

From Raw Data to Safety: Reducing Conservatism by Set Expansion

Mohammad Bajelani

Klaske van Heusden

School of Engineering

University of British Columbia,

3333 University Way,

Kelowna, BC Canada, V1V 1V7

MOHAMMAD.BAJELANI@UBC.CA

KLASKE.VANHEUSDEN@UBC.CA

Abstract

Safety filters have been proposed as a modular technique in response to safety concerns associated with learning-based algorithms. Generally, these filters heavily rely on the system’s model, which is contradictory if they are intended to enhance data-driven or end-to-end learning solutions. This paper extends our previous work, a purely Data-Driven Safety Filter (DDSF) based on Willems’ lemma, to extremely short-sighted and non-conservative settings. Specifically, we propose online and offline sample-based methods to expand the safe set of DDSF and reduce its conservatism. Since this method is defined in an input-output framework, it can systematically handle unknown time-delay LTI systems using only one batch of data. The proposed method is applied to a time-delay system under various settings to evaluate its performance. The simulation results validate the effectiveness of the set expansion algorithm in generating a notably large input-output safe set, resulting in safety filters that are not conservative, even with an extremely short prediction horizon.

Keywords: Data-Driven Safety Filter, Learning-based Control, Behavioural System Theory, Time-Delay Systems, Reinforcement Learning.

1. Introduction

Nowadays, learning-based and data-driven approaches outperform traditional controllers in terms of superior performance without the assistance of expert knowledge, especially in addressing hard-to-model problems (Brunke et al., 2022). Despite these advancements, integrating such methods poses challenges in ensuring real-world safety during the learning process (Hewing et al., 2020). Generally, the safety criteria are defined by input-output constraints, such as restricted torque and a pre-specified workspace in a robot manipulator (Wabersich and Zeilinger, 2021) or the load factor in flight envelope for an aircraft (Hsu et al., 2023). To respect these constraints during the learning process, safety filters have been proposed within the control community as a modular technique to ensure safety irrespective of the learning method, whether provided by a human or a reinforcement learning agent (Wabersich and Zeilinger, 2018). The main idea behind these safety filters is the mapping of potentially unsafe learning inputs to the nearest safe learning inputs while satisfying specific criteria, such as ensuring the system’s safety, i.e., forcing the system’s trajectory to live in some invariant sets, over a finite or infinite time horizon.

Currently, there are three main streams to designing safety filters (Wabersich et al., 2023), which all rely on accurate and explicit models. Using Control Barrier Functions (CBFs) allows for creating safety filters capable of gently adjusting a specified input control signal as the system approaches the boundary of a control invariant set (Ames et al., 2019). Despite CBFs’ proven effectiveness in applications requiring swift responses (Wang et al., 2018), a notable challenge arises in formulating

barrier functions or ensuring the feasibility of QP-based CBFs in the presence of input constraints (Molnar and Ames, 2023). Hamilton-Jacobi (HJ) reachability analysis is a tool for ensuring system safety by defining reachable sets in the state space. These sets represent areas where specific goals can be achieved, or safety conditions are satisfied (Bansal et al., 2017). However, a major drawback is its computational complexity, especially for higher-dimensional systems (Herbert et al., 2021). Model Predictive Safety Filters (MPSFs) use an explicit predictive model to evaluate the safety of incoming learning inputs and then map them to the closest safe actions by constructing the backup trajectories (Tearle et al., 2021). The primary challenges associated with this approach are the substantial computational requirements for online processing and the inherent complexity of robust design. For a comprehensive examination of their advantages and disadvantages and data-driven approaches for the modeling component, see the recent survey (Wabersich et al., 2023).

