse, the h obtained by minimizing (5) does not guarantee that a corresponding safe controller π exists. If we have an differential expression of dynamics f and replace $h(s(t + \Delta t))$ with $h(s(t) + f(s(t), \pi(s(t)))\Delta t)$,

the gradient flow can successfully reach $\pi(s)$ and update the controller parameter. However, this is not immediately possible because f is unknown by the black-box assumption.

A possible way to back-propagate gradient to π is to use a differentiable nominal model f_{nom} . There are many methods to obtain f_{nom} , such as fitting a neural network using sampled data from the real black-box system. We do not require f_{nom} to perfectly match the real dynamics f , because there will always exist an error between them. With f_{nom} , we can approximate \mathcal{L}_p as:

$$\begin{aligned} \dot{h}_2(s) &= \frac{h(s + f_{nom}(s, \pi(s))\Delta t) - h(s)}{\Delta t} \\ \mathcal{L}_{p2} &= \frac{1}{|\mathcal{S}_p \cap \mathcal{D}|} \sum_{s \in \mathcal{S}_p \cap \mathcal{D}} \max\left(0, -\dot{h}_2(s) - \alpha(h(s))\right), \end{aligned} \quad (6)$$

which is differentiable w.r.t. $\pi(s)$ because both h and f_{nom} are differentiable. The gradient of \mathcal{L}_{p2} can be back-propagated to the controller to update its parameters. However, whatever way we get f_{nom} , there still exists an error between the real dynamics f and f_{nom} , which means \dot{h}_2 is not a good approximation of \dot{h} and \mathcal{L}_{p2} is not the true value of \mathcal{L}_p . Using \mathcal{L}_{p2} , it is not guaranteed that the third CBF condition in (2) will be satisfied.

B. Learning Safe Control for Black-Box Dynamics

We present a novel re-formulation of \mathcal{L}_p that makes learning safe controller with CBF for black-box dynamical systems as easy as in white-box systems. The proposed formulation possesses two features: it enables the gradient to back-propagate to the controller in training, and offers an error-free approximation of \dot{h} .

Given state samples $s(t)$ and $s(t + \Delta t)$ from the trajectories of the real black-box dynamical system, where Δt is a sufficiently small time interval, we define $s_{nom}(t + \Delta t)$ as:

$$s_{nom}(t + \Delta t) = s(t) + f_{nom}(s(t), \pi(s(t)))\Delta t,$$

then construct $\bar{s}(t + \Delta t)$ as:

$$\bar{s}(t + \Delta t) = s_{nom}(t + \Delta t) + g(s(t + \Delta t) - s_{nom}(t + \Delta t)),$$

where $g(s) = s$ is an identity function but *without gradient*. We need to pretend that $g(s)$ is a constant and in back-propagation, the gradient on $g(s)$ cannot propagate to its argument s . In PyTorch [41], there is an off-the-shelf implementation of $g(s)$ as $g(s) = s.detach()$, which cuts off the gradient from g to s in back-propagation. Then we approximate \mathcal{L}_p using:

$$\begin{aligned} \dot{h}_3(s) &= \frac{h(\bar{s}(t + \Delta t)) - h(s(t))}{\Delta t} \\ \mathcal{L}_{p3} &= \frac{1}{|\mathcal{S}_p \cap \mathcal{D}|} \sum_{s \in \mathcal{S}_p \cap \mathcal{D}} \max\left(0, -\dot{h}_3(s) - \alpha(h(s))\right). \end{aligned} \quad (7)$$

Theorem 1: $\nabla_{\Theta} \mathcal{L}_{p3}$ exists and $\lim_{\Delta t \rightarrow 0} \mathcal{L}_{p3} = \mathcal{L}_p$. Namely, \mathcal{L}_{p3} is differentiable w.r.t. the controller parameter Θ , and \mathcal{L}_{p3} is an error-free approximation of \mathcal{L}_p as $\Delta t \rightarrow 0$.

