Robust Predictive Output-Feedback Safety Filter for Uncertain Nonlinear Control Systems

Lukas Brunke, Siqi Zhou, and Angela P. Schoellig

Abstract— In real-world applications, we often require reliable decision making under dynamics uncertainties using noisy high-dimensional sensory data. Recently, we have seen an increasing number of learning-based control algorithms developed to address the challenge of decision making under dynamics uncertainties. These algorithms often make assumptions about the underlying unknown dynamics and, as a result, can provide safety guarantees. This is more challenging for other widely used learning-based decision making algorithms such as reinforcement learning. Furthermore, the majority of existing approaches assume access to state measurements, which can be restrictive in practice. In this paper, inspired by the literature on safety filters and robust output-feedback control, we present a robust predictive output-feedback safety filter (RPOF-SF) framework that provides safety certification to an arbitrary controller applied to an uncertain nonlinear control system. The proposed RPOF-SF combines a robustly stable observer that estimates the system state from noisy measurement data and a predictive safety filter that renders an arbitrary controller safe by (possibly) minimally modifying the controller input to guarantee safety. We show in theory that the proposed RPOF-SF guarantees constraint satisfaction despite disturbances applied to the system. We demonstrate the efficacy of the proposed RPOF-SF algorithm using an uncertain mass-spring-damper system.

I. INTRODUCTION

In many practical settings including autonomous driving and other robotics applications, only noisy and high-dimensional measurements are available and controllers must be designed that provide desired safety guarantees despite not having perfect state information. However, many control techniques assume access to the full state and an exact description of the system dynamics and observation model, which may not be available.

Recently, learning-based controllers have gained interest in such settings as they can synthesize control policies from high-dimensional measurements. Learning-based controllers such as reinforcement learning (RL) agents can improve performance by leveraging past experiences from the interaction with an environment, e.g., by operating in the real world. These experiences generally also include failures, which could result in damage to the robot or of its surrounding.

As learning-based controllers often do not account for safety constraints, add-on safety filters have been proposed to decouple performance (learning-based controller) and

The authors are with the Dynamic Systems Lab (http://www.dynsyslab.org), Institute for Aerospace Studies, University of Toronto, Canada. The authors are also affiliated with the University of Toronto Robotics Institute and the Vector Institute for Artificial Intelligence, Toronto. Angela P. Schoellig is also with the Department of Electrical and Computer Engineering, Technical University of Munich, Germany. Emails: {lukas.brunke, siqi.zhou, angela.schoellig}@robotics.utias.utoronto.ca
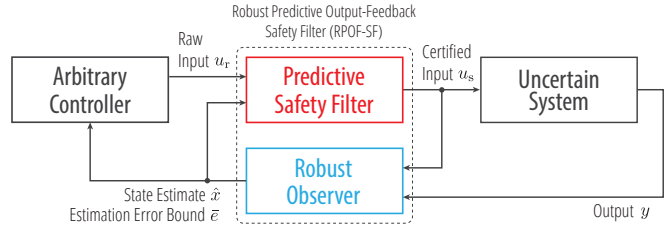
Fig. 1: A block diagram of the proposed robust predictive output-feedback safety filter (RPOF-SF). The proposed RPOF-SF consists of a robust observer (blue) and a predictive safety filter module (red). The robust observer provides an estimate of the system state from noisy output measurements, while the predictive safety filter modifies the input sent from an arbitrary controller with the goal to guarantee constraint satisfaction (i.e., to make it safe).

safety (safety filter). These safety filters certify the safety of control inputs from learning-based controllers and modify them if the safety filter determines the original control input would lead to safety violations. Current safety filter methods typically assume noisy state feedback. However, in practice, we must often rely on partial and noisy output measurements instead. If a safety filter does not account for measurement noise and/or state estimation errors present in the system, then the safety filter can generally not guarantee safety [1].

In this paper, building on the idea of robust output feedback model predictive control (MPC) [2] and model predictive safety certification (MPSC) [3], we derive a robust predictive output-feedback safety filter (RPOF-SF) that includes a predictive safety filter and a robustly stable observer (see Fig. 1). The proposed approach allows an uncertain system to operate safely under arbitrary control policies. We validate the proposed RPOF-SF on a numerical example[1].

II. RELATED LITERATURE

We summarize related literature on safety filters and robust output-feedback control in this section.

A. Safety Filters

Learning-based controllers such as reinforcement learning have been proposed to address the challenge of decision making under uncertainties. As compared to traditional model-based control techniques, learning-based approaches can be applied to a wider class of systems subject to

---

[1]The code is available at: https://github.com/utiasDSL/dsl__rpof__sf

model uncertainties [1]. Despite their flexibility, learning-based controllers often lack formal safety guarantees. In recent years, there have been several efforts from the control community targeted to address the problem of providing desired theoretical guarantees for learning-based controllers that are *not* initially designed to be safe. One stream of work in this area are safety filters. A safety filter minimally modifies any unsafe control inputs from the learning-based controller such that the system's state stays inside a safe set [1], which is either explicitly or implicitly defined.

Common techniques for explicitly defining safe sets include control barrier functions (CBF) [4]–[6] and Hamilton-Jacobi (HJ) reachability [7], [8], which are generally defined for continuous-time systems. Intuitively, the CBF framework provides a scalar condition for certifying control inputs to guarantee state constraint satisfaction, while the HJ reachability framework provides a means to compute a robust positive control invariant safe set, which is a set that is contained in a given state constraint set such that if the uncertain system starts inside the set, there exists a control law to keep the system inside the set despite disturbances. The explicit safe set characterizations from the CBF framework and the HJ reachability analysis provide the basis for designing safety filters to render a learning-based control system safe. Recently, CBF and HJ reachability are combined with learning to reduce the conservatism in the safety filter design for uncertain systems [5], [6], [8], [9].

An alternative to CBF and HJ reachability frameworks for safety certification is through the use of predictive filters, which do not assume a pre-computed safe set but determine the safe set implicitly via a MPC framework. A common safety filter that implements this idea is MPSC, which solves a finite-horizon constrained optimization problem with a discrete-time predictive model to prevent a learning-based controller from violating constraints [3]. These predictive filters typically require state-feedback and do not handle partial state measurements. To the best of our knowledge, a predictive safety filter for output-feedback systems has not been proposed.

B. Robust Output-Feedback Control

Due to their wide applicability and popularity, we focus on MPC techniques in this section. At every time step, MPC solves a finite-horizon optimal control problem subject to constraints, applies the first optimal control input, and replans at the next time step [10]. As MPC relies on its predictive model for optimal performance, disturbances acting on the system can lead to loss of feasibility of the optimization or constraint violation at the next time step. Under the assumption of full-state measurements, recent non-linear robust tube-based MPC schemes guarantee constraint satisfaction and feasibility for uncertain systems [11]–[13]. These schemes rely on a pre-stabilizing controller to stay close to a nominal trajectory and constraint tightening, such that no disturbance can lead to constraint violations. For linear systems with partial and noisy state measurements, there exist robust output-feedback MPC methods [14]–[17].

These approaches combine a robust MPC with a Luenberger observer and apply additional constraint tightening based on bounds on the estimation error. Recent work in [2] extends the nonlinear robust tube-based MPC scheme to the output-feedback setting. This approach designs a robustly stable observer that predicts verifiable error bounds on the estimation error. Additionally accounting for these bounds in the pre-stabilizing controller and constraint tightening yields constraint satisfaction for all future time steps.

There also exist min-max robust MPC formulations that simultaneously optimize for the optimal state estimate and control input sequence [18], [19]. While [19] is only applicable to linear systems, [18] handles nonlinear systems. However, these techniques typically have a high computational demand. This puts min-max methods at a disadvantage compared to tube-based methods for systems with complex dynamics requiring high control rates.

In this paper, we combine model predictive safety certification [3] with a recent framework for robust output-feedback MPC [2] to guarantee safe closed-loop operation for arbitrary control policies, which includes potentially unsafe control policies of a reinforcement learning agent.