It is critical to underscore that these methods depend on explicit models represented in state space, constituting a drawback for describing the uncertainty of the (multi-step) prediction (Köhler et al., 2022), as determining non-conservative bounds for state-space predictions is challenging. Furthermore, if these safety methods are intended to guarantee the safety of end-to-end learning algorithms, they should bypass the system identification process; otherwise, it is *counterproductive* to adopt an online learning algorithm to learn the system for which an accurate explicit model is available. In the indirect approach (system identification & control), uncertainty from the data must be transmitted through the system identification process, potentially causing mismatches with preferred uncertainty quantification for control design. In contrast, direct approaches (directly from data to control) seamlessly integrate uncertainty into control design, eliminating the need for complex uncertainty propagation and offering a transparent and effective strategy in data-driven control (Dörfler, 2023).

This paper proposes an entirely data-driven approach to determine safe terminal sets in the input-output framework. Combined with the recently introduced data-driven safety filter (DDSF) (Bajelani and van Heusden, 2023), this work achieves a non-conservative, short-sighted safety filter law directly from raw data. It should be noted that if an explicit model is available, terminal safe sets can be calculated. However, results for computation of terminal sets in the data-driven framework are limited to (trivial) solutions like an equilibrium point (Berberich et al., 2020a,c) or knowledge of the system’s lag (Berberich et al., 2021). Motivated by (Rosolia and Borrelli, 2017a,b), we extend the set expansion defined in the state-space framework to the input-output framework. By solving an offline optimization problem, the proposed method results in an input-output safe set, requiring only a single batch of data. For readability, noise-free measurements are used to develop the methods described in this paper. While the extension to data-driven predictive control that considers noise can be applied directly to the DDSF, the impact of noise or unmodeled dynamics on terminal sets has received little attention. The proposed data-driven set expansion offers a direct connection between noise and computed sets, extending the data-driven framework’s transparency to terminal sets.

The remaining sections of the paper are structured as follows: Section 2 covers preliminary material for the data-driven predictive safety filter based on Willems’ Lemma. Moving to Section 3, the formulation of DDSF with a sampled safe set is presented, demonstrating how the final set can be expanded using the input-output trajectories of the system or backup trajectories calculated by the DDSF. Section 4 provides simulation results for various settings, focusing on a second-order system with an unknown time delay by online and offline set expansion algorithms. Lastly, Section 5 comprises the discussion, concluding remarks, and potential avenues for future exploration.

2. Background

Consider a discrete-time LTI system as follows:

$$x(t+1) = Ax(t) + Bu(t - \tau_d), \quad y(t) = Cx(t) + Du(t - \tau_d), \quad (1)$$

where $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, $C \in \mathbb{R}^{p \times n}$, $D \in \mathbb{R}^{p \times m}$ are unknown matrices. The input, state, and output vectors are respectively denoted by $u(t) \in \mathbb{R}^m$, $x(t) \in \mathbb{R}^n$, $y(t) \in \mathbb{R}^p$ at time $t \in \mathbb{Z}_{\geq 0}$. Note that in addition to the matrices, the order of the system and the time delay $\tau_d \in \mathbb{Z}_{\geq 0}$ are unknown. It is also assumed that the pair of (A, B) and (A, C) are controllable and observable. The time-delay system (1) can be represented as an augmented delay-free system (2) with additional internal states $d(t) \in \mathbb{R}^{(k_d \times m) \times 1}$ as follows,

$$x_a(t+1) = A_a x_a(t) + B_a u(t), \quad y(t) = C_a x_a(t) + Du(t), \quad (2a)$$

$$x_a(t) := \begin{bmatrix} x(t) \\ d(t) \end{bmatrix}, \quad d(t) := \begin{bmatrix} u(t - \tau_d) \\ u(t - \tau_d + 1) \\ \vdots \\ u(t - 1) \end{bmatrix}, \quad (2b)$$

$$A_a := \begin{bmatrix} A_{n \times n} & B_{n \times m} & 0_{n \times (k_d - 1) \times m} \\ 0_{m \times (k_d - 1) \times (n + m)} & I_{m \times (k_d - 1) \times m} & 0_{m \times (k_d - 1) \times m} \\ 0_{m \times (n + m + k_d)} & 0 & 0 \end{bmatrix}, \quad (2c)$$

$$B_a := \begin{bmatrix} 0_{m \times m(k_d - 1)} \\ I_{m \times m} \end{bmatrix}, \quad C_a := \begin{bmatrix} C_{p \times n} \\ 0_{m \times k_d} \end{bmatrix}. \quad (2d)$$