Proof: Since f_{nom} is differentiable, s_{nom} and $\bar{s}(t + \Delta t)$ are differentiable w.r.t. π , \dot{h}_3 and \mathcal{L}_{p3} are also differentiable w.r.t. π and its parameter Θ . Thus, $\nabla_{\Theta} \mathcal{L}_{p3}$ exists. Furthermore, since

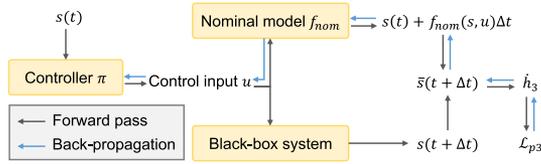


Fig. 2. Computational graph of \mathcal{L}_{p3} . The blue lines show how the gradient from \mathcal{L}_{p3} is back-propagated to the controller π despite that the black-box system itself is non-differentiable.

$\bar{s}(t + \Delta t) = s(t + \Delta t)$, \dot{h}_3 is an error-free approximation of the real \dot{h} when $\Delta t \rightarrow 0$. Thus \mathcal{L}_{p3} is also an error-free approximation of \mathcal{L}_p when $\Delta t \rightarrow 0$.

Note that Theorem 1 reveals the reason why the proposed SABLAS method can jointly learn the CBF and the safe controller for black-box dynamical system. **First**, since $\nabla_{\Theta} \mathcal{L}_{p3}$ exists, the gradient from \mathcal{L}_{p3} can be back-propagated to the controller parameters to learn a safe controller. On the contrary, \dot{h}_1 is not differentiable w.r.t. π so \mathcal{L}_{p1} cannot be used to train the controller. **Second**, since \mathcal{L}_{p3} is a good approximation of \mathcal{L}_p , minimizing \mathcal{L}_{p3} contributes to the minimization of \mathcal{L}_p and the satisfaction of the third CBF condition in (2). On the contrary, \dot{h}_2 is an inaccurate approximation of \dot{h} as we elaborated in Section IV-A. The construction of \mathcal{L}_{p3} incorporates the advantages of \mathcal{L}_{p1} and \mathcal{L}_{p2} , and avoids the disadvantages of \mathcal{L}_{p1} and \mathcal{L}_{p2} at the same time. The computational graph of \mathcal{L}_{p3} is illustrated in Fig. 2, which shows the forward pass and the backward gradient propagation from \mathcal{L}_{p3} to controller π .

Remark 1: One may argue that the gradient received by π via minimizing \mathcal{L}_{p3} is not exactly the gradient it should receive if we had a perfect differentiable model of the black-box system. Despite this, minimizing \mathcal{L}_{p3} directly contributes to the satisfaction of the third CBF condition in (2). A safe controller and its CBF can be found to keep the system within safe set.

Remark 2: Although the current formulation of \mathcal{L}_{p3} leads to promising performance in simulation as we will show in experiments, \mathcal{L}_{p3} requires further consideration in hardware experiments. Directly using the future state $s(t + \Delta t)$ to calculate the time derivative of h or s is not always desirable because noise will possibly dominate the numerical differentiation. When the noise dominates the time derivative of s or h , the training will have convergence issues. But a moderate noise is actually beneficial to training, because our optimization objective makes the CBF conditions hold even under noise disturbance, which increases the robustness of the trained CBF and controller. On physical robots where noise dominates the numerical differentiation, one can incorporate filtering techniques to mitigate the noise.

Combining with the loss functions \mathcal{L}_0 , \mathcal{L}_d and \mathcal{L}_g in (3) and (4), the total loss function can be formulated as:

$$\mathcal{L}(\Omega, \Theta) = \mathcal{L}_0(\Omega) + \mathcal{L}_d(\Omega) + \mathcal{L}_{p3}(\Omega, \Theta) + \lambda \mathcal{L}_g(\Theta), \quad (8)$$

where λ is a constant balancing the safety and goal-reaching objective. \mathcal{L} is minimized via stochastic gradient descent. Algorithm 1 summarizes the learning process of the safe controller and the corresponding safety certificate (CBF) for black-box

Algorithm 1: Learning Safe Controller With CBF for Black-Box Dynamical Systems.

