# Safety Filtering While Training: Improving the Performance and Sample Efficiency of Reinforcement Learning Agents

Federico Pizarro Bejarano<sup>®</sup>, Lukas Brunke<sup>®</sup>, and Angela P. Schoellig<sup>®</sup>

Abstract—Reinforcement learning (RL) controllers are flexible and performant but rarely guarantee safety. Safety filters impart hard safety guarantees to RL controllers while maintaining flexibility. However, safety filters can cause undesired behaviours due to the separation between the controller and the safety filter, often degrading performance and robustness. In this letter, we analyze several modifications to incorporating the safety filter in training RL controllers rather than solely applying it during evaluation. The modifications allow the RL controller to learn to account for the safety filter. This letter presents a comprehensive analysis of training RL with safety filters, featuring simulated and real-world experiments with a Crazyflie 2.0 drone. We examine how various training modifications and hyperparameters impact performance, sample efficiency, safety, and chattering. Our findings serve as a guide for practitioners and researchers focused on safety filters and safe RL.

*Index Terms*—Machine learning for robot control, reinforcement learning, robot safety, safety filtering.

# I. INTRODUCTION

**R** OBOTS are increasingly used for safety-critical applications such as autonomous driving [1] and surgery [2]. These tasks, characterized by complex cost functions and (possibly unknown) dynamics, are challenging for classical controllers [3]. This motivates learning-based controllers, especially reinforcement learning (RL) algorithms. Their ability to adapt to complex reward signals and unknown dynamics has led to superior performance in various domains [4]. However, a significant limitation of RL is the lack of safety guarantees [3]. This

Received 24 June 2024; accepted 10 November 2024. Date of publication 5 December 2024; date of current version 16 December 2024. This article was recommended for publication by Associate Editor Nadia Figueroa and Editor Jens Kober upon evaluation of the reviewers' comments. (*Corresponding author: Federico Pizarro Bejarano.*)

Federico Pizarro Bejarano is with the Learning Systems and Robotics Lab, University of Toronto, Toronto, ON M5S 1A1, Canada, and also with the University of Toronto Robotics Institute and the Vector Institute for Artificial Intelligence, Toronto, ON M5G 0C6, Canada (e-mail: federico.pizarrobejarano@robotics.utias.utoronto.ca).

Lukas Brunke and Angela P. Schoellig are with the Learning Systems and Robotics Lab, University of Toronto, Toronto, ON M5S 1A1, Canada, also with the University of Toronto Robotics Institute and the Vector Institute for Artificial Intelligence, Toronto, ON M5G 0C6, Canada, and also with the Technical University of Munich, Munich Institute for Robotics and Machine Intelligence (MIRMI), 80992 München, Germany (e-mail: lukas.brunke@robotics.utias.utoronto.ca; angela.schoellig@robotics.

This letter has supplementary downloadable material available at https://doi.org/10.1109/LRA.2024.3512374, provided by the authors.

Digital Object Identifier 10.1109/LRA.2024.3512374



Fig. 1. RL training with the safety filter modifications in red. An RL agent takes in the current state  $\mathbf{x}_k$  and outputs a (potentially unsafe) action  $\mathbf{u}_{uncert,k}$ . This action is passed to the safety filter, which outputs a minimally corrected safe action  $\mathbf{u}_{cert,k}$ . This safe action is applied to the environment, and the difference between the safe and unsafe action is used to penalize the reward received by the RL agent. When an episode is completed, the environment is reset to a safe state using the safety filter.

is undesirable for deployment in safety-critical scenarios despite promising results.

Safety filters can ensure that the RL controllers operate within defined constraints while minimally interfering with their operation. They determine whether uncertified (i.e., potentially unsafe) inputs from the controller will violate the constraints [5], [6]. If so, the filter determines the minimal deviation from the input that results in constraint satisfaction.

However, adding a safety filter changes how the controller interacts with the environment. By incorporating the safety filter in training, the RL algorithm can train on the system on which it will be evaluated (i.e., the system certified by the safety filter) [7]. This way, it can learn to account for the safety filter, maximizing rewards on the desired system. Additionally, using the safety filter during training can eliminate constraint violations during training, allowing RL controllers to be trained directly on physical systems. Finally, adding the safety filter significantly improves convergence as the RL controller will not explore unsafe states or cause complete failures (e.g., a drone crashing) and instead focuses on the desired safe states [7].

*Contributions:* In this letter, we analyze three modifications to the training process of any RL controller through the incorporation of a safety filter (see Fig. 1). We use model predictive safety

2377-3766 © 2024 IEEE. All rights reserved, including rights for text and data mining, and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.



Fig. 2. Experimental setup on a crazyflie 2.0 drone. A PPO controller trained without a safety filter or constraint violation penalization (blue) tries to track a reference trajectory (black), but unforeseen interactions with the safety filter cause poor tracking performance. However, when trained with a safety filter (green, with  $\alpha = 1$ ), the behaviour is smoother and more performant. The constraints are shown in red.

filters (MPSFs), but the training modifications can be equally applied with other types of safety filters. The training modifications are demonstrated in simulation using the safe-controlgym [8] on a quadrotor system and on a real-world Crazyflie 2.0 quadrotor (see Fig. 2). We found that the modifications significantly improve sample efficiency, eliminate constraint violations during training, and improve final performance and reduce chattering [9] on the certified system. Our code can be found at https://tinyurl.com/sf-train-code, and a video of our experiment can be found at https://tinyurl.com/sf-train-video.

## II. RELATED WORK

## A. Safety Filters

Safety filters can provide any controller with hard safety guarantees by determining the minimal input adjustments (corrections) necessary to ensure constraint satisfaction [5], [6]. Control barrier functions (CBFs) identify a safe set as the superlevel set of a function and guarantee safety through an online optimization that keeps the system in the safe set [10]. Generating CBFs requires knowledge of the system dynamics or offline data, they are typically defined in continuous time, and they do not guarantee input-constraint satisfaction [3]. Other methods of defining a safe set, such as Hamilton-Jacobi reachability analysis [6], also provide hard safety guarantees but can be equally difficult or computationally expensive to determine. Model predictive safety filters (MPSFs) leverage model predictive control (MPC) to assess the safety of proposed inputs [6], [11]. By simulating input trajectories using a system dynamics model over an MPC horizon, the MPSF can guarantee that a safe trajectory exists from the current state to a safe terminal set, guaranteeing safety and recursive feasibility. MPSFs tend to have higher online computation demands than other safety filters but require less offline computation.