Notation: The non-negative real numbers are $\mathbb{R}_{\geq 0}$. The set of integers in the interval $[a, b] \subset \mathbb{R}$ is $\mathbb{I}_{[a,b]}$, and the set of integers in the interval $[a, \infty) \subset \mathbb{R}$ is $\mathbb{I}_{\geq a}$. The class $\mathcal{K}$ denotes continuous functions $\gamma : [0, a) \to \mathbb{R}_{\geq 0}$ with $a > 0$, $\gamma(0) = 0$, and γ strictly monotonically increasing. The class $\mathcal{K}_{\infty}$ denotes functions $\gamma \in \mathcal{K}$ with $a = \infty$ that satisfy $\lim_{r \to \infty} \gamma(r) = \infty$.

III. PROBLEM STATEMENT

This paper is concerned with safely controlling an uncertain nonlinear constrained system with multiple inputs and multiple outputs, where not all states are measured and there is measurement noise. We consider the following uncertain nonlinear discrete-time system

$$\begin{aligned} x_{k+1} &= f_w(x_k, u_k, w_k) = f(x_k, u_k) + Ew_k \\ y_k &= h_w(x_k, u_k, w_k) = h(x_k, u_k) + Fw_k \,, \end{aligned} \quad (1)$$

where $k \in \mathbb{I}_{\geq 0}$ is the the discrete time step, $x_k \in \mathbb{X} \subseteq \mathbb{R}^{n_x}$ is the state, $u_k \in \mathbb{U} \subseteq \mathbb{R}^m$ is the control input, $w_k \in \mathbb{W} \subset \mathbb{R}^q$ is the disturbance with $\mathbb{W}$ being a bounded disturbance set containing zero, and $y_k \in \mathbb{Y} \subseteq \mathbb{R}^{n_y}$ is the noisy output. As $\mathbb{W}$ is bounded, there exists a scalar bound on the disturbance $\bar{w} \geq 0$ such that $\|w_k\| \leq \bar{w}$ for all $k \in \mathbb{I}_{\geq 0}$. The state dynamics $f_w : \mathbb{X} \times \mathbb{U} \times \mathbb{W} \to \mathbb{R}^{n_x}$ and output equation $h_w : \mathbb{X} \times \mathbb{U} \times \mathbb{W} \to \mathbb{R}^{n_y}$ are continuous functions. The nominal dynamics and output equations are given by $f(x, u) = f_w(x, u, 0)$ and $h(x, u) = h_w(x, u, 0)$, respectively. The matrices $E \in \mathbb{R}^{n_x \times q}$ and $F \in \mathbb{R}^{n_y \times q}$ map disturbances to states and outputs, respectively.

The goal of this work is to design a robust predictive output-feedback safety filter that guarantees safe operation of the system under arbitrary and potentially unsafe control policies. Safe operation of the system is defined by staying inside user-defined state and input constraints $(x_k, u_k) \in \mathbb{Z}$ for all time steps $k \in \mathbb{I}_{\geq 0}$. In this paper, we leverage

techniques from a recent robust output-feedback MPC framework [2] to develop a model predictive safety certification scheme for uncertain nonlinear systems without full-state feedback and noisy measurements.

IV. BACKGROUND

In this section, we introduce the necessary background for our proposed robust predictive output-feedback safety filter, which we will refer to as the RPOF-SF. First, we introduce a recent framework for robust output-feedback MPC [2]. This framework leverages incremental Lyapunov functions, which provide scalar propagation laws for over-approximations of the estimation and prediction errors. These propagation laws allow the design of robustly stable observers and can be efficiently evaluated in the MPC optimization. Second, we present a predictive safety filter inspired by MPC as in [3]. We leverage these ideas in our RPOF-SF to guarantee safe operation of dynamics systems with constraint satisfaction under bounded uncertainties in the dynamics and output.

A. Robust Output-Feedback MPC

We first introduce recent results on robust output-feedback MPC for nonlinear uncertain systems using online estimation error bounds [2].

Incremental Lyapunov Functions: Incremental Lyapunov functions provide a way of analyzing the stability between two trajectories from the same system with the same control input but different initial conditions. Intuitively, incremental stability between two trajectories means that both trajectories converge to the same trajectory. The notion of incremental stability can be extended to *(i)* incremental input/output-to-state stability (i-IOSS) by considering systems with inputs and outputs or *(ii)* incremental input-to-state stabilizability by considering a system with an additional feedback policy [2]. We begin with the definitions of detectability and stabilizability in terms of incremental Lyapunov functions.

Detectability for nonlinear systems can be described by an i-IOSS Lyapunov function [20]. As we will show below, the i-IOSS Lyapunov function provides us with the tool to synthesize robustly stable observers with guaranteed estimation error bounds.

Definition 1 (i-IOSS Lyapunov function [20]): A function $V_{\mathrm{d}} : \mathbb{X} \times \mathbb{X} \to \mathbb{R}_{\geq 0}$ (with subscript d for detectability) is called an (exponential-decrease) i-IOSS Lyapunov function if there exist lower and upper bounds on V_{d} with $\alpha_{\mathrm{d,l}}, \alpha_{\mathrm{d,u}} \in \mathcal{K}_{\infty}$, bounding functions $\sigma_{\mathrm{d},w}, \sigma_{\mathrm{d},y} \in \mathcal{K}$, and a decay rate $\rho_{\mathrm{d}} \in (0,1)$ such that

$$\alpha_{\mathrm{d,l}}(\|x - \tilde{x}\|) \leq V_{\mathrm{d}}(x, \tilde{x}) \leq \alpha_{\mathrm{d,u}}(\|x - \tilde{x}\|) \quad (2\mathrm{a})$$

$$\begin{aligned} V_{\mathrm{d}}(x^+, \tilde{x}^+) \leq \rho_{\mathrm{d}} V_{\mathrm{d}}(x, \tilde{x}) + \sigma_{\mathrm{d},w}(\|w - \tilde{w}\|) + \\ \sigma_{\mathrm{d},y}(\|y - \tilde{y}\|), \end{aligned} \quad (2\mathrm{b})$$

where $x^+{=}f_w(x, u, w)$, $\tilde{x}^+{=}f_w(\tilde{x}, u, \tilde{w})$, $y{=}h_w(x, u, w)$, $\tilde{y}{=}h_w(\tilde{x}, u, \tilde{w})$ for all $x, \tilde{x} \in \mathbb{X}$, $u \in \mathbb{U}$, and $w, \tilde{w} \in \mathbb{W}$.

Intuitively, such an i-IOSS Lyapunov function defines the convergence of two trajectories x and $\tilde{x}$ under the same control inputs u and the respective disturbances w and $\tilde{w}$. In Eqn. 2b, the decay rate ρ_{d} defines how quickly the two

trajectories converge, while $\sigma_{\mathrm{d},w}$ and $\sigma_{\mathrm{d},y}$ bound the effect of the disturbances and output, respectively. Given this notion of stability, we can then bound a trajectory $\{x_i\}_{i=0}^H$ of length $H \in \mathbb{N}$ using the knowledge of another trajectory $\{\tilde{x}_i\}_{i=0}^H$.

Incremental input-to-state stability (i-ISS) studies the stability of trajectories in the presence of disturbances w [13]. The notion of incremental input-to-state stability (i-ISS) can be extended to incremental input-to-state stabilizability by considering a control Lyapunov function (CLF). As we will see next, we can conveniently use this condition in a robust MPC framework to determine constraint tightening and thereby provide constraint satisfaction guarantees.