- 1: **Input:** The set of initial condition \mathcal{S}_0 , the set of dangerous state \mathcal{S}_d , controller π and CBF h with randomly initialized parameters Θ and Ω , nominal dynamics f_{nom} , loss function $\mathcal{L}(\Omega, \Theta)$, the number of episodes K in data collection, the number of training iterations N .
 - 2: **for** $i = 1, 2, \dots, N$ **do**
 - 3: Initialize state samples dataset $\mathcal{D} \leftarrow \emptyset$
 - 4: **for** $j = 1, 2, \dots, K$ **do**
 - 5: $s(0) \leftarrow \text{SAMPLE}(\mathcal{S}_0)$
 - 6: $\mathcal{D}_j \leftarrow \text{RUN}(\pi(\cdot; \Theta), s(0))$, $\mathcal{D} \leftarrow \mathcal{D} \cup \mathcal{D}_j$
 - 7: **end for**
 - 8: $\Omega, \Theta \leftarrow \text{UPDATE}(\Omega, \Theta, \mathcal{D}, \mathcal{L})$
 - 9: **end for**
 - 10: **return** CBF parameter Ω and controller parameter Θ
-

dynamical systems. The RUN function runs the black-box system using the controller π under initial condition $s(0)$, and returns the trajectory data. The UPDATE function updates parameters Ω, Θ by minimizing $\mathcal{L}(\Omega, \Theta)$ using the state samples in \mathcal{D} via gradient descent.

V. EXPERIMENT

The primary objective of our experiment is to examine the effectiveness of the proposed method in terms of safety and goal-reaching when controlling black-box dynamical systems. We will conduct comprehensive experiments on two simulation environments illustrated in Fig. 1(a) and (b), and compare with state-of-the-art learning-based control methods for black-box systems.

A. Task Environment Description

a) *Drone control in a city (CityEnv):* In our first case study, we consider the package delivery task in a city using drones, as is illustrated in Fig. 1(a). There is one controlled drone and 1024 non-player character (NPC) drones that are not controlled by our controller. In each simulation episode, each drone is assigned a sequence of randomly selected goals to visit. The aim of our controller is to make sure the controlled drone reach its goals while avoiding collision with NPC drones at any time. A reference trajectory $s_{ref}(t)$, $t \in [0, T]$ will be given, which sequentially connects the goals and avoid collision with buildings. The reference trajectory can be generated by any off-the-shelf single-agent path planning algorithm. We use FACTEST [42] in our implementation, and other options such as RRT [43] are also suitable. The reference path planner does not need to consider the dynamic obstacles, such as the moving NPCs in our experiment. A nominal controller π_{norm} will also be given, which outputs control commands that drive the drone to follow the reference trajectory. However, π_{norm} is purely for goal-reaching and does not consider safety. The CityEnv has two modes: with **static** NPCs and **moving** NPCs. If the NPCs

are static, they will constantly stay at their initial locations. If the NPCs are moving, they will follow pre-planned trajectories to sequentially visit their goals. The drone model is with state space $[x, y, z, v_x, v_y, v_z, \theta_x, \theta_y]$, where θ_x and θ_y are roll and pitch angles. The control inputs are the angular acceleration of θ_x, θ_y and the vertical thrust. The underlying model dynamics is from [17] and assumed unknown to the controller and CBF in our experiment.

b) Ship control in a valley (ValleyEnv): In our second case study, we consider task of controlling a ship in valley illustrated in Fig. 1(b). There are one controlled ship and 32 NPC ships. The number of NPCs in ValleyEnv is less than CityEnv because ships are large in size and inertia, and hard to maneuver in dense traffic. Also, different from the 3D CityEnv, ValleyEnv is in 2D, which means the agents have fewer degrees of freedom to avoid collision. The initial location and goal location of each ship is randomly initialized at the beginning of each episode. The aim of our controller is to ensure the controlled ship reach its goal and avoid collision with NPC ships. Similar to CityEnv, a reference trajectory and nominal controller will be provided. There are also two modes in ValleyEnv, including the static and moving NPC mode, as is in CityEnv. The ship model is with state space $[x, y, \theta, u, v, \omega]$, where θ is the heading angle, u, v are speed in the ship body coordinates, and ω is the angular velocity of the heading angle. The ship model is from Section IV-B of [44] and is unknown to the controller and CBF.