Controllers are typically designed without consideration for the safety filter and thus will not choose optimal actions for the certified system. This may result in diminished performance and chattering [6], [9], where the controller persistently attempts to violate constraints, only to be continually stopped by the safety filter, resulting in oscillatory, inefficient behaviour.

## B. Reinforcement Learning

Reinforcement learning (RL) is a framework for sequential decision-making that aims to find a policy that maximizes the sum of attained rewards. RL typically assumes that the underlying dynamics function and reward signal are unknown. RL's lack of explicit assumptions and constraints limits safety but promotes expressiveness [3], [12].

Several strategies have been proposed to develop controllers with the expressiveness offered by RL while providing safety guarantees. We will focus on those that aim toward state and input constraint satisfaction.

One approach is to use a safety layer to convert an optimal but potentially unsafe action produced by the RL policy to the closest safe action with respect to the state constraints [13], [14], which is similar to the safety filters discussed in Section II-A. However, the optimization only attempts to prevent constraint violations in the next time step rather than guaranteeing safety for all future times.

Constrained Markov decision processes (CMDPs) [15] extend Markov decision processes (MDPs) with constraints. However, most of these approaches suffer from significant computational complexity, confining demonstrations to naive tasks [3]. One approach to solve the CMDP problem is to transform the problem into an unconstrained optimization using Lagrangian methods and use RL as a subroutine in the primal-dual updates [16].

While these methods encourage safety, they do not have the hard guarantees often required in safety-critical applications [3], [17]. The absence of a model of the system also tends to limit the sample efficiency of these methods [3], [18]. Additionally, these methods do not prevent constraint violations during training and may need to violate the constraints to learn a safe policy [16], making them unsuitable for training RL controllers directly on physical systems.

# C. Safety Filtering During Training

Filtering the actions of an RL agent during training has been approached using several methods [7]. State-wise safe sets are used in [17], [19], [20], CBFs are used in [18], [21], [22], and [11], [23], [24] use MPSFs. To encourage the RL agent to pursue safe actions, the reward signal has been penalized by a constant penalty [7], [24], a penalty proportional to the correction [11], [22], or a filter-specific penalty [25]. Differentiable safety layers maintain safety in training and evaluation while jointly optimizing the controller and safety filter [26], [27], but require additional assumptions such as perfect system knowledge. Training with a safety filter has been found to improve sample efficiency as the RL agent only explores safe states and occasionally improve the final performance [7], [19], [23].

This letter uses model predictive safety filters, which do not require an explicit representation of the safe set. However, the modifications described in Section V can also be applied to other safety filters [7]. Unlike approaches that aim to phase out the safety filter during training [23], [24], our focus is on the final *certified* performance. We examine the impact of each modification individually and in combination, alongside an analysis of how varying the correction penalty weight compares to constraint violation penalties [19], [28]. This work aims to compile, formally describe, and analyze training modifications with a safety filter, providing a resource on including safety filters into any RL training procedure.

## III. PROBLEM FORMULATION

We consider a discrete, time-invariant system given by

$$\mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k),\tag{1}$$

where  $\mathbf{x}_k \in \mathbb{X}$  represents the system state at time step  $k, \mathbf{u}_k \in \mathbb{U}$  denotes the control input, and  $\mathbf{f}$  encapsulates the system dynamics. The system is subject to known state and input constraints  $\mathbf{x} \in \mathbb{X}_c$ ,  $\mathbf{u} \in \mathbb{U}_c$ , where  $\mathbb{X}_c \subset \mathbb{X} \subset \mathbb{R}^n$  is closed and  $\mathbb{U}_c \subset \mathbb{U} \subset \mathbb{R}^m$  is compact.

We assume that we only have access to a nominal model f and compact uncertainty set W such that

$$\mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{w}(\mathbf{x}_k, \mathbf{u}_k), \qquad (2)$$

where  $\mathbf{w}(\mathbf{x}, \mathbf{u}) \in \mathbb{W} \subset \mathbb{R}^n$ ,  $\forall \mathbf{x} \in \mathbb{X}_c$ ,  $\forall \mathbf{u} \in \mathbb{U}_c$ .

A safety filter, which has access to the nominal model  $\mathbf{f}$ , takes in the current state of the system  $\mathbf{x}_k$  and input  $\mathbf{u}_{\text{uncert},k} = \pi_{\text{uncert}}(\mathbf{x}_k)$ , and uses the nominal system  $\mathbf{\bar{f}}$  and knowledge of the uncertainty set  $\mathbb{W}$  to determine if the proposed input is safe (i.e., satisfies state and input constraints, and will not lead to a constraint violation in the future). If it is not safe, it will find a safe input  $\mathbf{u}_{\text{cert},k}$  that minimally modifies the uncertified input. This certified input is the input applied to the system (i.e.,  $\mathbf{u}_k = \mathbf{u}_{\text{cert},k}$ ).

## IV. BACKGROUND

## A. Model Predictive Safety Filters

MPSFs can be based on any MPC framework, inheriting the safety and recursive feasibility guarantees of the underlying MPC [9]. A general formulation of MPSFs based on robust tube-based MPC is provided below to motivate MPSFs. The specific implementation used for experiments is the robust non-linear MPC described in [29]. This is further elaborated on in Section VI.

Definition 1 (Robust pos. control inv. (RPCI) set [30]): A set  $\mathbb{P} \subseteq \mathbb{X}_c$  is robust positively control invariant (RPCI) for the system in (2) with a controller  $\pi_{\text{RPCI}} : \mathbb{P} \to \mathbb{U}_c$  if  $\forall \mathbf{x} \in \mathbb{P}$  and  $\forall \mathbf{w} \in \mathbb{W}$ :

$$\bar{\mathbf{f}}(\mathbf{x}, \pi_{\mathrm{RPCI}}(\mathbf{x})) + \mathbf{w} \in \mathbb{P}.$$

Assumption 1: There exists a terminal set  $\mathbb{X}_{term} \subset \mathbb{X}_c$  and a terminal controller  $\pi_{term} : \mathbb{X}_{term} \to \mathbb{U}_c$  such that  $\mathbb{X}_{term}$  is a RPCI set for the system in (2) under  $\pi_{term}$ .