Definition 2 (i-ISS CLF [2]): A function $V_{\mathrm{s}} : \mathbb{X} \times \mathbb{X} \to \mathbb{R}_{\geq 0}$ (with subscript s for stabilizability) is called an (exponential-decrease) i-ISS CLF if there exist lower and upper bounds on V_{s} with $\alpha_{\mathrm{s,l}}, \alpha_{\mathrm{s,u}} \in \mathcal{K}_{\infty}$, bounding functions $\sigma_{\mathrm{s},w}, \sigma_{\pi} \in \mathcal{K}$, decay rate $\rho_{\mathrm{s}} \in (0,1)$, and a control law $\pi : \mathbb{X} \times \mathbb{X} \times \mathbb{U} \to \mathbb{U}$ such that

$$\alpha_{\mathrm{s,l}}(\|x - \tilde{x}\|) \leq V_{\mathrm{s}}(x, \tilde{x}) \leq \alpha_{\mathrm{s,u}}(\|x - \tilde{x}\|) \quad (3\mathrm{a})$$

$$V_{\mathrm{s}}(x^+, \tilde{x}^+) \leq \rho_{\mathrm{s}} V_{\mathrm{s}}(x, \tilde{x}) + \sigma_{\mathrm{s},w}(\|w - \tilde{w}\|) \quad (3\mathrm{b})$$

$$\|\pi(\tilde{x}, x, u) - u\| \leq \sigma_{\pi}(\|x - \tilde{x}\|), \quad (3\mathrm{c})$$

where $x^+{=}f_w(x, u, w)$, $\tilde{x}^+{=}f_w(\tilde{x}, \pi(\tilde{x}, x, u), \tilde{w})$ for all $x, \tilde{x} \in \mathbb{X}$, $u \in \mathbb{U}$, and $w, \tilde{w} \in \mathbb{W}$.

The i-ISS CLF describes the convergence of two trajectories x and $\tilde{x}$ under respective disturbances w and $\tilde{w}$ achieved by the additional feedback π. Similar to the notion of i-ISS, the decay rate ρ_{s} in Eqn. 3b defines how quickly the two trajectories converge, while $\sigma_{\mathrm{s},w}$ bounds the effect of the disturbance. The condition in Eqn. 3b allows us to efficiently propagate the set of states that are reachable at future time steps.

We note that there exist various methods in the literature for the synthesis of such incremental Lyapunov functions, see [21]–[23]. By making a few simplifying assumptions (e.g., continuously differentiable dynamics, quadratic incremental Lyapunov function, linear $\mathcal{K}$ functions), we can rewrite the conditions for the i-IOSS Lyapunov function and i-ISS CLF as linear matrix inequalities (LMIs). Furthermore, as discussed in [23], by using a convexifying or a gridding approach, we can turn the synthesis problem into a finite semi-definite program (SDP), which can be solved to achieve desired decay rates ρ_{d} and ρ_{s}. The existence of the incremental Lyapunov functions is checked prior to closed-loop execution, and the constants and the class $\mathcal{K}$ functions are pre-computed based on either Eqn. 2 or Eqn. 3. As we will see in Eqn. 12, in a robust output-feedback MPC formulation, we can replace the incremental Lyapunov conditions by a set of scalar conditions using the decay rates $\rho_{\mathrm{d}}, \rho_{\mathrm{s}}$ and the class $\mathcal{K}$ functions. These scalar conditions only minimally increase the computational demand compared to a nominal MPC.

State Estimation Error Bounds: One key ingredient in a robust output-feedback MPC framework is a robustly stable observer. We next introduce Luenberger-like observers and moving horizon estimators (MHE) with valid online estimation error bounds. We consider the Luenberger-like

observer as a back-up observer if the robustness conditions for the MHE are not satisfied.

Inspired by Luenberger observers for linear systems, a nonlinear observer is given by

$$\hat{x}_{k+1} = f(\hat{x}_k, u_k) + \hat{L}(\hat{x}_k, u_k, y_k) = \hat{f}(\hat{x}_k, u_k, y_k), \quad (4)$$

with the estimated state $\hat{x}_k \in \mathbb{R}^{n_x}$ and the estimator correction $\hat{L} : \mathbb{X} \times \mathbb{U} \times \mathbb{Y} \to \mathbb{X}$ that satisfies $\hat{L}(\hat{x}, u, h(\hat{x}, u)) = 0$.

Assumption 1 (Robustly stable observer [2]): There exist an incremental Lyapunov function $V_\mathrm{o} : \mathbb{X} \times \mathbb{X} \to \mathbb{R}_{\geq 0}$ (with subscript o for observer) with lower and upper bounds $\alpha_{\mathrm{o,l}}, \alpha_{\mathrm{o,u}} \in \mathcal{K}_\infty$, bounding functions $\sigma_{\mathrm{o,}w}, \sigma_{\mathrm{o,L}}, \sigma_{\mathrm{o,L},w} \in \mathcal{K}$, and decay rate $\rho_\mathrm{o} \in (0, 1)$, such that

$$\alpha_{\mathrm{o,l}}(\|x - \hat{x}\|) \leq V_\mathrm{o}(x, \hat{x}) \leq \alpha_{\mathrm{o,u}}(\|x - \hat{x}\|) \quad (5a)$$

$$V_\mathrm{o}(\hat{f}(\hat{x}, u, y), f(x, u, w)) \leq \rho_\mathrm{o} V_\mathrm{o}(x, \hat{x}) + \sigma_{\mathrm{o,}w}(\|w\|) \quad (5b)$$

$$\|\hat{L}(\hat{x}, u, y)\| \leq \sigma_{\mathrm{o,L}}(V_\mathrm{o}(x, \hat{x})) + \sigma_{\mathrm{o,L},w}(\|w\|), \quad (5c)$$

where $y = h_w(x, u, w)$ for all $x, \hat{x} \in \mathbb{X}$, $u \in \mathbb{U}$, and $w \in \mathbb{W}$.

To simplify the following discussion, we assume that the incremental Lyapunov function is identical to the i-IOSS Lyapunov function with $V_\mathrm{o} = V_\mathrm{d}$. Sufficient conditions for V_o being an i-IOSS Lyapunov function can be found in [2] and require affine disturbances in the state dynamics and output equation (as assumed in Eqn. 1) and a quadratic V_o.

Given an upper bound $\bar{e}_0$ on the initial estimation error e_0 such that $0 \leq e_0 \leq \bar{e}_0$ and the robust stable observer described by the i-IOSS Lyapunov function V_o, the error bound can be recursively updated either offline or online

$$\bar{e}_{k+1,\mathrm{offline}} = \rho_\mathrm{o} \bar{e}_k + \sigma_{\mathrm{o,}w}(\|\bar{w}\|), \quad (6)$$
$$\bar{e}_{k+1,\mathrm{online}} = \rho_\mathrm{d} \bar{e}_k + \sigma_{\mathrm{d,}w}(\bar{w} + \|\hat{w}_k\|) + \sigma_{\mathrm{d,}y}(\|\hat{y}_k - y_k\|),$$

where $\hat{w}_k = E^\mathsf{T} (EE^\mathsf{T})^{-1} \hat{L}(\hat{x}_k, u_k, y_k)$ with E defined in Eqn. 1 and $\hat{y}_k = h(\hat{x}_k, u_k)$. The online update incorporates the latest output measurement y_k, while the offline update requires no additional data and relies entirely on the pre-computed decay rate ρ_o, $\sigma_{\mathrm{o,}w}$, and $\bar{w}$. See [2] for a proof of the recursive updates in Eqn. 6.

The (exponential-decay) i-IOSS Lyapunov function V_o with $V_\mathrm{o}(\hat{x}_0, x_0) \leq \bar{e}_0$ guarantees that for all $k \in \mathbb{I}_{\geq 0}$, $\hat{x}_k, \bar{e}_k$ from Eqn. 4 and Eqn. 6 satisfy

$$V_\mathrm{o}(x_k, \hat{x}_k) \leq \bar{e}_k \quad (7a)$$

$$\|\hat{L}(\hat{x}_k, u_k, y_k)\| \leq \sigma_{\mathrm{o,L}}(\bar{e}_k) + \sigma_{\mathrm{o,L},w}(\|\bar{w}\|), \quad (7b)$$

$$\bar{e}_{k+i} \leq \rho_\mathrm{o}^i \bar{e}_k + \frac{1 - \rho_\mathrm{o}^i}{1 - \rho_\mathrm{o}} \sigma_{\mathrm{o,}w}(\|\bar{w}\|), \forall i \in \mathbb{I}_{\geq 0}. \quad (7c)$$

See [2] for the proof. The conditions in Eqn. 7 provide desirable properties for the state estimation since they yield valid bounds on the estimation error and their prediction.