B. Evaluation Criteria

Three evaluation criteria are considered. **Relative safety rate** measures the improvement of safety comparing to a nominal controller that only targets at goal-reaching but not safety. To formally define the relative safety rate, we first consider the absolute safety rate α as $\alpha = \frac{1}{T} \int_0^T \mathbb{I}(s(t) \notin \mathcal{S}_d) dt$, which measures the proportion of time that the system stays outside the dangerous set. Given two control policies π_1 and π_2 with absolute safety rate α_1 and α_2 , the relative safety rate of π_1 w.r.t. π_2 is defined as $\beta_{12} = \frac{\alpha_1 - \alpha_2}{1 - \alpha_2} \in (-\infty, 1]$. If $\beta_{12} = 0$, then control policy π_1 does not have any improvement over π_2 in terms of safety. If $\beta_{12} = 1$, then π_1 completely guarantees safety of the system in $t \in [0, T]$. In our experiment, π_1 is the controller to be evaluated, and π_2 is the nominal controller π_{norm} that only accounts for goal-reaching without considering safety. **Task completion rate** is defined as the success rate of reaching the goal state before timeout. **Tracking error** is the average deviation of the system's state trajectory comparing to a pre-planned reference trajectory $s_{ref}(t), t \in [0, T]$ as $\gamma = \frac{1}{T} \int_0^T \|s(t) - s_{ref}(t)\|_2^2 dt$. Note that we do not assume s_{ref} always stay outside the dangerous set.

C. Baseline Approaches

In terms of safe control for black-box systems, the most recent state-of-the-art approaches are safe reinforcement learning (safe RL) algorithms. We choose three safe RL algorithms for comparison: **CPO** [26] is a general-purpose policy optimization algorithm for black-box systems that maximizes the expected reward while satisfying the safety constraints. **PPO-Safe** is

a combination of PPO [3] and RCPO [28]. It uses PPO to maximize the expected cumulative reward while leveraging the Lagrangian multiplier update rule in RCPO to enforce the safety constraint. **TRPO-Safe** is a combination of TRPO [27] and RCPO [28]. The expected reward is maximized via TRPO and the safety constraints are imposed using the Lagrangian multiplier in RCPO.

D. Implementation and Training

Both the controller π and CBF h are multi-layer perceptrons (MLP) with architecture adopted from Section IV-B of [17]. π and h not only take the state of the controlled agent as input, but also the states of 8 nearest NPCs that the controlled agent can observe. In Algorithm 1, we choose $K = 1, N = 2000$. The total number of state samples collected during training is 10^6 . In UPDATE of the algorithm, we use the Adam [45] optimizer with learning rate 10^{-4} and batch size 1024. The gradient descent runs for 100 iterations in UPDATE. The nominal model dynamics are fitted from trajectory data in simulation. We used 10^4 state samples to fit a linear approximation of the drone dynamics, and 10^5 samples to fit a non-linear 3-layer MLP as the ship dynamics.

In training the safe RL methods, the reward in every step is the negative distance between the system's current state and the goal state, and the cost is 1 if the system is within the dangerous set \mathcal{S}_d and 0 otherwise. The threshold for expected cost is set to 0, which means we wish the system never enter the dangerous set (never reach a state with a positive cost). During training, the agent runs the system for 10^7 timesteps in total and performs 2000 policy updates. In each policy update, 100 iterations of gradient descent are performed. The implementation of the safe RL methods is based on [46].

All the methods are trained with static NPCs and tested on both static and moving NPCs. We believe this can make the testing more challenging and examine the generalization capability of the tested methods in different scenarios. All the agents are assigned random initial and goal locations in every simulation episode, which prevents the learned controller from overfitting a single configuration.