An MPSF solves an optimization problem at each time step for an optimal input sequence over the next H time steps, where  $H \in \mathbb{N}$  is the horizon. At each time step k, the MPSF only applies the first input from the resulting input sequence. The optimization problem for a general robust, tube-based multi-step MPSF can be stated as follows [9]:

$$\min_{\mathbf{u}_{:k}, \mathbb{X}_{:k}} \sum_{j=0}^{M-1} w(j) \| \pi_{\text{uncert}}(\mathbf{z}_{j|k}) - \mathbf{u}_{j|k} \|_2^2$$
(3a)

s.t. 
$$\mathbf{x}_k \in \mathbb{X}_{0|k}, \, \mathbf{z}_{0|k} = \mathbf{x}_k,$$
 (3b)

$$\mathbb{X}_{i+1|k} \supseteq \mathbf{\Phi}(\mathbb{X}_{i|k}, \mathbf{u}_{i|k}, \mathbb{W}), \tag{3c}$$

$$\mathbf{z}_{i+1|k} = \mathbf{\bar{f}}(\mathbf{z}_{i|k}, \mathbf{u}_{i|k}), \tag{3d}$$

$$\mathbb{X}_{i|k} \subseteq \mathbb{X}_{c},\tag{3e}$$

$$\mathbf{u}_{i|k} \in \mathbb{U}_{\mathbf{c}}, \ \forall \ i = 0, \dots, H-1,$$
(3f)

$$\mathbb{X}_{H|k} \subset \mathbb{X}_{\text{term}},\tag{3g}$$

where  $\mathbf{u}_{i|k}$  is the input at the (k + i)th time step computed at time step k, M is the filtering horizon,  $w(\cdot) : \mathbb{N}_0 \to \mathbb{R}^+$  calculates the weights associated with the *j*th correction,  $\pi_{\text{uncert}}$  is the uncertified controller,  $\mathbf{z}_{j|k}$  is the estimated future state at the (k + j)th time step computed at time step k,  $\mathbb{X}_{i|k}$  is the set of possible states at the (k + i)th time step computed at time step k, and the evolution of the system is  $\mathbb{X}_{i+1|k} \supseteq \Phi(\mathbb{X}_{i|k}, \mathbf{u}_{i|k}, \mathbb{W}) =$  $\{\bar{\mathbf{f}}(\mathbf{x}, \mathbf{u}_{i|k}) + \mathbf{w} \mid \forall \mathbf{x} \in \mathbb{X}_{i|k}, \mathbf{w} \in \mathbb{W}\}$ . The future actions of the controller  $\pi_{\text{uncert}}(\mathbf{z}_{j|k})$  are approximated using [9].

In [9], the norm of the rate of change of the inputs metric is proposed. This measures how much the input varied during an experiment. Consider the applied inputs  $\mathbf{u}_k$  for k = 0, ..., K - 1. First, we take the numerical derivative  $\delta \mathbf{u}_k = (\mathbf{u}_k - \mathbf{u}_{k-1})/\delta t$  for k = 1, ..., K - 1 (where  $\delta t$  is the length of a time step) and stack them into a matrix  $\Delta \mathbf{u} = [\delta \mathbf{u}_1, ..., \delta \mathbf{u}_{K-1}]$ . Then the norm of the rate of change of the inputs is  $\|\Delta \mathbf{u}\|_F$ , where  $\|\cdot\|_F$  is the Frobenius norm. We will use this metric to evaluate chattering and jerkiness.

#### B. Training Reinforcement Learning Algorithms

RL algorithms learn a policy to maximize the accumulated rewards by interacting with the environment. RL assumes the control problem is a Markov decision process (MDP). An MDP comprises a state space  $\mathbb{X}$ , an input (action) space  $\mathbb{U}$ , a set of starting states  $\mathbb{S} \subset \mathbb{X}$ , a transition model  $\mathcal{T} : \mathbb{X} \times \mathbb{U} \to \mathbb{X}$ , and a per-step reward function  $\mathcal{R} : \mathbb{X} \times \mathbb{U} \to \mathbb{R}$  [12]. RL algorithms typically assume the dynamics  $\mathcal{T}$  and reward function  $\mathcal{R}$  are unknown [3].

RL aims to find the optimal policy  $\pi^* : \mathbb{X} \to \mathbb{U}$  which maximizes the discounted accumulated reward (i.e., the return) from any starting state  $\mathbf{x}_0 \in \mathbb{S}$  [12], [17]:

$$G_{\pi}(\mathbf{x}_0) = \mathcal{R}(\mathbf{x}_0, \pi(\mathbf{x}_0)) + \gamma \mathcal{R}(\mathbf{x}_1, \pi(\mathbf{x}_1)) + \dots$$
(4)

$$=\sum_{k=0}^{\infty}\gamma^{k}\mathcal{R}(\mathbf{x}_{k},\pi(\mathbf{x}_{k})),$$
(5)

where  $G_{\pi}(\mathbf{x}_0)$  is the return starting from state  $\mathbf{x}_0$  under the policy  $\pi$ ,  $\gamma \in (0, 1]$  is a discount factor, and  $\mathbf{x}_{k+1} = \mathcal{T}(\mathbf{x}_k, \pi(\mathbf{x}_k)), \forall k \ge 0.$ 

The agent improves the policy by interacting with the environment. At each time step k, the agent receives an observation  $\mathbf{x}_k \in \mathbb{X}$ , takes an action  $\mathbf{u}_k \in \mathbb{U}$ , transitions to a next state  $\mathbf{x}_{k+1} = \mathcal{T}(\mathbf{x}_k, \mathbf{u}_k)$ , and receives a reward  $r_k = \mathcal{R}(\mathbf{x}_k, \mathbf{u}_k)$ . These four values form a quadruple  $(\mathbf{x}_k, \mathbf{u}_k, \mathbf{x}_{k+1}, r_k)$ , which is used to improve the policy. On-policy RL algorithms require the quadruples to be from the current policy. In contrast, off-policy algorithms can learn from data generated by other policies [12].

Training is typically divided into episodes, where a starting state  $\mathbf{x}_0$  is sampled from S. The policy controls the system for a number of steps up to a maximum K, or until another condition is met [12]. Then, the system is reset to a new starting point for the next episode.