In the following, we introduce the MHE that determines the state estimate over a backward finite-time horizon $M > 0$. The MHE is considered as the dual of MPC for state estimation and solves a finite-horizon optimization to determine an optimized state estimate using past input and output data.

Assumption 2 (Continuous i-IOSS Lyapunov function): There exists a function $\sigma_\mathrm{d} \in \mathcal{K}$, such that for any $x, \hat{x}, \tilde{x} \in \mathbb{X}$, the i-IOSS Lyapunov function V_d satisfies

$$|V_\mathrm{d}(\hat{x}, x) - V_\mathrm{d}(\tilde{x}, x)| \leq \sigma_\mathrm{d}(\|\hat{x} - \tilde{x}\|). \quad (8)$$

In the following, we use the notation $\hat{x}_{i|k}$ to denote the open-loop prediction from state $\hat{x}_k$ at time step $k + i$. At time step k, the MHE with the backward horizon $M_k = \min\{k, M\}$ and the past initial state estimate $\hat{x}_{k-M_k}$ is given by the following nonlinear program [2]:

$$J^*_{M_k,\mathrm{MHE}}(\hat{x}_{k-M_k}) =$$

$$\min_{\hat{w}_{-M_k:-1|k}, \hat{x}_{-M_k|k}} \sum_{j=1}^{M_k} \rho_\mathrm{d}^{j-1}(\sigma_{\mathrm{d,}w}(\bar{w} + \|\hat{w}_{-j|k}\|) +$$

$$\sigma_{\mathrm{d,}y}(\|\hat{y}_{-j|k} - y_{k-j}\|)) +$$

$$\rho_\mathrm{d}^{M_k} \bar{e}_{k-M_k} + \quad (9)$$

$$\rho_\mathrm{d}^{M_k} \sigma_\mathrm{d}(\|\hat{x}_{-M_k|k} - \hat{x}_{k-M_k}\|)$$

$$\text{s.t.} \quad \hat{x}_{-j+1|k} = f_w(\hat{x}_{-j|k}, u_{k-j}, \hat{w}_{-j|k}),$$

$$\hat{y}_{-j|k} = h_w(\hat{x}_{-j|k}, u_{k-j}, \hat{w}_{-j|k}),$$

$$\forall j \in \mathbb{I}_{[1,M_k]}.$$

The resulting state estimate and estimation error bound can be determined by propagating the minimizers of Eqn. 9 $\hat{w}^*_{-M_k:-1|k}, \hat{x}^*_{-M_k|k}$ until $\hat{x}^*_{0|k}$ is reached:

$$\hat{x}_{k,\mathrm{MHE}} = \hat{x}^*_{0|k},$$
$$\hat{e}_{k,\mathrm{MHE}} = J^*_{M_k,\mathrm{MHE}}(\hat{x}_{k-M_K}). \quad (10)$$

As the MHE requires additional assumptions to be robustly stable, we can instead check conditions Eqn. 7b and Eqn. 7c on the estimation error for the MHE update to guarantee robustness of the MHE update. In the following, we assume that the resulting estimation error bound from the MHE is valid according to Eqn. 7b and Eqn. 7c. In practice, these conditions must be checked at every step. If the conditions are not met, the backup Luenberger-like observer is used.

Prediction Error Bounds: In this part, we highlight the use of incremental Lyapunov functions for tube-based predictions of the combined error bounds originating from the estimation error and the additive disturbance on the dynamics. This requires the analysis of the incremental stability between the true state x, the state estimate $\hat{x}$, and a nominal state $\bar{x}$.

Under the assumption that V_s is an i-ISS CLF, there also exist functions $\sigma_{\mathrm{s,o}}, \sigma_{\mathrm{s,o},w} \in \mathcal{K}$ such that:

$$V_\mathrm{s}(\bar{x}^+, \hat{x}^+) \leq \rho_\mathrm{s} V_\mathrm{s}(\bar{x}, \hat{x}) + \sigma_{\mathrm{s,o}}(V_\mathrm{o}(\hat{x}, x)) + \sigma_{\mathrm{s,o},w}(\bar{w}), \quad (11)$$

where $u = \pi(\hat{x}, \bar{x}, \bar{u})$, $y = h_w(x, u, w)$, $\bar{x}^+ = f(\bar{x}, \bar{u})$, and $\hat{x}^+ = \hat{f}(\hat{x}, u, y)$ (see [2] for a proof of this implication). The conditions on V_s can be used to predict the combined errors from the estimation and the disturbed state dynamics.

Tube-based Robust Output-feedback MPC: The robust output-feedback MPC minimizes a user-defined continuous stage cost $\ell : \mathbb{X} \times \mathbb{U} \times \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \to \mathbb{R}$, which allows for additional arguments to also include the bounds on the estimation and prediction errors. Robust constraint satisfaction

is guaranteed via the proper design of a continuous terminal cost $V_\mathrm{f} : \mathbb{X} \times \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \to \mathbb{R}$ and a terminal constraint set $\mathbb{X}_\mathrm{f} \subseteq \mathbb{X} \times \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}$. At time step k, for a given state estimate from the Luenberger-like observer Eqn. 6 and Eqn. 4 or the MHE Eqn. 10, the robust output-feedback MPC open-loop optimization problem over a finite horizon $N \in \mathbb{I}_{\geq 1}$ is:

$$J^*_{N,\mathrm{MPC}}(\hat{x}_k, \bar{e}_k) = \tag{12a}$$

$$\min_{\bar{u}_{0:N-1|k}, \bar{x}_{0|k}} \sum_{i=0}^{N-1} l(\bar{x}_{i|k}, \bar{u}_{i|k}, \bar{e}_{i|k}, \bar{s}_{i|k}) + \tag{12b}$$

$$V_\mathrm{f}(\bar{x}_{N|k}, \bar{e}_{N|k}, \bar{s}_{N|k}) \tag{12c}$$

$$\text{s.t.} \quad \bar{s}_{0|k} = V_\mathrm{s}(\bar{x}_{0|k}, \hat{x}_k) , \tag{12d}$$

$$\bar{x}_{i+1|k} = f(\bar{x}_{i|k}, \bar{u}_{i|k}), \forall i \in \mathbb{I}_{[0,N-1]} , \tag{12e}$$

$$\bar{e}_{j|k} = \frac{1 - \rho_\mathrm{d}^j}{1 - \rho_\mathrm{d}} \sigma_{\mathrm{o},w}(\bar{w}) + \rho_\mathrm{d}^j \bar{e}_k , \forall j \in \mathbb{I}_{[0,N]} , \tag{12f}$$

$$\bar{s}_{i+1|k} = \rho_\mathrm{s} \bar{s}_{i|k} + \sigma_{\mathrm{s},\mathrm{o}}(\bar{e}_{i|k}) + \sigma_{\mathrm{s},\mathrm{o},w}(\bar{w}) , \tag{12g}$$

$$(x_{i|k}, \pi(\hat{x}_{i|k}, \bar{x}_{i|k}, \bar{u}_{i|k})) \in \mathbb{Z} , \tag{12h}$$

$$V_\mathrm{s}(\bar{x}_{i|k}, \hat{x}_{i|k}) \leq \bar{s}_{i|k} , \forall x_{i|k}, \hat{x}_{i|k} , \tag{12i}$$

$$V_\mathrm{o}(\hat{x}_{i|k}, x_{i|k}) \leq \bar{e}_{i|k} , \forall x_{i|k}, \hat{x}_{i|k} , \tag{12j}$$

$$(\bar{x}_{N|k}, \bar{e}_{N|k}, \bar{s}_{N|k}) \in \mathbb{X}_\mathrm{f} . \tag{12k}$$

Eqn. 12c guarantees that the initial nominal state is inside the set around the state estimate given by V_s. The constraints in Eqn. 12e, Eqn. 12f, and Eqn. 12g specify the nominal dynamics, estimation error bound and prediction error bound propagation, respectively. Eqn. 12h guarantees constraint satisfaction for the true state and control input. Proper containment inside the tubes given by V_s and V_o for the nominal and true state is achieved by Eqn. 12i and Eqn. 12j. Finally, the terminal constraint is satisfied with Eqn. 12k.