In addition, the input-output admissible set is represented by a set of polytopes in the form of

$$u(t) \in \mathcal{U} := \{u \in \mathbb{R}^m \mid A_u u < b_u, A_u \in \mathbb{R}^{n_u \times m}, b_u \in \mathbb{R}^{n_u}\}, \quad (3a)$$

$$y(t) \in \mathcal{Y} := \{y \in \mathbb{R}^p \mid A_y y < b_y, A_y \in \mathbb{R}^{n_y \times p}, b_y \in \mathbb{R}^{n_y}\}, \quad (3b)$$

for all $t \in \mathbb{Z}_{\geq 0}$. The primary objective of a safety filter is to modify any (potentially) unsafe learning inputs, $u_t \in \mathbb{R}^m$, as minimally as possible while ensuring that the system (1) remains within the defined bounds $(\mathcal{U}, \mathcal{Y})$ defined in (3) for infinite time. Throughout this paper, it is assumed that a single noise-free input-output measured trajectory of the system (1) in the form of (4) is available, which is sufficiently exciting in the sense of Definition 3.

$$u_{[0, N_0 - 1]}^d = [u_0^\top, \dots, u_{N_0 - 1}^\top]^\top, \quad y_{[0, N_0 - 1]}^d = [y_0^\top, \dots, y_{N_0 - 1}^\top]^\top. \quad (4)$$

To build an implicit model from the pre-recorded trajectory (4), we reshape these two sequences into the form of Hankel matrices shown by (5). These matrices provide an implicit model to represent the input-output behavior of the system (1).

$$H_L(u) = \begin{bmatrix} u_0 & u_1 & \dots & u_{N_0 - L} \\ u_1 & u_2 & \dots & u_{N_0 - L + 1} \\ \vdots & \vdots & \ddots & \vdots \\ u_{L-1} & u_L & \dots & u_{N_0 - 1} \end{bmatrix}, \quad H_L(y) = \begin{bmatrix} y_0 & y_1 & \dots & y_{N_0 - L} \\ y_1 & y_2 & \dots & y_{N_0 - L + 1} \\ \vdots & \vdots & \ddots & \vdots \\ y_{L-1} & y_L & \dots & y_{N_0 - 1} \end{bmatrix}, \quad (5)$$

where $H_L(u) \in \mathbb{R}^{(mL) \times (N_0 - L + 1)}$ and $H_L(y) \in \mathbb{R}^{(pL) \times (N_0 - L + 1)}$. The raw data stored in the stacked Hankel matrices in (8) is used to directly parameterize any input-output trajectory of the system (1) using the fundamental Lemma proposed by Jan Willems in (Willems et al., 2005).

Definition 1 (System's Lag) $\underline{l} = l(A, C) + k_d$ denotes the lag of the system (1), in which $l(A, C)$ is the smallest integer that can make the observability matrix full rank.

$$l(A, C) := \left(C, CA, \dots, CA^{l-1} \right). \quad (6)$$

Definition 2 (Extended State (Berberich et al., 2021)) For some integers $T_{ini} \geq \underline{l}$, the extended state ξ_t at time t is defined as follows

$$\xi_t := \begin{bmatrix} u_{[t-T_{ini}, t-1]} \\ y_{[t-T_{ini}, t-1]} \end{bmatrix} \in \mathbb{R}^{(m+p)T_{ini} \times 1}. \quad (7)$$

where $u_{[t-T_{ini}, t-1]}$ and $y_{[t-T_{ini}, t-1]}$ denote the last T_{ini} input and output measurements, respectively.