#### V. TRAINING MODIFICATIONS

We consider three modifications to the training of RL algorithms [7]. These can be combined or used separately and can be applied to any RL controller and any safety filter, not only MPSFs.

Authorized licensed use limited to: The University of Toronto. Downloaded on December 19,2024 at 20:48:54 UTC from IEEE Xplore. Restrictions apply.

## A. Filtering Training Actions

During RL training, the controller generates uncertified actions  $\mathbf{u}_{\text{uncert},k} \in \mathbb{U}$ . By applying the safety filter  $\mathbf{u}_{\text{cert},k} = \pi_{\text{SF}}(\mathbf{x}_k, \mathbf{u}_{\text{uncert},k})$  to certify these actions, safety is guaranteed during training.

For on-policy algorithms, the uncertified action generated by the RL controller must be the one buffered for policy improvement. Thus the buffered quadruple is  $(\mathbf{x}_k, \mathbf{u}_{uncert,k}, \mathbf{f}(\mathbf{x}_k, \mathbf{u}_{cert,k}), \mathcal{R}(\mathbf{x}_k, \mathbf{u}_{cert,k}))$  [7]. Effectively, the RL algorithm operates in a safe environment protected by the safety filter rather than in the true, unsafe environment. For off-policy algorithms, this same approach can be done, or the safety filter can be considered an expert, and thus the safe action  $\mathbf{u}_{cert,k}$  can be buffered rather than the unsafe action. A detailed comparison of the learning quadruples and their effect on training can be found in [7].

This modification aims to maximize the reward on the final certified system, prevent violations during training (assuming that the episodes start from safe states, see Section V-C), and improve sample efficiency by focusing the training on the safe areas. However, the RL agent may grow dependent on the safety filter and thus not perform well without the safety filter [17], [22], [24]. This is not a problem if the safety filter is used during evaluation, as assumed in this letter, and can be mitigated using the following approach [22].

## **B.** Penalizing Corrections

Penalizing constraint violations is a common approach to encouraging safety in RL [28]. This is typically done by introducing a constant penalty  $\beta \ge 0$  to the reward when a constraint is violated [19], [24], or replacing the reward entirely with the negative penalty  $-\beta$  [17]. With the safety filter, we can penalize corrections rather than constraint violations. Corrections indicate that the proposed action was unsafe before a constraint violation, and the magnitude of the correction measures how unsafe the action was.

In [7], [24], the reward is penalized by a constant penalty when the safety filter is activated. However, this does not incorporate how unsafe the proposed action is. In [11], [22], the reward is penalized by the magnitude of the correction. Filter-specific penalties have been proposed [25] but cannot be applied with other safety filters. The general framework for utilizing the safety filter to penalize the reward is:

$$\begin{aligned} &\mathcal{R}^{\alpha}(\mathbf{x}_{k},\mathbf{u}_{\text{uncert},k},\mathbf{u}_{\text{cert},k}) \\ &= \mathcal{R}(\mathbf{x}_{k},\mathbf{u}_{\text{applied},k}) - \alpha \mathcal{P}(\mathbf{u}_{\text{uncert},k},\mathbf{u}_{\text{cert},k}) \end{aligned}$$

where  $\alpha$  weighs the correction penalty,  $\mathcal{P}$  defines the penalty based on the correction, and  $\mathbf{u}_{applied,k}$  is the action applied to the system ( $\mathbf{u}_{uncert,k}$  in standard RL training, and  $\mathbf{u}_{cert,k}$  if filtering the training actions as discussed in the previous section). After experimentation, we found that penalizing the magnitude of the correction  $\mathcal{P} = \|\mathbf{u}_{uncert,k} - \mathbf{u}_{cert,k}\|_2^2$  [11], [22] led to the best results in our experiments.

This modification encourages safety and minimizes corrections but does not enforce safety during training or maximize the true reward. This reward-shaping modification should converge to an optimal, safe solution if combined with the previous section's filtering training actions modification. This is because an optimal RL policy will never propose an action that causes a correction (as it would be penalized and corrected); thus, every action will be safe. However, in practice, the weight  $\alpha$  will affect the behaviour of the resulting controller, and there are no guarantees that the converged policy will always produce safe actions [12], [22].

## C. Safely Resetting the Environment

Sample efficiency can be improved by using the safety filter to avoid initiating an episode in an unsafe state [31]. Often, there is no prescribed starting state, and thus  $\mathbb{S} := \mathbb{X}_c$ . However, even if  $\mathbf{x}_0 \in \mathbb{X}_c$ , this does not guarantee that there exists an input trajectory that will keep the system safe for all future time. Using the safety filter, safety can be guaranteed by restricting  $\mathbb{S}$  to a *H*-step robust positively control invariant (*H*-RPCI) set [31].

An *H*-RPCI set guarantees that the system can be safely driven to a RPCI set (see Definition 1) in *H* steps. We wish to find the *H*-RPCI set that can be driven to the terminal set  $X_{\text{term}}$ . The *H*-RPCI set can be explicitly computed and sampled from, but this can be difficult [32]. Instead, we will sample  $x_0 \sim S$ , then determine the feasibility of certifying an input from that state [31]. If the safety filtering optimization is feasible,  $x_0$  is within the *H*-RPCI set. If infeasible, another starting state is randomly generated until a feasible starting state is found.

This modification ensures that episodes only commence from certifiably safe states, improving sample efficiency as the RL training focuses on feasible trajectories [31]. When combined with filtering training actions (Section V-A), this ensures the constraints are never violated during training.

## VI. EXPERIMENTAL RESULTS

To determine the efficacy of the training modifications, we ran experiments in the safe learning-based control simulation environment safe-control-gym [8] and on a real quadrotor, the Crazyflie 2.0 (see Fig. 2). The underlying MPC is the robust nonlinear MPC detailed in [29], which assumes that the system is incrementally stabilizable and that the model mismatch between the nominal system and the real system is contained within a known compact set  $\mathbb{W}$ . The upper bound on the model mismatch  $w_{\max} = \max_{k \in \{0,...,K-1\}} \|\mathbf{x}_{k+1} - \mathbf{z}_{1|k}\|_2$ was found experimentally by comparing the true states  $x_{k+1}$  and the states predicted by the nominal model  $\mathbf{z}_{1|k}$  while executing trajectories for system identification (further described below). We define  $\mathbb{W} = \{ \mathbf{w} \in \mathbb{R}^n \mid ||\mathbf{w}||_2 \le w_{\max} \}$ . The terminal set is determined from the uncertainty set  $\mathbb{W}$  and a hand-tuned LQR terminal controller  $\pi_{\text{term}}$  [29]. We use acados [33] to efficiently solve the MPC optimization at each time step.