The minimizers of Eqn. 12 are $\bar{u}^*_{0:N-1|k}, \bar{x}^*_{0|k}$. At each time step, only the first optimal control input $\bar{u}^*_{0|k}$ is applied to the system through $u_k = \pi(\hat{x}_k, \bar{x}^*_{0|k}, \bar{u}^*_{0|k})$. Then the open-loop optimization problem is solved again at the next time step using an updated state estimate $\hat{x}_{k+1}$ and error bound $\bar{e}_{k+1}$ over a shifted time horizon.

Properly designed terminal ingredients provide constraint satisfaction and recursive feasibility of the robust output-feedback MPC.

Assumption 3 (Terminal constraint set [2]): There exists a control law $\pi_\mathrm{f} : \mathbb{X} \to \mathbb{U}$ such that for all $(\bar{x}, \bar{e}, \bar{s}) \in \mathbb{X}_\mathrm{f}$ and $x, \hat{x} \in \mathbb{X}$ satisfying $V_\mathrm{s}(\bar{x}, \hat{x}) \leq \bar{s}$ and $V_\mathrm{o}(\hat{x}, x) \leq \bar{e}$, and for all $\bar{s}^+, \bar{e}^+ \in \mathbb{R}_{\geq 0}$ satisfying $\bar{s}^+ \leq \rho_\mathrm{s} \bar{s} + \sigma_{\mathrm{s},\mathrm{o}}(\bar{e}) + \sigma_{\mathrm{s},\mathrm{o},w}(\bar{w})$ and $\bar{e}^+ \leq \rho_\mathrm{o} \bar{e} + \sigma_{\mathrm{o},w}(\bar{w})$, it holds that

$$(\bar{x}^+, \bar{e}^+, \bar{s}^+) \in \mathbb{X}_\mathrm{f} , \quad (x, \pi(\hat{x}, \bar{x}, \bar{u})) \in \mathbb{Z} , \tag{13a}$$

where $\bar{x}^+ = f(\bar{x}, \bar{u})$ and $\bar{u} = \pi_\mathrm{f}(\bar{x})$.

Recursive feasibility and boundedness of the closed-loop cost can be shown using a terminal constraint set satisfying Assumption 3 and some additional assumptions on the stage cost ℓ and the terminal cost function V_f [2].

Under the assumption of continuous constraint functions and the initial nominal state being equal to the initial state estimate, $\bar{x}^*_{0|k} = \hat{x}_k$ (this implies $\bar{s}^*_{0|k} = 0$), the constraint tightening can be pre-computed. The resulting computational demand is comparable to a nominal MPC [2].

We use the robust output-feedback MPC framework as the basis for our proposed safety filter to provide safety guarantees for constrained nonlinear systems with bounded estimation errors.

B. Model Predictive Safety Certification

The model predictive safety certification (MPSC) method is inspired by model predictive control techniques. Instead of providing an optimal control input, the MPSC takes an arbitrary control input and certifies the provided input either as safe or unsafe [3]. If the control input is considered as safe, it passes through the MPSC without modification. Otherwise, the control input is deemed unsafe and is minimally modified (according to a distance measure) such that the control input is safe. This relies on techniques from MPC to guarantee constraint satisfaction for all time steps. Guarantees are provided through the existence of a safe set $\mathbb{S}_\mathrm{f}$:

Assumption 4 (Safe set): There exists a control policy $\pi_\mathrm{safe} : \mathbb{X} \to \mathbb{U}$ such that for all $(\bar{x}_k, \bar{e}_k, \bar{s}_k) \in \mathbb{S}_\mathrm{f}$ and for all $j \in \mathbb{I}_{\geq 0}$ and all $x_{k+j}, \hat{x}_{k+j} \in \mathbb{X}$ satisfying $V_\mathrm{s}(\bar{x}_{k+j}, \hat{x}_{k+j}) \leq \bar{s}_{k+j}$ and $V_\mathrm{o}(\hat{x}_{k+j}, x_{k+j}) \leq \bar{e}_{k+j}$, it holds that

$$(x_{k+j}, \pi(\hat{x}_{k+j}, \bar{x}_{k+j}, \bar{u}_{k+j})) \in \mathbb{Z} \tag{14}$$

with $\bar{x}_{k+j+1} = f(\bar{x}_{k+j}, \bar{u}_{k+j})$, $\bar{u}_{k+j} = \pi_\mathrm{safe}(\bar{x}_{k+j})$ and with any $\bar{s}_{k+j+1}, \bar{e}_{k+j+1} \in \mathbb{R}_{\geq 0}$ satisfying $\bar{s}_{k+j+1} \leq \rho_\mathrm{s} \bar{s}_{k+j} + \sigma_{\mathrm{s},\mathrm{o}}(\bar{e}_{k+j}) + \sigma_{\mathrm{s},\mathrm{o},w}(\bar{w})$ and $\bar{e}_{k+j+1} \leq \rho_\mathrm{o} \bar{e}_{k+j} + \sigma_{\mathrm{o},w}(\bar{w})$.

The assumption of the existence of safe set $\mathbb{S}_\mathrm{f}$ is less restrictive than the assumption of a robustly positive invariant terminal set $\mathbb{X}_\mathrm{f}$ as in Assumption 3. However, guaranteeing constraint satisfaction for all future time steps requires additional attention (cf. proof for Theorem 1).

We present the idea of MPSC using nominal dynamics and full-state measurements, such that $x_k = \bar{x}_k$ and $u_k = \pi(\hat{x}_k, \bar{x}_k, \bar{u}_k) = \pi(\bar{x}_k, \bar{x}_k, \bar{u}_k) = \bar{u}_k$. The safety filter framework can be extended to uncertain systems using a formulation as in robust MPC. The main difference between the MPSC and a nominal MPC is the objective function. For the MPSC, the objective is the squared error between the first optimal control input $u^*_{0|k}$ and the provided control input at the current state $\tilde{u} = \pi_\mathcal{L}(x_k)$, where $\pi_\mathcal{L} : \mathbb{X} \times \mathbb{U}$ is an arbitrary and potentially unsafe control policy. Then the open-loop optimization problem for the MPSC at time step k under the assumption of zero disturbances and perfect state measurements is given by:

$$J^*_{N,\mathrm{MPSC}}(x_k) = \min_{u_{0:N-1|k}} \|\pi_\mathcal{L}(x_k) - u_{0|k}\|_2^2$$

$$\text{s.t.} \quad \forall i \in \mathbb{I}_{[0,N-1]} ,$$
$$x_{i+1|k} = f(x_{i|k}, u_{i|k}) , \tag{15}$$
$$(x_{i|k}, u_{i|k}) \in \mathbb{Z} ,$$
$$x_{0|k} = x_k ,$$
$$x_{N|k} \in \mathbb{S}_\mathrm{f} .$$

The certified control inputs that can be safely applied to the system are given by the minimizer $u_k = u^*_{0|k}$ in an MPC fashion. Proofs for constraint satisfaction for all future time steps using an adaptive horizon approach and guarantees for handling uncertain state dynamics can be found in [3].

Previous results assume full-state measurements without noise. In this work, we extend the model predictive safety certification optimization problem to handle uncertain output measurements using state estimates with valid error bounds.