Utilizing sufficiently long past input-output measurements, known as the extended state, enables the determination of the internal state of the underlying and unknown system (Coulson et al., 2019). The minimally required length of the extended state is dependent on the system's lag (Definition 1).

Definition 3 (Persistently Excitation (Berberich et al., 2020a)) Let the Hankel matrix's rank be $\text{rank}(H_L(u)) = mL$, then $u \in \mathbb{R}^m$ represents a persistently exciting signal of order L .

Behavioral system theory views dynamical systems differently than classical system theory. It focuses on systems as input-output trajectories defined in the signal subspaces (Markovsky et al., 2006). It has been demonstrated that if a finite-time trajectory of an LTI system is available and the input is persistently exciting, all the trajectories can be parameterized linearly by combining the columns of the Hankel matrix. This concept is known as the fundamental lemma introduced by Jan Willems (Willems et al., 2005). It also enables us to predict input-output trajectories of the true system (1) directly from raw data without requiring an explicit model. Using the implicit model derived from the fundamental lemma, we avoid the inevitable under- or over-modeling associated with explicit methods. For a recent overview of implicit vs explicit perspectives, see (Dörfler, 2023; Markovsky and Dörfler, 2021).

Theorem 4 (Fundamental Lemma (Berberich et al., 2020b)) Let u^d be persistently exciting of order $L + n$, and $\{u_k^d, y_k^d\}_{k=0}^{N_0-1}$ a trajectory of system G . Then, $\{\bar{u}_k, \bar{y}_k\}_{k=0}^{N_0-1}$ is a trajectory of system G if and only if there exists $\alpha \in \mathbb{R}^{N_0-L+1}$ such that

$$\begin{bmatrix} H_L(u^d) \\ H_L(y^d) \end{bmatrix} \alpha = \begin{bmatrix} \bar{u} \\ \bar{y} \end{bmatrix}. \quad (8)$$

Remark 1 The application of model-based safety filters poses challenges for systems with unknown time delays. The input-output framework in behavioral system theory inherently addresses this, emphasizing essential input-output measurements and an overestimation of the system's lag (6). Given the equivalent augmented system representation that treats the delayed inputs as internal states, denoted by $d(t)$ in (2), the input-output behavior of the implicit model provided in Theorem (4) and systems (1-2) are equivalent.

Definition 5 (Invariant Set) A set $\Xi \subseteq \mathbb{R}^{mT_{ini}+pT_{ini}}$ is said to be a control invariant set for the system (1) subjected to constraints in (3) from the input-output perspective, if

$$\xi(t) \in \Xi \quad \Rightarrow \quad \exists u(t) \in \mathcal{U} \text{ such that } \xi(t+1) \in \Xi, \quad \forall t \in \mathbb{Z}_{\geq t}.$$

Definition 6 (Input-output Equilibrium Point) *The extended state ξ^s is an input-output equilibrium point of system (1) if it is defined by the sequence $\{u_k, y_k\}_{k=0}^{T_{\text{ini}}-1}$ with a constant value $(u_k, y_k) = (u^s, y^s)$.*

Note that the equilibrium point in Definition 6 represents a special case of an invariant set as defined in Definition 5. In essence, if the system's input and output can be maintained at the equilibrium point for more than T_{ini} steps, all underlying states of the system (1) are fixed at their equilibrium values. The next section uses backup trajectories and sampled extended states to construct the final safe set. The backup trajectories are generated by the optimization problem (12) that steers the system to a final safe set, and sampled extended states are input-output measurements.