During training, we do not use early stopping (i.e., ending an episode when a constraint is violated as used in [17]), as it was found to prevent convergence when not filtering actions. The MPSF has a filtering horizon M = 1 during training and M = 2during evaluation, which was found to reduce chattering in all approaches while minimally increasing computation. The state constraints were softened to allow for feasibility at all time steps, which was found to improve simulation results and allow for the successful execution of real-world experiments. This was done by incorporating slack variables and penalizing their magnitude with linear and quadratic terms.

Proximal policy optimization (PPO) [34] was used as the RL controller due to its popularity in robotics tasks [8]. Each approach is trained five times with different seeds to determine the variability in the results. Safe explorer [13] and constrained policy optimization [16], two popular safe model-free RL approaches, were tested and found to not converge for performed

Authorized licensed use limited to: The University of Toronto. Downloaded on December 19,2024 at 20:48:54 UTC from IEEE Xplore. Restrictions apply.

tasks, likely due to the tasks' high dimensionality and difficulty in ensuring safety.

The reward signal for the RL algorithm at time step k is  $r_k = \exp(-2\|\mathbf{p}_k - \mathbf{p}_{\text{ref},k}\|_2^2)$ , where  $\mathbf{p}_k$  is the current position of the drone and  $\mathbf{p}_{\text{ref},k}$  is the reference position at time step k. During training, we use two modified reward signals. For the baseline approaches we penalize constraint violations by a constant penalty,  $r_k^\beta = r_k - \mathbf{1}_{\text{viol}}\beta$ , where  $\mathbf{1}_{\text{viol}}$  is an indicator function for violations, equal to 1 when a violation occurs and is otherwise zero, and  $\beta \ge 0$  is the constraint violation penalty. For the approaches which penalize corrections, we use  $r_k^\alpha = r_k - \alpha \|\mathbf{u}_{\text{uncert},k} - \mathbf{u}_{\text{cert},k}\|_2^2$  [11], [22], where  $\alpha$  is the correction penalization weight. The maximum reward achievable is  $r_k = r_k^\alpha = r_k^\beta = 1$ , for a total undiscounted return equal to the maximum number of iterations. The return in this section always refers to the undiscounted return (i.e.,  $\gamma = 1$ ) without using any penalities (i.e., using  $r_k$  as the reward).

## A. Ablation Study of the Training Modifications

Every combination of the modifications was trained and evaluated to analyze the effects of the training modifications. Additionally, some combinations are similar to the related works that only filter actions (e.g., [17], [21], [23]) and filter actions as well as penalize corrections [24], [25]; thus, this ablation study also serves to better analyze the related works. In the following section "Std. ( $\beta$ =0)" refers to the baseline with none of the training modifications nor constraint violation penalties. The other approaches are combinations of the training modifications: FA = Filtering Actions, PC = Penalizing Corrections, SR = Safe Reset.

The controllers were evaluated on a simulation of a Crazyflie 2.0 using the safe-control-gym [8]. The state is  $\mathbf{x} = [x, \dot{x}, y, \dot{y}, z, \dot{z}, \phi, \theta, \psi, \dot{\phi}, \dot{\theta}, \dot{\psi}]^{\mathsf{T}}$ , where the x - y - z values are the position,  $\dot{x} - \dot{y} - \dot{z}$  are the velocities,  $\phi - \theta - \psi$  are the roll-pitch-yaw angles, and  $\dot{\phi} - \dot{\theta} - \dot{\psi}$  are the body rates of those angles. The inputs are the thrusts applied to the four rotors. The nominal equations of motion can be found in [35]. These equations were discretized to get the nominal discrete model of the 3D quadrotor, which was very accurate, with  $w_{\text{max}} = 0.0205$ . This uncertainty was determined by starting at random states in  $\mathbb{X}_c$ , executing random actions in  $\mathbb{U}_c$ , and measuring the model mismatch. The MPSF had an MPC horizon of H = 20 and a frequency of 50 Hz. The correction penalization weight is set to  $\alpha = 1$  when using the correction penalization approach. The constraint violation penalty is set to  $\beta = 0$  in all cases.

The trajectory tracking task consists of tracking a figure-eight reference with an amplitude of 2 m in the x dimension, 0.5 m in the y dimension, 1m in the z dimension, which must be flown once in 5 s. The total number of iterations, and thus the maximum return, is 50 Hz  $\cdot$  5 s = 250. The position is constrained to be 5% smaller than the full extent of the trajectory, the velocities are constrained to be within  $[-2 \text{ ms}^{-1}, 2 \text{ ms}^{-1}]$ , the angles were constrained to be within  $[-2 \text{ rads}^{-1}, 2 \text{ ms}^{-1}]$ .

To find suitable starting points for evaluating the final trained models, certifiably safe states were found using the MPSF (using the process described in Section V-C). One hundred starting points were generated, and each algorithm was tested using these starting states. The starting states and the different training seeds are the source of randomness.



Fig. 3. The return for the ablation study of the training modifications.



Fig. 4. The rate of change of the inputs (see [9]) for the ablation study of the training modifications.



Fig. 5. The return versus the number of training steps on the ablation study of the training modifications.

From Fig. 6, we note that penalizing the corrections reduces the number of constraint violations during training, lowers the rate of change of the inputs (see Fig. 4), and increases the return (see Fig. 3), although this depends on the correction penalization weight  $\alpha$  (see next section). Secondly, we find that the safe reset modification significantly improves convergence (see Fig. 5) and evaluation return. Finally, filtering the actions reduces the constraint violations to nearly zero, especially when partnered with the safe reset approach. Most importantly, combining all the modifications leads to the best return and convergence and the least constraint violations during training. Additionally, it results in the third lowest rate of change of the inputs.

The standard approach "Std.  $\beta=0$ " is significantly faster than the approaches using the modifications (see Fig. 7), followed by the approach that only uses the safe reset modification, "SR". The approach with all modifications "FA,PC,SR" is five

TABLE I Results for the Ablation on the Correction Penalization Weight  $\alpha$  and the Constraint Violation Penality  $\beta$ .