E. Experimental Results

a) Safety and goal-reaching performance: Results are shown in Fig. 3. Among the compared methods, our method is the only one that can reach a high task completion rate and relative safety rate at the same time. For other methods such as TRPO-Safe, when the controlled drone or ship is about to hit the NPCs, the learned controller tend to brake and decelerate. Thus, the agent is less likely to reach its goal when a simulation episode ends. The task completion rate and safety rate are opposite to each other for CPO, PPO-Safe and TRPO-Safe. On the contrary, the controller obtained by our method can maneuver smoothly among NPCs without severe deceleration. This enables the controlled agent to reach the goal location on time. Our method can also keep a relatively low tracking error, which means the difference between actual trajectories and reference trajectories is small.

b) Generalization capability to unseen scenarios: As is stated in Section V-D, the NPCs are static in training and can be either

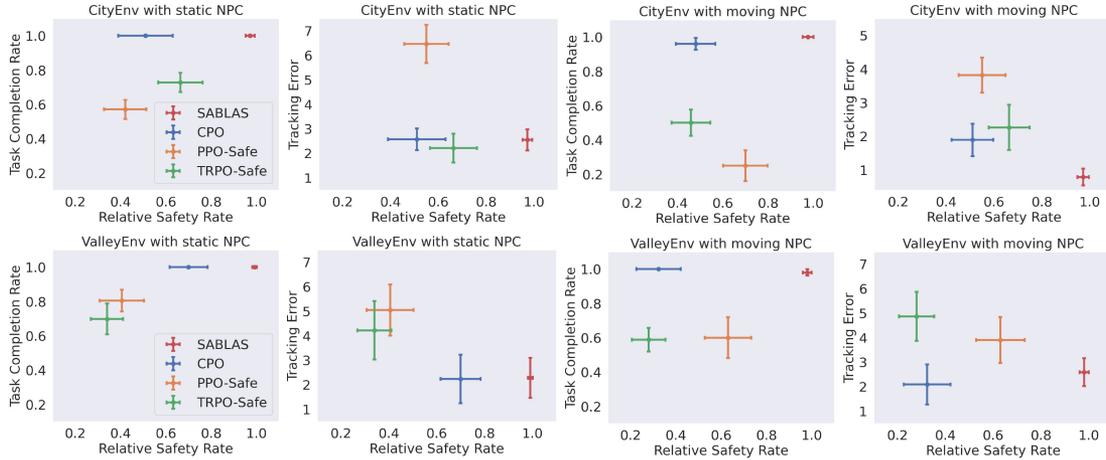


Fig. 3. Quantitative results of relative safety rate, task completion rate and tracking error. The center points are mean and error bars represent standard deviations. Results are collected over 50 random runs on each method. The first row is for drone control in CityEnv and the second row is for ship control in ValleyEnv. Our method reaches a high task completion rate and relative rate at the same time, while keeping a low tracking error.

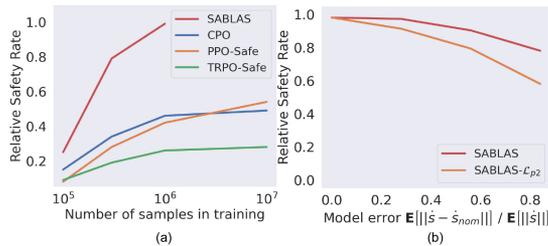


Fig. 4. (a) Comparison of sampling efficiency. With only 1/10 samples, SABLAS achieves a nearly perfect relative safety rate. (b) Influence of model error on safety performance. SABLAS- \mathcal{L}_{p2} uses (6) instead of the proposed (7) as loss function. SABLAS is tolerant to large error between nominal and real models while keeping a high safety rate.

static or moving in testing. Fig. 3 also demonstrates that our method has promising generalization capability across different training and testing scenarios.

c) Sampling efficiency: In Fig. 4(a), we show the safety performance under different sizes of the training set. The results are averaged over the drone and ship control tasks. SABLAS only needs around 1/10 of the samples required by the compared methods to achieve a nearly perfect relative safety rate. Note that SABLAS requires an extra 10^4 to 10^5 samples to fit the nominal dynamics, but this does not change the fact that the total number of samples needed by SABLAS is much fewer than the baselines.