3. Data-Driven Safety Filter with Sampled Terminal Sets

This section introduces an extended version of DDSF based on behavioral system theory, as proposed in (Bajelani and van Heusden, 2023). More specifically, the prior method assumed that the terminal safe set should align with the system's equilibrium point, resulting in an excessively conservative system, particularly when dealing with a short prediction horizon. Consequently, the safety filter prevents learning by over-correcting the learning input. To tackle this challenge, this paper proposes two sample-based approaches motivated by (Rosolia and Borrelli, 2017b; Wabersich and Zeilinger, 2018). The first approach is designed for online implementation, which only needs new experiment data. In contrast, the second one allows the offline calculation of data-driven safe sets using a single dataset. The sampled safe sets are represented in (9a) for the online approach at time t and (9b) for the offline approach at iteration K , respectively.

$$\bar{\Xi}_f^t = \text{Conv}\left(\bigcup_{t=0}^{t-1} \xi_t\right), \quad (9a)$$

$$\bar{\Xi}_f^K = \text{Conv}\left(\bigcup_{l=0}^{K-1} \bigcup_{j=0}^{N-1} \bar{\xi}_j^l\right). \quad (9b)$$

where $\text{Conv}(\cdot)$ is the convex hull, ξ_t is the extended state containing past input-output measurements at time t , $\bar{\xi}_j^l$ is the extended state provided by the l^{th} element of backup trajectory at iteration j , K is the number of iteration, and N is the prediction horizon. Using the definition of the convex hull, the terminal sets (9a-9b) can be written in the form of (10a-10b) for the online and offline approaches.

$$\bar{\Xi}_f^t = \left\{ \sum_{t=0}^{t-1} \alpha_t \xi_t : \alpha_t \geq 0, \sum_{t=0}^{t-1} \alpha_t = 1 \right\}, \quad (10a)$$

$$\bar{\Xi}_f^K = \left\{ \sum_{l=0}^{K-1} \sum_{j=0}^{N-1} \alpha_j^l \bar{\xi}_j^l : \alpha_j^l \geq 0, \sum_{j=0}^{N-1} \alpha_j^l = 1 \right\}. \quad (10b)$$

For linear systems, if the input-output constraints $(\mathcal{U}, \mathcal{Y})$ are polytopes, then the safe set is convex. Therefore, the safe set can be reconstructed by sampling any safe trajectories and their convex hull. It should be noted that a safe trajectory can be sampled by actual measurements or backup trajectories, leading to algorithms 1 and 2, respectively.

Algorithm 1 Set Expansion (Online)

Input: $H_L(u^d), H_L(y^d), N_p, T_{ini}, t^*$, and $\Xi^{t=0}$
Output: Ξ_f

 initialization **while true do**

Solve problem (12) for $u_l(t)$
 Apply the safe input to the system (1) and
 update the initial condition
 Expand the safe set Ξ^t using (9a)
if $\Xi^{t-1} \approx \Xi^t$ & $t \geq t^*$ **then**
 | break the while loop
end
 $t \Rightarrow t + 1$

end

a. t^* is the minimum number of time steps.

Algorithm 2 Set Expansion (Offline)

Input: $H_L(u^d), H_L(y^d), N_p, T_{ini}, l^*$, and $\Xi^{l=0}$.

Output: Ξ_f

 initialization **while true do**

Solve problem (12) for $u_l(l)$
 Choose an element of backup trajectory as
 the initial condition
 Expand the safe set Ξ^l using (9b)
if $\Xi^{l-1} \approx \Xi^l$ & $l \geq l^*$ **then**
 | break the while loop
end
 $l \Rightarrow l + 1$

end

b. l^* is the minimum number of iterations.