Metric	Std. ( $\beta$ =0)	Std. ( $\beta$ =0.01)	Std. ( $\beta$ =0.1)	Std. ( $\beta$ =1)	Safe ( $\alpha$ =0.1)	Safe ( $\alpha$ =1)	Safe ( $\alpha$ =10)	Safe ( $\alpha$ =100)
Return	$200.2 \pm 17.1$	$208.3 \pm 15.5$	$210.3 \pm 14.4$	$199.1 \pm 15.6$	$212.6 \pm 13.8$	$\textbf{214.1} \pm \textbf{14.0}$	$210.7 \pm 14.0$	$202.4 \pm 13.4$
Return when uncertified	$222.1 \pm 13.7$	$\textbf{229.0}\pm\textbf{10.1}$	$222.2 \pm 10.9$	$206.8 \pm 14.6$	$11.2 \pm 15.7$	$31.2 \pm 44.8$	$210.7 \pm 50.0$	$202.3 \pm 19.1$
Input rate of change $[m s^{-1}]$	$16.4 \pm 17.0$	$10.0 \pm 13.8$	$4.7\pm2.9$	$3.2 \pm 3.0$	$9.4 \pm 2.9$	$7.5 \pm 3.3$	$3.8 \pm 3.0$	$3.0\pm3.1$
Training constraint violations [%]	$82.8 \pm 6.6$	$78.0 \pm 5.4$	$67.0 \pm 4.2$	$74.5 \pm 9.9$	$0.23 \pm 0.02$	$0.22 \pm 0.03$	$0.22 \pm 0.02$	$0.22~\pm~0.02$
Training time per step [ms]	$2.3 \pm 0.3$	$2.7\pm0.2$	$2.8\pm0.4$	$2.8\pm0.2$	$14.6 \pm 1.2$	$11.6 \pm 1.5$	$13.5 \pm 1.0$	$12.6 \pm 1.1$

The bold values are the columns that get the best value in that row.



Fig. 6. The number of constraint violations versus the number of training steps in the ablation study of the training modifications.



Fig. 7. The time required to execute one training step in the ablation study of the training modifications.

times slower than training without any modifications. However, the approaches trained with the action filtering and safe reset modifications reached 80% of their final return ten times faster than the standard training and surpassed a return of 200 in 2.5 times fewer environment interactions (see Fig. 5). This sample efficiency is important when environment interactions are slow or expensive, such as training on a real system. It should also be noted that the training time depends on the type of safety filter and its implementation. A control barrier function (CBF) based safety filter can be significantly faster than an MPSF [5], [6] and could equally be used with the training modifications. However, designing a CBF safety filter has other challenges, like determining the control barrier function, which remains an active area of study [3].

#### B. Ablation Study of Reward Penalties

To study the effects of reward penalties, we conducted experiments with various values of the correction penalization weight  $\alpha$  and the constraint violation penalty  $\beta$ . We compare the safe approach (all of the training modifications together,



Fig. 8. The return versus the number of training steps on the ablation study of reward penalties.

denoted "Safe") with  $\alpha \in \{0.1, 1, 10, 100\}$  to the standard training (none of the training modifications, denoted "Std.") with  $\beta \in \{0, 0.01, 0.1, 1\}$ . These experiments were performed with the same setup as the ablation study of the training modifications. The results are summarized in Table I.

In the experiments (evaluated with the safety filter), we note that as  $\alpha$  and  $\beta$  increase, the rate of change of the inputs decreases. The return for the safe approach is maximized by  $\alpha = 1$ , while the return is maximized for the baseline with  $\beta = 0.1$ . The baseline did not converge with  $\beta > 1$ .

The safe approach increases the return by up to 7.3% compared to the standard approaches with  $\beta = 0$ , and up to 2.0% compared to the best performing standard approach,  $\beta = 0.1$ . The rate of change of the inputs is reduced by up to 75% compared to the baseline with  $\beta = 0$ , and up to 7.8% compared to the smoothest baseline,  $\beta = 1$ . This demonstrates that the safe training approach effectively reduces chattering and jerky motions due to undesirable interplay between the safety filter and the controller and increases the controller's performance when certified.

Further, let us consider the convergence of the safe approach. From Fig. 8, the fastest converging safe approach ("Safe  $\alpha =$ 10") reached 80% of its final performance seven times faster than the fastest converging standard training ("Std.  $\beta = 0.01$ ") and surpassed a return of 200 with 1.8 times fewer environment interactions. We found that as the complexity and difficulty in ensuring the safety of the experiment increased, the benefit of the safe approach in terms of convergence increased, and the convergence of the baseline approaches slowed significantly. Additionally, from Table I, when using the safe approach, only 0.22% of the training steps violate constraints. In contrast, the baseline with  $\beta = 0$  and  $\beta = 0.1$  (the  $\beta$  value resulting in the least constraint violations) violate the constraints 82.8% and 67.0% of the time, respectively. The safe approaches' small number of constraint violations is due to softening the constraints in the MPSF formulation, which improved training efficiency. There were no constraint violations during the evaluation.

 TABLE II

 Results for the Real Experiments of five Trials Flying a Crazyflie 2.0 on a Figure-Eight Path

Metric	Std. ( $\beta$ =0)	Std. ( $\beta$ =1)	Safe ( $\alpha$ =0.1)	Safe ( $\alpha$ =1)	Safe ( $\alpha$ =10)
Return	$263.5\pm7.7$	$311.7 \pm 4.4$	$314.2 \pm 9.7$	$\textbf{316.1} \pm \textbf{1.4}$	$305.5 \pm 0.8$
Input rate of change $[m s^{-1}]$	$237.3 \pm 15.9$	$143.8\pm5.0$	$234.8 \pm 13.0$	$199.9 \pm 6.0$	$154.0 \pm 8.7$
Training constraint violations [%]	$47.7 \pm 14.6$	$14.3 \pm 12.8$	$11.9 \pm 10.9$	$6.3 \pm 7.5$	$\textbf{2.7} \pm \textbf{3.0}$
Training time per step [ms]	$14.4 \pm 0.1$	$14.2\pm0.1$	$16.4 \pm 0.1$	$16.4 \pm 0.1$	$16.1 \pm 0.1$

The bold values are the columns that get the best value in that row.