V. ROBUST PREDICTIVE OUTPUT-FEEDBACK SAFETY FILTER (RPOF-SF)

In this section, we present our proposed RPOF-SF (see Fig. 1). The RPOF-SF takes an input from arbitrary state-feedback control policies and either certifies the input as safe or modifies the desired control input minimally (according to the ℓ_2-norm) to still guarantee safety for uncertain constrained nonlinear systems.

A. Overview of the RPOF-SF Algorithm

As we are dealing with noisy and partial state measurements, we leverage the state estimation techniques from the previous section to determine state estimates with valid error bounds. To simplify the formulation, we assume that the MHE is a robustly stable observer. In practice, this requires checking the conditions Eqn. 7b and Eqn. 7c. The estimation error bound at time step $k \in \mathbb{I}_{\geq 1}$ is computed as

$$\bar{e}_k = \min\{\bar{e}_{k,\text{offline}}, \bar{e}_{k,\text{online}}, \bar{e}_{k,\text{MHE}}\}, \qquad (16)$$

to determine the smallest upper bound and the state estimate is then either given by the associated MHE or the Luenberger-like observer with

$$\hat{x}_k = \begin{cases} \hat{x}_{k,\text{MHE}} & \text{if } \bar{e}_k = \bar{e}_{k,\text{MHE}}, \\ \hat{f}(\hat{x}_{k-1}, u_{k-1}, y_{k-1}) & \text{otherwise}. \end{cases} \qquad (17)$$

The pair of state estimate $\hat{x}_k$ and estimation error bound $\bar{e}_k$ are then used to initialize our predictive safety filter.

By combining the robust output-feedback MPC and the MPSC, we propose the following open-loop optimization problem for the robust output-feedback predictive safety filter:

$$J^*_{N,\text{RPOF}-\text{SF}}(\hat{x}_k, \bar{e}_k) =$$
$$\min_{\bar{u}_{0:N-1|k}, \bar{x}_{0|k}} \|\pi_{\mathcal{L}}(\hat{x}_k) - \pi(\hat{x}_k, \bar{x}_{0|k}, \bar{u}_{0|k})\|^2_2 \qquad (18a)$$

$$\text{s.t.} \quad \bar{s}_{0|k} = V_{\text{s}}(\bar{x}_{0|k}, \hat{x}_k), \qquad (18b)$$

$$\bar{x}_{i+1|k} = f(\bar{x}_{i|k}, \bar{u}_{i|k}), \forall i \in \mathbb{I}_{[0,N-1]}, \qquad (18c)$$

$$\bar{e}_{j|k} = \frac{1-\rho_{\text{d}}^j}{1-\rho_{\text{d}}}\sigma_{\text{o},w}(\bar{w}) + \rho_{\text{d}}^j \bar{e}_k, \forall j \in \mathbb{I}_{[0,N]}, \qquad (18d)$$

$$\bar{s}_{i+1|k} = \rho_{\text{s}}\bar{s}_{i|k} + \sigma_{\text{s},\text{o}}(\bar{e}_{i|k}) + \sigma_{\text{s},\text{o},w}(\bar{w}), \qquad (18e)$$

$$(x_{i|k}, \pi(\hat{x}_{i|k}, \bar{x}_{i|k}, \bar{u}_{i|k})) \in \mathbb{Z}, \qquad (18f)$$

$$V_{\text{s}}(\bar{x}_{i|k}, \hat{x}_{i|k}) \leq \bar{s}_{i|k}, \forall x_{i|k}, \hat{x}_{i|k}, \qquad (18g)$$

$$V_{\text{o}}(\hat{x}_{i|k}, x_{i|k}) \leq \bar{e}_{i|k}, \forall x_{i|k}, \hat{x}_{i|k}, \qquad (18h)$$

$$(\bar{x}_{N|k}, \bar{e}_{N|k}, \bar{s}_{N|k}) \in \mathbb{S}_{\text{f}}. \qquad (18i)$$

In the proposed RPOF-SF we use the objective function from Eqn. 15 in Eqn. 18a to certify arbitrary control policies and substitute the terminal set $\mathbb{X}_{\text{f}}$ for the more general safe set $\mathbb{S}_{\text{f}}$ in Eqn. 18i.

The minimizers of Eqn. 18 are denoted by $\bar{u}^*_{0:N-1|k}, \bar{x}*_{0|k}$. We apply $u_k = \pi(\hat{x}_k, \bar{x}^*_{0|k}, \bar{u}^*_{0|k})$ to the system and solve the open-loop optimization again at the next time step with a receding horizon. If the problem is infeasible at the next time step, which is possible due to the more general assumption on the safe set $\mathbb{S}_{\text{f}}$, the procedure in Alg. 1 must be followed to guarantee constraint satisfaction for all future time steps. If $\pi_{\mathcal{L}}(\hat{x}_k)$ is a feasible control input then $u_k = \pi_{\mathcal{L}}(\hat{x}_k)$. Otherwise u_k is the minimal modification of $\pi_{\mathcal{L}}(\hat{x}_k)$ such that the optimization in Eqn. 18 is feasible.

The proposed algorithm is given in Alg. 1, where the additional subscript $\tilde{N}$ in lines 9 and 10 of Alg. 1 indicates that this nominal state or control input is the result of an optimization with reduced horizon $\tilde{N} < N$.

The formulation can be further simplified as in the robust output-feedback MPC with pre-computed constraint tightening and tube size 0 at the initial nominal state [24].

B. RPOF-SF Constraint Satisfaction Guarantees

By leveraging the results in [2], [3], we show that the proposed RPOF-SF guarantees constraint satisfaction for all time steps.

Theorem 1: Let Assumptions 1 and 4 hold. Suppose that the system in Eqn. 1 admits an (exponential-decay) i-IOSS Lyapunov function (Def. 1) and an (exponential decrease) i-ISS CLF (Def. 2). If the optimization in Eqn. 18 is feasible at $k = 0$ with a valid initial estimation error bound $V_{\text{o}}(\hat{x}_0, x_0) \leq \bar{e}_0$, then the RPOF-SF described in Alg. 1 guarantees constraint satisfaction for all $k \in \mathbb{I}_{\geq 0}$. Furthermore, if the safe set $\mathbb{S}_{\text{f}}$ also satisfies Assumption 3, then the optimization in Eqn. 18 is recursively feasible.

Proof: Suppose that Eqn. 18 is feasible at time step $k \in \mathbb{I}_{\geq 0}$. The sequence of nominal optimal control inputs is $\{\bar{u}^*_{0|k}, \ldots, \bar{u}^*_{N-1|k}\}$ and the initial optimal nominal state

Algorithm 1 Robust predictive output-feedback safety filter

1: **while** True **do**
2:    Update $\hat{x}_k$ and $\bar{e}_k$ using Eqn. 16 and Eqn. 17, respectively.
3:    **if** 18 is feasible for horizon N **then**
4:       $k_{\text{feasible}} \leftarrow k$ ▷ Update feasible time step
5:       $u_k \leftarrow \pi(\hat{x}_k, \bar{x}^*_{0|k}, \bar{u}^*_{0|k})$
6:    **else**
7:       **if** $k < N + k_{\text{feasible}}$ **then**
8:          Solve 18 for horizon $\tilde{N} := N - (k - k_{\text{feasible}})$
9:          $u_k \leftarrow \pi(\hat{x}_k, \bar{x}^*_{0|k,\tilde{N}}, \bar{u}^*_{0|k,\tilde{N}})$
10:          $\bar{x}_k \leftarrow \bar{x}^*_{0|k,\tilde{N}}$
11:       **else**
12:          $u_k \leftarrow \pi(\hat{x}_k, \bar{x}_k, \pi_{\text{safe}}(\bar{x}_k))$
13:          $\bar{x}_k \leftarrow f(\bar{x}_k, \pi_{\text{safe}}(\bar{x}_k))$
14:       **end if**
15:    **end if**
16:    Apply the certified control input u_k
17:    $k \leftarrow k+1$
18: **end while**

is $\bar{x}^*_{0|k}$. The true state and input satisfy $(x_k, u_k) \in \mathbb{Z}$ with $u_k = \pi(\hat{x}_k, \bar{x}^*_{0|k}, \bar{u}^*_{0|k})$ due to the over-approximation of the error bounds $\bar{e}_k$ and $\bar{s}^*_{0|k}$ guaranteed by the i-IOSS Lyapunov function and the i-ISS CLF under Assumption 1. Then we can safely apply the control input u_k.