If more measurements or backup trajectories are collected, the final set Ξ_f can be extended (9a-9b). Hence, the size of the sampled safe set does not decrease with the addition of more data, indicating that if the first safe set is not empty, then the subsequent ones are also not empty, as described as follows,

$$\bar{\Xi}_f^{t-1} \subseteq \bar{\Xi}_f^t, \quad (11a)$$

$$\bar{\Xi}_f^{K-1} \subseteq \bar{\Xi}_f^K. \quad (11b)$$

A visualization of the evolution of these sets for the online approach is shown in Figure (1). By comparing Figures (1A) and (1B), it can be seen that when the final safe set is expanded, conservatism is reduced and performance is improved. This allows the learning algorithm to explore the space more, and the learning inputs are less altered by the safety filter. A significant advantage of the online approach (1) is the possibility of collecting a sample safe set after each time step without collecting a complete iteration as proposed in (Rosolia and Borrelli, 2017a). Furthermore, the offline approach enables us to expand the safe set by using only one batch of data without running any experiments. The formulation of DDSF with the sampled safe set, defined by (10a-10b), is as follows:

$$\min_{\alpha(t), \bar{u}(t), \bar{y}(t)} \|\bar{u}_0(t) - u_l(t)\|_R^2 \quad (12a)$$

$$\text{s.t.} \quad \begin{bmatrix} \bar{u}(t) \\ \bar{y}(t) \end{bmatrix} = \begin{bmatrix} H_L(u^d) \\ H_L(y^d) \end{bmatrix} \alpha(t), \quad (12b)$$

$$\begin{bmatrix} \bar{u}_{[-T_{ini}, -1]}(t) \\ \bar{y}_{[-T_{ini}, -1]}(t) \end{bmatrix} = \begin{bmatrix} u_{[t-T_{ini}, t-1]} \\ y_{[t-T_{ini}, t-1]} \end{bmatrix}, \quad (12c)$$

$$\bar{\xi}^{N-1} \in \Xi_f, \quad (12d)$$

$$\bar{u}_k(t) \in \mathcal{U}, \quad \bar{y}_k(t) \in \mathcal{Y}, \quad k \in \{0, \dots, N-1\}. \quad (12e)$$




Figure 1: A visualization of the final set is shown. The solid black line shows the realized system trajectory. The dashed lines are the system’s backup trajectories at time t . The blue polytopes show the convex hulls provided by extended states’ realized trajectory. (A) The final set is the system’s equilibrium point as proposed in (Bajelani and van Heusden, 2023). (B) The terminal set is expanded by the algorithm (1) at times $t - 2$, $t - 1$, and t .

By minimizing the cost function (12a), the first input $u_0(t)$ becomes the nearest safe input to the potentially unsafe input u_l while adhering to constraints (12b-12e). The implicit model for prediction, driven by measured raw data (4), adheres to (12b) based on Theorem 4. Considering past input-output measurements, the optimization problem’s initial condition is fixed using (12c). Ξ_f in (12d) is defined as the terminal set for recursive feasibility and building backup trajectories, representing the sampled safe set calculated by the algorithm (1) at time t as Ξ_f^{t-1} or the algorithm (2) at iteration K as Ξ_f^K . For $t = 1$ or $l = 1$, Ξ_f^0 is assumed to be a known equilibrium point for the system (1). Additionally, we consider that input-output constraints are defined by (12e). The following assumptions proposed in (Bajelani and van Heusden, 2023) are also adopted.

Assumption 1 (Prediction Horizon Length) *The prediction horizon N is greater than $T_{ini} \geq l$.*

Assumption 2 (Persistent Excitation) *The stacked Hankel matrix defined in Definition 8 is PE of order $L = N + 2T_{ini}$ in the sense of Definition 3.*

Assumption 3 (Terminal Safe Set) *The equilibrium point of the system (1) defined in Definition 6 is known, belongs to admissible sets $(\mathcal{U}, \mathcal{Y})$, and used as the initial safe set of algorithms (1-2).*

Theorem 7 (Recursive Feasibility of DDSF with sampled safe set - online algorithm) *Let assumptions (1-3) hold, $T_{ini} \geq l$, $\Xi_f^{t_0}$ be a non-empty set. Then, the DDSF optimization problem (12) is feasible for all $t > t_0$, if it is feasible at $t = t_0$.*