Fig. 9. The number of constraint violations during evaluation when not using a safety filter for the ablation study of reward penalties. With the safety filter, there are no constraint violations.

When the safety filter is removed during evaluation, all the approaches violate the constraints (see Fig. 9), although increasing  $\alpha$  or  $\beta$  leads to fewer constraint violations. The return of approaches trained without the safety filter is reduced with the addition of the safety filter, demonstrating that they do not operate optimally when the safety filter is applied in evaluation (see "Return when uncertified" in Table I). However, the approaches trained with the safety filter have the same or better return with the safety filter. This highlights the importance of not removing the safety filter, even when penalizing corrections, as safety guarantees are lost and performance may decrease.

## C. Real Experiments

The safe approach was tested on a real Crazyflie 2.0 quadrotor with a desired trajectory of a figure-eight in the x - y plane with an amplitude of 2 m, which must be completed twice in 15 s. The controllers were trained in a simulation of the real Crazyflie based on the sim-to-real stack from [36], simulating both its dynamics and internal software. Additionally, the controllers were trained with domain randomization and added Gaussian noise to the dynamics and observations to improve their ability to generalize from the simulated environment to the real one. The trained controllers were then evaluated on the real Crazyflies.

We consider the state  $\mathbf{x} = [x, \dot{x}, y, \dot{y}, \phi, \theta]^{\mathsf{T}} \in \mathbb{R}^6$ . The x and y positions of the drone are constrained to be within [-0.95 m, 0.95 m], the velocities of the drone in x and y are constrained to be within  $[-2 \text{ ms}^{-1}, 2 \text{ ms}^{-1}]$ , and the roll  $\phi$  and pitch  $\theta$  angles are constrained to be within [-0.25 rad, 0.25 rad]. The input  $\mathbf{u} = [\phi, \theta]^{\mathsf{T}} \in \mathbb{R}^2$  are the desired roll and pitch angles which the quadrotor tracks using an internal PID controller. The commanded angles are constrained to be within [-0.25 rad, 0.25 rad]. The closed-loop dynamics in the x - y plane can be approximated as a linear system. The linear system was identified from the uncertified flights of the baseline approaches, using simulation data during training and real data during real evaluation.

The maximum model mismatch was significantly higher than for the simulation experiments, with  $w_{\text{max}} = 0.1267$ . However, there were a significant number of outliers, so rather than using the maximum error, we used the mean error plus three standard deviations,  $w_{\sigma=3} = 0.0515$ . A motion capture system measures the quadrotor pose during the real experiment.

The safe approaches (all of the training modifications together, denoted "Safe") with  $\alpha \in \{0.1, 1, 10\}$  and the baselines (none of the training modifications, denoted "Std.") with  $\beta \in \{0, 1\}$  are tested five times each with an MPC horizon of H = 10 and a control frequency of 25 Hz. The total number of iterations, and thus the maximum return, is 25 Hz · 15 s = 375. Every test starts at the same initial position, the origin of the x - y plane, and a height of 1 m.

As seen in Table II, the safe approach reduces the rate of change of the inputs by up to 35% compared to the  $\beta = 0$  baseline and increases the total return by up to 20%. Compared to the  $\beta = 1$  baseline, the return is increased by up to 1.5%, and the norm of the rate of change of the inputs is increased by 7%. This is likely due to the simulation experiments and the conservativeness of the  $\beta = 1$  baseline. Interestingly, the safe approach provides a greater proportional increase in the return of real experiments compared to simulation. This may be due to the increased conservativeness of the safety filter, which causes a larger distribution shift between the certified and uncertified environments.

Once again, the safe approach significantly decreases constraint violations in training despite softened constraints, domain randomization, and added Gaussian noise. We reduce constraint violations by over 17 times compared to the  $\beta = 0$  baseline and five times compared to the  $\beta = 1$  baseline. The baseline approaches execute training steps 13% faster than the safe approaches, compared to five times faster in the simulation sections (see Fig. 7 and Table I). This is because the simulation used in this section includes the internal software of the Crazyflies [36], significantly increasing the time per training step. This highlights the importance of improving sample efficiency.

### VII. CONCLUSION

This letter analyzes three methods of incorporating safety filters in the training of model-free RL algorithms: filtering the training actions, penalizing the reward by the magnitude of the corrections, and resetting episodes to safe starting states. The presented modifications require little tuning and have been designed to be easily incorporated into the training of any RL algorithm. In simulation and real experiments done using model predictive safety filters, we found that the training modifications improve performance and sample efficiency, reduce chattering, and nearly eliminate training-time constraint violations. These modifications further leverage safety filters, allowing RL algorithms to be effectively applied to safety-critical systems.

Authorized licensed use limited to: The University of Toronto. Downloaded on December 19,2024 at 20:48:54 UTC from IEEE Xplore. Restrictions apply.