d) Effect of model error: We investigate the influence of the model error $\|\dot{s} - \dot{s}_{nom}\|$ between the real dynamics and the nominal dynamics on the safety performance. We change the modeling error of the drone model and test the learning controller on CityEnv with static NPCs. We also perform an ablation study where we use \mathcal{L}_{p2} in (6) instead of \mathcal{L}_{p3} in (7) as loss function. The red curve in Fig. 4(b) show that SABLAS is tolerant to large model errors while exhibiting a promising safety rate. In our previous experiments, the model error $e = \mathbf{E}[\|\dot{s} - \dot{s}_{nom}\|] / \mathbf{E}[\|\dot{s}\|]$ is always less than 0.2. We did not encounter any difficulty fitting a nominal model with

empirical error $e \leq 0.2$. The orange curve in Fig. 4(b) shows that if we use \mathcal{L}_{p2} in (6), the trained controller will have a worse performance in terms of safety rate. This is because \mathcal{L}_{p2} only uses the nominal dynamics to calculate the loss, without leveraging the real black-box dynamics.

F. Discussion on Limitation

The main limitation of the proposed approach is that it cannot guarantee the satisfaction of the CBF conditions in (2) in the entire state space. Even if we minimize \mathcal{L}_0 , \mathcal{L}_d and \mathcal{L}_{p3} to 0 during training, the CBF conditions may still be occasionally violated during testing. After all, the training samples are finite and cannot cover the continuous state space. If the testing distribution and training distribution are the same, one can leverage the Rademacher complexity to give an error bound that the CBF conditions are violated, as is in Appendix B of [17]. But if the testing distribution is different from training, it is still unclear to derive the generalization error of the CBF conditions. To train CBF and controller that provably satisfy the CBF conditions, one can also use verification tools to find the counterexamples in the state space that violates the CBF conditions and add those counterexamples to the training set [23], [24]. The process is finished when no more counterexample can be found. However, the time complexity of the verification makes it not applicable for large and expressive neural networks. Also, the error between the nominal and real dynamics will have a negative impact on the safety performance. These limitations are left for future work.

VI. CONCLUSION AND FUTURE WORKS

We presented SABLAS, a general-purpose safe controller learning approach for black-box systems, which is supported by the theoretical guarantees from control barrier function theory, and at the same time is strengthened using a novel learning structure so it can directly learn the policies and barrier certificates for black-box dynamical systems. Simulation results show that SABLAS indeed provides a systematic way of learning safe

control policies with a great improvement over safe RL methods. For future works, we plan to study SABLAS on multi-agent systems, especially with adversarial players.

ACKNOWLEDGMENT

This letter solely reflects the opinions and conclusions of its authors and not DSTA Singapore or the Singapore Government.