Proof. At $t = t_0$ the solution of optimization problem (12) gives an input-output trajectory from the initial condition $\xi_{t=t_0}$ to the terminal set $\Xi_f^{t=t_0-1}$, as it is assumed the problem (12) is feasible at $t = t_0$. At $t = t_0 + 1$, a new backup trajectory is calculated and the terminal set is expanded to $\Xi_f^{t=t_0}$ by measuring $\xi_{t=t_0}$. Based on (11), the terminal safe set is expanded implying $\Xi^{t=t_0-1} \subseteq \Xi^{t=t_0}$. Because the feasible set of the system (1) is a convex set, it follows that the convex hull of all the realized extended states (or backup trajectories) is a subset of this set. Suppose a system’s trajectory enters this set. In that case, it will remain there permanently under the policy of (12), thus demonstrating that all sampled safe sets are invariant and a subset of the feasible set. Therefore, by relying on induction, we can conclude that the problem (12) remains feasible for all $t > t_0$.

Remark 2 *An under-approximation of the safe set may be necessary to expedite the optimization problem (12) for high-order systems. Specifically, when the system's lag is overestimated, the solution yields extremely high-dimensional polytopes, as the number of constraints depends on estimating the system's lag.*

4. Simulation Example

To demonstrate the effectiveness of the proposed method, a time-delay second-order system is considered, as presented in (13). This system allows us to visualize the sampled terminal sets.

$$x(t+1) = \begin{bmatrix} 1 & -0.1 \\ 0 & 1 \end{bmatrix} x(t) + \begin{bmatrix} 0 \\ 0.1 \end{bmatrix} u(t - \tau_d), \quad y(t) = \begin{bmatrix} 1 & 0 \end{bmatrix} x(t). \quad (13)$$

where the input-output constraints are assumed as follows,

$$u(t) \in \mathcal{U} := \{u \in \mathbb{R} \mid -1 \leq u \leq +1\}, \quad (14a)$$

$$y(t) \in \mathcal{Y} := \{y \in \mathbb{R} \mid -1 \leq y \leq +1\}. \quad (14b)$$

To show the effect of the sampled safe set by the online and offline set expansion algorithms (1-2) and the overestimation of the system's lag, two case studies are investigated. The impact of the sampled terminal set on the performance of a DDSF is demonstrated in section (4.1) for three safe sets, equilibrium point, and safe sets provided by online and offline algorithms when the system is subjected to delayed input. In section (4.2), it is assumed that there is no delayed input, but the effect of overestimating the system's lag is illustrated. Note that DDSFs are designed to keep input-output trajectories safe regardless of the control input; hence, we utilized unsafe control inputs to destabilize the system in this section. The oscillatory behavior of the output in the safe set, shown in light green, results from the sinusoidal learning inputs and the correction provided by the DDSF.

4.1. First study: Sampled safe sets vs Equilibrium point safe set

For this study, a short prediction horizon $N_p = 6$, dead-time $\tau_d = 1$, and the corresponding known system's lag $T_{ini} = 3$ are selected. The learning input is also assumed to be a sinusoidal signal bounded in $[-1, +1]$. The equilibrium point $(u^s, y^s) = (0, 0)$ is used as the initial terminal safe set, and algorithms (1-2) are implemented using the CasADi toolbox (Andersson et al., 2019). The learning and safe inputs, along with the system's response for terminal sets, are illustrated in Figure (2). It is evident that employing the sampled safe set as the terminal condition of (12) reduces the conservatism of DDSF. Furthermore, it can be seen that at time $t = 50 [s]$, the online approach can reach the boundary of the defined safe set. However, the convergence of the algorithms must be determined by the convergence of the safe set. The final safe sets for the online (1) and offline (2) algorithms are illustrated in Figure (3) at different times and iterations. It must be noted that the final learned safe set provided by algorithms (1) and (2) are the same. In other words, both the online (1) and offline (2) algorithms yield the same safety filters after expanding the safe set.