#### REFERENCES

- K. Burnett et al., "Zeus: A system description of the two-time winner of the collegiate SAE AutoDrive competition," *J. Field Robot.*, vol. 38, pp. 139–166, 2021.
- [2] A. A. Shvets, A. Rakhlin, A. A. Kalinin, and V. I. Iglovikov, "Automatic instrument segmentation in robot-assisted surgery using deep learning," in in *Proc. 17th IEEE Int. Conf. Mach. Learn. Appl.*, 2018, pp. 624–628.
- [3] L. Brunke et al., "Safe learning in robotics: From learning-based control to safe reinforcement learning," *Annu. Rev. Control Robot., Auton. Syst.*, vol. 5, no. 1, pp. 411–444, 2022.
- [4] D. Silver et al., "Mastering the game of go with deep neural networks and tree search," *Nat.*, vol. 529, pp. 484–489, 2016.
- [5] K.-C. Hsu, H. Hu, and J. F. Fisac, "The safety filter: A unified view of safety-critical control in autonomous systems," *Annu. Rev. Control Robot.*, *Auton. Syst.*, vol. 5, pp. 47–72, 2023.
- Auton. Syst., vol. 5, pp. 47–72, 2023.
  [6] K. P. Wabersich et al., "Data-driven safety filters: Hamilton-jacobi reachability, control barrier functions, and predictive methods for uncertain systems," *IEEE Control Syst. Mag.*, vol. 43, no. 5, pp. 137–177, Oct. 2023.
- [7] H. Krasowski, J. Thumm, M. Müller, L. Schäfer, X. Wang, and M. Althoff, "Provably safe reinforcement learning: Conceptual analysis, survey, and benchmarking," *Trans. Mach. Learn. Res.*, 2023.
  [8] Z. Yuan et al., "Safe-control-gym: A unified benchmark suite for safe
- [8] Z. Yuan et al., "Safe-control-gym: A unified benchmark suite for safe learning-based control and reinforcement learning in robotics," *IEEE Robot. Automat. Lett.*, vol. 7, no. 4, pp. 11142–11149, Oct. 2022.
- [9] F. P. Bejarano, L. Brunke, and A. P. Schoellig, "Multi-step model predictive safety filters: Reducing chattering by increasing the prediction horizon," in *Proc. 62nd IEEE Conf. Decis. Control*, 2023, pp. 4723–4730.
- [10] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *Proc. 18th IEEE Eur. Control Conf.*, 2019, pp. 3420–3431.
- [11] K. P. Wabersich and M. N. Zeilinger, "A predictive safety filter for learningbased control of constrained nonlinear dynamical systems," *Automatica*, vol. 129, 2021, Art. no. 109597.
- [12] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. Cambridge, MA, USA: MIT Press, 2018.
- [13] G. Dalal, K. Dvijotham, M. Vecerik, T. Hester, C. Paduraru, and Y. Tassa, "Safe exploration in continuous action spaces," 2018, arXiv:1801.08757.
- [14] T.-H. Pham, G. De Magistris, and R. Tachibana, "Optlayer practical constrained optimization for deep reinforcement learning in the real world," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2018, pp. 6236–6243.
- [15] E. Altman, Constrained Markov Decision Processes. London, U.K.: Routledge, 2021.
- [16] J. Achiam, D. Held, A. Tamar, and P. Abbeel, "Constrained policy optimization," in *Proc. 34th Int. Conf. Mach. Learn.*, 2017, pp. 22–31.
  [17] Z. Li, U. Kalabić, and T. Chu, "Safe reinforcement learning: Learning
- [17] Z. Li, U. Kalabić, and T. Chu, "Safe reinforcement learning: Learning with supervision using a constraint-admissible set," in *Proc. IEEE Amer. Control Conf.*, 2018, pp. 6390–6395.
- [18] 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, pp. 3387–3395.

- [19] N. Kochdumper, H. Krasowski, X. Wang, S. Bak, and M. Althoff, "Provably safe reinforcement learning via action projection using reachability analysis and polynomial zonotopes," *IEEE Open J. Control Syst.*, vol. 2, pp. 79–92, 2023.
- [20] W. Zhao, Y. Sun, F. Li, R. Chen, T. Wei, and C. Liu, "Learn with imagination: Safe set guided state-wise constrained policy optimization," 2023, arXiv:2308.13140.
- [21] Y. Emam, P. Glotfelter, Z. Kira, G. Notomista, and M. Egerstedt, "Safe reinforcement learning using robust control barrier functions," *IEEE Robot. Automat. Lett.*, 2022.
- [22] X. Wang, "Ensuring safety of learning-based motion planners using control barrier functions," *IEEE Robot. Automat. Lett.*, vol. 7, no. 2, pp. 4773–4780, Apr. 2022.
- [23] A. Norouzi, S. Shahpouri, D. Gordon, M. Shahbakhti, and C. R. Koch, "Safe deep reinforcement learning in diesel engine emission control," J. Syst. Control Eng., vol. 237, no. 8, pp. 1440–1453, 2023.
- [24] K. Dunlap, M. Mote, K. Delsing, and K. L. Hobbs, "Run time assured reinforcement learning for safe satellite docking," *J. Aerosp. Inf. Syst.*, vol. 20, no. 1, pp. 298–313, 2023.
- [25] S. Pfrommer, T. Gautam, A. Zhou, and S. Sojoudi, "Safe reinforcement learning with chance-constrained model predictive control," in *Proc. Learn. Dyn. Control Conf.*, 2022, pp. 1–13.
- [26] W. Xiao et al., "BarrierNet: Differentiable control barrier functions for learning of safe robot control," *IEEE Trans. Robot.*, vol. 39, no. 3, pp. 2289–2307, Jun. 2023.
- [27] S. Yang, S. Chen, V. M. Preciado, and R. Mangharam, "Differentiable safe controller design through control barrier functions," *IEEE Contr. Syst. Lett.*, vol. 7, pp. 1207–1212, 2023.
- [28] A. Y. Ng, D. Harada, and S. Russell, "Policy invariance under reward transformations: Theory and application to reward shaping," in *Int. Conf. Mach. Learn.*, vol. 99, 1999, pp. 278–287.
- [29] J. Köhler, R. Soloperto, M. A. Müller, and F. Allgöwer, "A computationally efficient robust model predictive control framework for uncertain nonlinear systems–extended version," *IEEE Trans. Automat. Control*, vol. 66, no. 2, pp. 794–801, Feb. 2021.
- [30] J. B. Rawlings, D. Q. Mayne, and M. Diehl, *Model Predictive Control: Theory, Computation, and Design*, 2nd ed. Madison, WI, USA: Nob Hill Publishing, 2017.
- [31] X. Wang and M. Althoff, "Safe reinforcement learning for automated vehicles via online reachability analysis," *IEEE Trans. Intell. Veh.*, early access: Dec. 19, 2023, doi: 10.1109/TIV.2023.3344226.
- [32] F. Blanchini, "Set invariance in control," *Automatica*, vol. 35, no. 11, pp. 1747–1767, 1999.
- [33] R. Verschueren et al., "Acados: A modular open-source framework for fast embedded optimal control," *Math. Program. Computation*, vol. 14, no. 1, pp. 147–183, 2021.
- [34] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," 2017, arXiv:1707.06347.
- [35] C. E. Luis and J. L. Ny, "Design of a trajectory tracking controller for a nanoquadcopter," École Polytechnique de Montréal, Tech. Rep., 2016. [Online]. Available: https://publications.polymtl.ca/36013/
- [36] S. Teetaert et al., "A remote sim2real aerial competition: Fostering reproducibility and solutions' diversity in robotics challenges," 2023, arXiv:2308.16743.