REFERENCES

- [1] Z. Qin, K. Fang, Y. Zhu, L. Fei-Fei, and S. Savarese, "KETO: Learning keypoint representations for tool manipulation," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2020, pp. 7278–7285.
- [2] Z. Qin, Y. Chen, and C. Fan, "Density constrained reinforcement learning," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 8682–8692.
- [3] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *Comput. Res. Repository (CoRR)*, vol. abs/1707.06347, 2017.
- [4] T. P. Lillicrap *et al.*, "Continuous control with deep reinforcement learning," in *Proc. 4th Int. Conf. Learn. Representations*, 2016.
- [5] R. Freeman and P. V. Kokotovic, *Robust Nonlinear Control Design: State-Space and Lyapunov Techniques*. Berlin, Germany: Springer, 2008.
- [6] A. Isidori, E. Sontag, and M. Thoma, *Nonlinear Control Systems*, vol. 3. Berlin, Germany: Springer, 1995.
- [7] W. Lohmiller and J.-J. E. Slotine, "On contraction analysis for non-linear systems," *Automatica*, vol. 34, no. 6, pp. 683–696, 1998.
- [8] I. R. Manchester and J.-J. E. Slotine, "Control contraction metrics: Convex and intrinsic criteria for nonlinear feedback design," *IEEE Trans. Autom. Control*, vol. 62, no. 6, pp. 3046–3053, Jun. 2017.
- [9] D. Sun, S. Jha, and C. Fan, "Learning certified control using contraction metric," in *Proc. Conf. Robot Learn.*, 2020.
- [10] H. Tsukamoto and S.-J. Chung, "Neural contraction metrics for robust estimation and control: A convex optimization approach," *IEEE Control Sys. Letters (L-CSS)*, vol. 5, no. 1, pp. 211–216, 2021.
- [11] A. D. Ames, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs with application to adaptive cruise control," in *Proc. IEEE 53rd Annu. Conf. Decis. Control*, 2014, pp. 6271–6278.
- [12] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *Proc. 18th Eur. Control Conf.*, 2019, pp. 3420–3431.
- [13] U. Borrmann, L. Wang, A. D. Ames, and M. Egerstedt, "Control barrier certificates for safe swarm behavior," *IFAC-Papers-OnLine*, vol. 48, no. 27, pp. 68–73, 2015.
- [14] Y. Chen, A. Singletary, and A. D. Ames, "Guaranteed obstacle avoidance for multi-robot operations with limited actuation: A control barrier function approach," *IEEE Contr. Syst. Lett.*, vol. 5, no. 1, pp. 127–132, Jan. 2021.
- [15] R. Cheng, G. Orosz, R. M. Murray, and J. W. Burdick, "End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks," in *Proc. AAAI Conf. Artif. Intell.*, 2019, vol. 33, pp. 3387–3395.
- [16] R. Cheng, M. J. Khojasteh, A. D. Ames, and J. W. Burdick, "Safe multi-agent interaction through robust control barrier functions with learned uncertainties," in *Proc. 59th IEEE Conf. Decision Control (CDC)*, 2020, pp. 777–783.
- [17] Z. Qin, K. Zhang, Y. Chen, J. Chen, and C. Fan, "Learning safe multi-agent control with decentralized neural barrier certificates," in *Proc. Int. Conf. Learn. Representations*, 2021.
- [18] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Trans. Autom. Control*, vol. 62, no. 8, pp. 3861–3876, Aug. 2017.
- [19] A. Papachristodoulou and S. Prajna, "A tutorial on sum of squares techniques for systems analysis," in *Proc. Amer. Control Conf.*, 2005, pp. 2686–2700.
- [20] U. Topcu, A. Packard, and P. Seiler, "Local stability analysis using simulations and sum-of-squares programming," *Automatica*, vol. 44, no. 10, pp. 2669–2675, 2008.
- [21] J. Kapinski, J. V. Deshmukh, S. Sankaranarayanan, and N. Aréchiga, "Simulation-guided lyapunov analysis for hybrid dynamical systems," in *Proc. 17th Int. Conf. Hybrid Syst., Comput. Control*, 2014, pp. 133–142.
- [22] H. Ravanbakhsh and S. Sankaranarayanan, "Learning control lyapunov functions from counterexamples and demonstrations," *Auton. Robots*, vol. 43, no. 2, pp. 275–307, 2019.
- [23] Y.-C. Chang, N. Roohi, and S. Gao, "Neural lyapunov control," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 3245–3254.
- [24] H. Dai, B. Landry, L. Yang, M. Pavone, and R. Tedrake, "Lyapunov-stable neural-network control," *Robot. Sci. Syst.*, 2021.