If the optimization in Eqn. 18 is feasible at time $(k+1)$, then constraint satisfaction for the true state $(x_{k+1}, u_{k+1}) \in \mathbb{Z}$ is again guaranteed by the over-approximation of the error bounds and we apply u_{k+1}.

If the optimization is infeasible at time $(k+1)$, we can construct a feasible solution using a previous feasible solution for a reduced horizon $(N-1)$. This is guaranteed to be feasible since we already know that there exists at least one feasible nominal control input sequence with $\{\bar{u}^*_{1|k}, \ldots, \bar{u}^*_{N-1|k}\}$ and the initial nominal state $\bar{x}_{0|k+1} = \bar{x}^*_{1|k}$. Feasibility of the optimization in Eqn. 18 with horizon $(N-1)$ again guarantees $(x_{k+1}, u_{k+1}) \in \mathbb{Z}$ with $u_{k+1} = \pi(\hat{x}_{k+1}, \bar{x}^*_{0|k+1,N-1}, \bar{u}^*_{0|k+1,N-1})$ and we apply u_{k+1}.

If the optimization with the full horizon N is consecutively infeasible for the next $N-1$ steps, then feasibility of the optimization problem with horizon of length 1 at the previous time step guarantees that $(\bar{x}_{1|k+N-1,1}, \bar{e}_{1|k+N-1,1}, \bar{s}_{1|k+N-1,1}) \in \mathbb{S}_{\text{safe}}$. Then by Assumption 4, all future states and inputs guarantee $(x_{k+N+i}, u_{k+N+i}) \in \mathbb{Z}$ for all $i \in \mathbb{I}_{\geq 0}$.

In case the safe set $\mathbb{S}_{\text{f}}$ satisfies Assumption 3, we recover recursive feasibility and constraint satisfaction and require no adaptation of the horizon, see [2]. ∎

The proof above shows that, despite the uncertainty in dynamics and output equations, the proposed RPOF-SF is still able to achieve constraint satisfaction at every time step. This result enables the certification of control inputs from arbitrary control policies.

VI. NUMERICAL EXAMPLE

In this section, we demonstrate that RPOF-SF achieves constraint satisfaction for uncertain nonlinear systems despite an arbitrary control policy. We apply the proposed RPOF-SF on the following simulated nonlinear mass-spring-damper system from [25] with an output measurement of the first element of the state, similar to [24]:

$$\dot{x}_1 = x_2 , \; \dot{x}_2 = \frac{1}{M}\big(- k_0 \exp(-x_1)\, x_1 - h_d x_2 + u \big),$$
$$y = x_1 ,$$

(19)

where $M = 1$, $k_0 = 0.33$, and $h_d = 1.1$. The discrete-time model is determined based on a fourth-order Runge-Kutta method with sampling time $\Delta t = 0.25$ s, which yields the discrete-time dynamics $f(x, u)$. We consider additive disturbances with $Ew = \begin{bmatrix} \Delta t \cdot w_1 & \frac{\Delta t}{M} w_2 \end{bmatrix}^{\mathsf{T}}$ and $Fw = w_3$, where $w \in \mathbb{W} = \{w \in \mathbb{R}^3 : \|w\|_\infty \leq 0.01\}$. We enforce the following constraints point-wise in time $\mathbb{Z} = [-0.85, 0.85] \times [-2, 2] \times [-6, 6]$.

We derive LMIs for the incremental Lyapunov functions assuming quadratic incremental Lyapunov functions, linear $\mathcal{K}$ and $\mathcal{K}_\infty$ functions, and linear feedback π, similar to [23]. Using a gridding approach, we can solve the resulting SDPs over the entire constraint set for the discrete-time system since we assume global properties. We run this computation offline in Matlab using YALMIP [26] and the solvers MOSEK [27] and SDPT3 [28]. This yields $\rho_{\text{d}} = 0.74$, $\rho_{\text{o}} = 0.67$, $\rho_{\text{s}} = 0.78$ and $\sigma_{\text{o},w}(a) = 2.25a$, $\sigma_{\text{s},\text{o}}(a) = 1.04a$, and $\sigma_{\text{s},w}(a) = 2.23a$.

For simplicity, we design a safe set $\mathbb{S}_{\text{f}}$ defined by a quadratic Lyapunov function V_{f}, that also satisfies Assumption 3. The safe set and the associated controller are determined as in [10] and the size of the safe set is determined online based on the error bounds as in [24]. Since the constraints are polytopic, we use the constraint tightening strategy from [2].

The RPOF-SF uses a horizon $N = 40$ and the MHE uses a backward horizon $M = 10$. The uncertified control inputs at every time step k are obtained from an arbitrary sinusoidal control input signal. We assume $\hat{x}_0 = x_0 \iff \bar{e}_0 = 0$ with the initial state $x_0 = \begin{bmatrix} 0.79 & 0.7 \end{bmatrix}^{\mathsf{T}}$. We implement the online computation (see Alg. 1) in Matlab using Casadi [29] and solve the MPC and MHE using IPOPT [30], with the MHE optimization limited to 1e3 iterations.

The closed-loop behavior of the proposed RPOF-SF is shown in Fig. 2. This highlights the successful certification of an arbitrary control policy, which would otherwise lead to constraint violations. The proposed RPOF-SF is able to modify the control inputs at every time step to achieve constraint satisfaction despite state disturbances, measurement uncertainties, and estimation errors. We emphasize that the safety filter allows safe operation even in proximity to a constraint boundary due to the over-approximated error bounds.
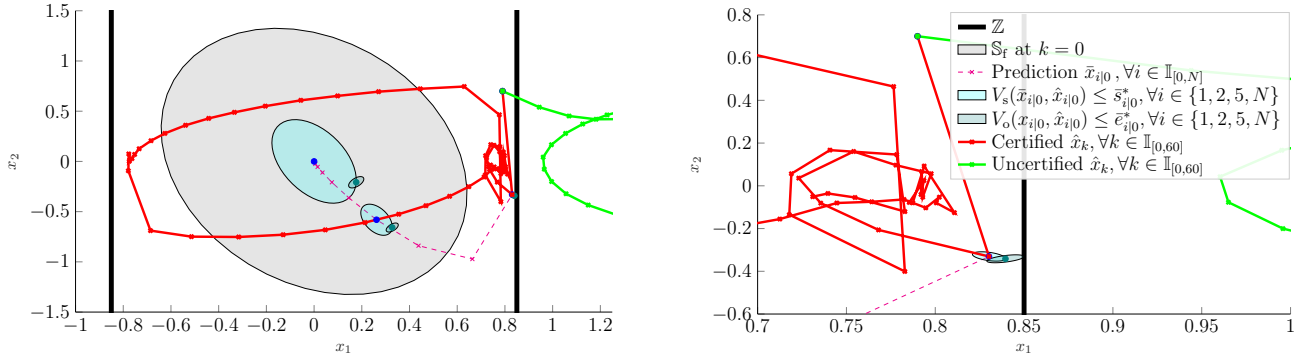
Fig. 2: Closed-loop behavior of the proposed RPOF-SF for 60 time steps. The constraint boundaries (black) and closed-loop trajectories of the estimated state for *(i)* the unfiltered control input $u_k = \pi_{\mathcal{L}}(k)$ (green) and *(ii)* the certified control input as provided through the RPOF-SF (red) are shown. The first open-loop trajectory (magenta, dashed) starting from $x_0 = \hat{x}_0 = \begin{bmatrix} 0.79 & 0.7 \end{bmatrix}^{\mathsf{T}}$ and the predicted ellipsoids for the estimation error bound given by V_{o} (teal) and the prediction error bound given by V_{s} (blue) for the predicted time steps $k = \{1, 2, 5, N = 40\}$ are also displayed. The terminal set (gray) is shown for the first open-loop prediction. **Left:** The figure on the left shows that the unfiltered control input immediately leads to constraint violations. In contrast, the certified control inputs from our robust predictive output-feedback safety filter achieve constraint satisfaction for all time steps. **Right:** The figure on the right shows a close-up of (a). This highlights constraint satisfaction of the RPOF-SF. Furthermore, the worst-case ellipsoid for V_{o} (teal) touches the constraint, which validates the constraint tightening.