- [25] S. Dean, A. J. Taylor, R. K. Cosner, B. Recht, and A. D. Ames, "Guaranteeing safety of learned perception modules via measurement-robust control barrier functions," in *Proc. Conf. Robot Learn.*, 2020.
- [26] J. Achiam, D. Held, A. Tamar, and P. Abbeel, "Constrained policy optimization," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 22–31.
- [27] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, "Trust region policy optimization," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 1889–1897.
- [28] C. Tessler, D. J. Mankowitz, and S. Mannor, "Reward constrained policy optimization," in *Proc. 7th Int. Conf. Learn. Representations*, 2019.
- [29] X. Chen and E. Hazan, "Black-box control for linear dynamical systems," in *Proc. 34th Conf. Learn. Theory*, 2021, pp. 1114–1143.
- [30] M. Fliess, C. Join, and H. Sira-Ramirez, "Complex continuous nonlinear systems: Their black box identification and their control," *IFAC Proceedings Volumes*, vol. 39, no. 1, pp. 416–421, 2006.
- [31] A. Levant, "Practical relative degree in black-box control," in *Proc. IEEE 51st IEEE Conf. Decis. Control*, 2012, pp. 7101–7106.
- [32] W. Grathwohl, D. Choi, Y. Wu, G. Roeder, and D. Duvenaud, "Back-propagation through the void: Optimizing control variates for black-box gradient estimation," in *Proc. 6th Int. Conf. Learn. Representations*, 2018.
- [33] F. Djeumou, A. P. Vinod, E. Goubault, S. Putot, and U. Topcu, "On-the-fly control of unknown smooth systems from limited data," in *Proc. Amer. Control Conf.*, 2021, pp. 3656–3663.
- [34] F. Djeumou, A. Zutshi, and U. Topcu, "On-the-fly, data-driven reachability analysis and control of unknown systems: An F-16 aircraft case study," in *Proc. 24th Int. Conf. Hybrid Syst., Comput. Control*, 2021, pp. 1–2.
- [35] T.-Y. Yang, J. Rosca, K. Narasimhan, and P. J. Ramadge, "Projection-based constrained policy optimization," in *Proc. 8th Int. Conf. Learn. Representations*, 2020.
- [36] C. Sun, D.-K. Kim, and J. P. How, "FISAR: Forward invariant safe reinforcement learning with a deep neural network-based optimizer," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2021, pp. 10617–10624.
- [37] J. Tordesillas, B. T. Lopez, and J. P. How, "FASTER: Fast and safe trajectory planner for flights in unknown environments," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2019, pp. 1934–1940.
- [38] J. Tordesillas and J. P. How, "MADER: Trajectory planner in multi-agent and dynamic environments," *IEEE Trans. Robot.*, to be published, doi: [10.1109/TRO.2021.3080235](https://doi.org/10.1109/TRO.2021.3080235).
- [39] J. J. Choi, D. Lee, K. Sreenath, C. J. Tomlin, and S. L. Herbert, "Robust control barrier-value functions for safety-critical control," *Comput. Res. Repository (CoRR)*, vol. abs/2104.02808, 2021.
- [40] F. Castañeda, J. J. Choi, B. Zhang, C. J. Tomlin, and K. Sreenath, "Pointwise feasibility of Gaussian process-based safety-critical control under model uncertainty," *Comput. Res. Repository (CoRR)*, vol. abs/2106.07108, 2021.
- [41] A. Paszke *et al.*, "PyTorch: An imperative style, high-performance deep learning library," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, vol. 32, pp. 8026–8037.
- [42] C. Fan, K. Miller, and S. Mitra, "Fast and guaranteed safe controller synthesis for nonlinear vehicle models," in *Computer Aided Verification*, S. K. Lahiri and C. Wang, Eds., Cham, Switzerland: Springer, 2020, pp. 629–652.
- [43] S. M. LaValle *et al.*, "Rapidly-exploring random trees: Progress and prospects," *Algorithmic Comput. Robot., New Directions*, vol. 5, pp. 293–308, 2001.
- [44] T. I. Fossen, "A survey on nonlinear ship control: From theory to practice," *IFAC Proceedings Volumes*, vol. 33, no. 21, pp. 1–16, 2000.
- [45] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. 3rd Int. Conf. Learn. Representations (ICLR)*, 2015.
- [46] A. Ray, J. Achiam, and D. Amodei, "Benchmarking safe exploration in deep reinforcement learning," in *Proc. Deep Reinforcement Learn. Workshop, Conf. Neural Inf. Process. Sys. (NeurIPS)*, vol. 7, 2019.