VII. CONCLUSIONS

In this paper, we proposed a robust predictive output-feedback safety filter (RPOF-SF) for certifying arbitrary control inputs applied to disturbed nonlinear systems without full-state measurements. The efficacy of the proposed RPOF-SF approach for guaranteeing constraint satisfaction under uncertainties is proved in theory and demonstrated using a mass-spring-damper system as a numerical example.

As future work, we plan to generalize the proposed RPOF-SF approach to a broader class of uncertain systems by incorporating probabilistic learning techniques and use the proposed approach to guide reinforcement learning to improve sampling efficiency.

REFERENCES

[1] L. Brunke, M. Greeff, A. W. Hall, Z. Yuan, S. Zhou, J. Panerati, and A. P. Schoellig, "Safe learning in robotics: From learning-based control to safe reinforcement learning," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, 2022.

[2] J. Köhler, M. A. Müller, and F. Allgöwer, "Robust output feedback model predictive control using online estimation bounds," 2021. [Online]. Available: https://arxiv.org/abs/2105.03427

[3] K. P. Wabersich and M. N. Zeilinger, "A predictive safety filter for learning-based control of constrained nonlinear dynamical systems," *Automatica*, vol. 129, pp. 1–13, 2021.

[4] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *Proc. of the European Control Conference (ECC)*, 2019, pp. 3420–3431.

[5] A. J. Taylor, A. Singletary, Y. Yue, and A. D. Ames, "A control barrier perspective on episodic learning via projection-to-state safety," *IEEE Control Systems Letters (L-CSS)*, vol. 5(3), pp. 1019–1024, 2020.

[6] L. Brunke, S. Zhou, and A. P. Schoellig, "Barrier Bayesian linear regression: Online learning of control barrier conditions for safety-critical control of uncertain systems," in *Proc. of the Conference on Learning for Dynamics and Control (L4DC)*, 2022, accepted.

[7] I. Mitchell, A. Bayen, and C. Tomlin, "A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games," *IEEE Transactions on Automatic Control*, vol. 50(7), pp. 947–957, 2005.

[8] J. F. Fisac, N. F. Lugovoy, V. Rubies-Royo, S. Ghosh, and C. J. Tomlin, "Bridging Hamilton-Jacobi safety analysis and reinforcement learning," in *Proc. of the International Conference on Robotics and Automation (ICRA)*, 2019, pp. 8550–8556.

[9] J. H. Gillula and C. J. Tomlin, "Guaranteed safe online learning via reachability: tracking a ground target using a quadrotor," in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, 2012, pp. 2723–2730.

[10] J. Rawlings, D. Mayne, and M. Diehl, *Model Predictive Control: Theory, Computation, and Design*. Nob Hill Publishing, LLC, 01 2017.

[11] J. Köhler, R. Soloperto, M. A. Müller, and F. Allgöwer, "A computationally efficient robust model predictive control framework for uncertain nonlinear systems," *IEEE Transactions on Automatic Control*, vol. 66(2), pp. 794–801, 2021.

[12] S. Singh, B. Landry, A. Majumdar, J.-J. E. Slotine, and M. Pavone, "Robust feedback motion planning via contraction theory," *International Journal of Robotics Research*, 2019, submitted.

[13] F. Bayer, M. Bürger, and F. Allgöwer, "Discrete-time incremental iss: A framework for robust nmpc," in *Proc. of the European Control Conference (ECC)*, 2013, pp. 2068–2073.

[14] D. Q. Mayne, S. V. Raković, R. Findeisen, and F. Allgöwer, "Robust output feedback model predictive control of constrained linear systems," *Automatica*, vol. 42(7), pp. 1217–1222, 2006.

[15] M. Kögel and R. Findeisen, "Robust output feedback mpc for uncertain linear systems with reduced conservatism," *Proc. of the IFAC World Congress*, vol. 50(1), pp. 10 685 – 10 690, 2017.

[16] J. Lorenzetti and M. Pavone, "A simple and efficient tube-based robust output feedback model predictive control scheme," *Proc. of the European Control Conference (ECC)*, pp. 1775–1782, 2020.

[17] L. Brunke, S. Zhou, and A. P. Schoellig, "RLO-MPC: Robust learning-based output feedback mpc for improving the performance of uncertain systems in iterative tasks," in *Proc. of the IEEE Conference on Decision and Control (CDC)*, 2021, pp. 2183–2190.

[18] D. A. Copp and J. P. Hespanha, "Simultaneous nonlinear model predictive control and state estimation," *Automatica*, vol. 77, pp. 143 – 154, 2017.

[19] J. Löfberg, "Towards joint state estimation and control in minimax mpc," *Proc. of the IFAC World Congress*, vol. 35(1), pp. 273 – 278, 2002.

[20] D. A. Allan, J. Rawlings, and A. R. Teel, "Nonlinear detectability and incremental input/output-to-state stability," *SIAM Journal on Control and Optimization (SICON)*, vol. 59(4), pp. 3017–3039, 2021.

[21] P. Zhao, A. Lakshmanan, K. Ackerman, A. Gahlawat, M. Pavone, and N. Hovakimyan, "Tube-certified trajectory tracking for nonlinear systems with robust control contraction metrics," *IEEE Robotics and Automation Letters (RA-L)*, vol. 7(2), pp. 5528–5535, 2022.

[22] P. J. Koelewijn and R. Tóth, "Incremental stability and performance analysis of discrete-time nonlinear systems using the lpv framework," in *Proc. of the IFAC Workshop on Linear Parameter Varying Systems (LPVS)*, vol. 54(8), 2021, pp. 75–82.

[23] J. Köhler, M. A. Müller, and F. Allgöwer, "A nonlinear model predictive control framework using reference generic terminal ingredients," *IEEE Transactions on Automatic Control (TAC)*, vol. 65(8), pp. 3576–3583, 2020.

[24] J. Köhler, F. Allgöwer, and M. A. Müller, "A simple framework for nonlinear robust output-feedback mpc," in *Proc. of the European Control Conference (ECC)*, 2019, pp. 793–798.

[25] L. Magni, G. De Nicolao, R. Scattolini, and F. Allgöwer, "Robust model predictive control for nonlinear discrete-time systems," *International Journal of Robust and Nonlinear Control*, vol. 13(3-4), pp. 229–246, 2003.

[26] J. Löfberg, "Yalmip : A toolbox for modeling and optimization in matlab," in *Proc. of the CACSD Conference*, Taipei, Taiwan, 2004.

[27] M. ApS, *The MOSEK optimization toolbox for MATLAB manual. Version 9.0.*, 2019. [Online]. Available: http://docs.mosek.com/9.0/toolbox/index.html

[28] K. C. Toh, M. J. Todd, and R. H. Tütüncü, "Sdpt3 — a matlab software package for semidefinite programming, version 1.3," *Optimization Methods and Software*, vol. 11(1-4), pp. 545–581, 1999.

[29] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi – A software framework for nonlinear optimization and optimal control," *Mathematical Programming Computation*, vol. 11(1), pp. 1–36, 2019.

[30] A. Wächter and L. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical Programming*, vol. 106, pp. 25–57, 2006.