4.2. Second study: Overestimation of system's lag

In this part, it is assumed that $\tau_d = 0$, the algorithm (2) is employed for $T_{ini} \in \{2, 3\}$ and the learning input is a PRBS signal. The input-output behavior is depicted in Figure (4). It is clear that by overestimating the system's lag, the safety and performance of DDSF have not been affected, as the input-output behavior is the same.

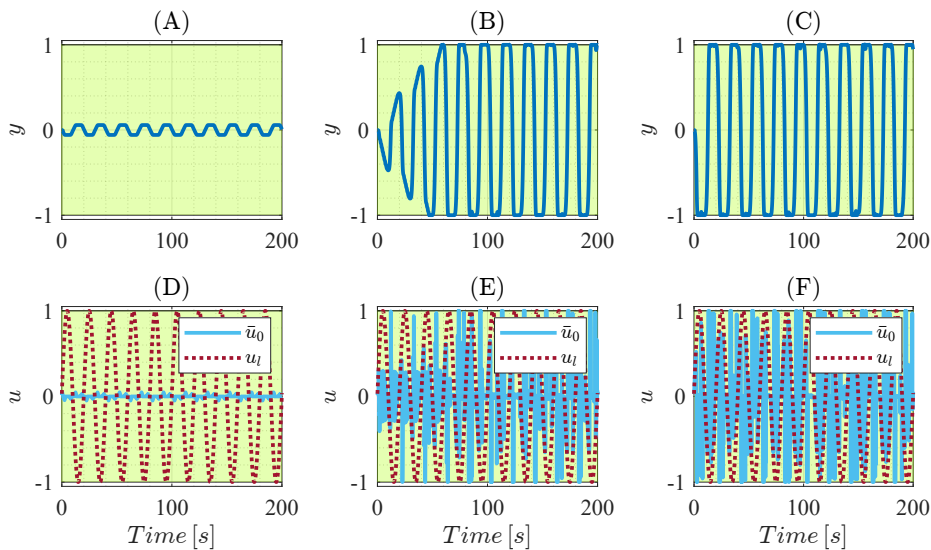


Figure 2: First case study: Input-output trajectories of system (13) under the policy of (12) for three different safe sets: the system’s equilibrium point, subplots (A) and (D), the sampled safe set computed by the online algorithm (1), subplots (B) and (E), and the sampled safe set computed by the offline algorithm (2), subplots (C) and (F).

5. Conclusion

This paper introduced online and offline set expansion algorithms based on sampled data to alleviate the potential conservatism associated with the Data-Driven Safety Filter (DDSF) if implemented using short prediction horizons. Specifically, we demonstrated that input-output data can be used to compute safe sets through extended state and backup trajectories. With these contributions, all stages of the design of a data-driven safety filter, from the raw data to the safety filter and safe sets, are purely data-driven, with no explicit modeling requirement. Importantly, the proposed input-output framework can handle unknown time delays by simply overestimating the system’s lag. In contrast, time delays are rarely considered in model-based safety filters in the state-space framework, for which the exact delay would need to be known.

We showed that safe sets can be computed offline using an implicit model derived from limited measured data. If the exact model is known, computing these sets is straightforward using backward-reachable sets. Otherwise, characterizing the impact of parametric or unmodeled uncertainty on the size and shape of the safe set is challenging. Due to the multi-step prediction property of the data-driven method proposed here, the effect of measurement noise on the computed safe set is direct. Furthermore, with the proposed method, no unmodeled dynamics exist as long as a sufficiently exciting dataset is available and sufficiently long past input-output measurements are considered. A potential avenue for future research involves exploring the impact of noise-polluted measurements and disturbances on the estimation of safe sets and the prediction of backup trajectories. Furthermore, applying this methodology to controllers could be investigated, wherein backup trajectories and expanded terminal sets could be employed to design the terminal component of a data-driven predictive controller.

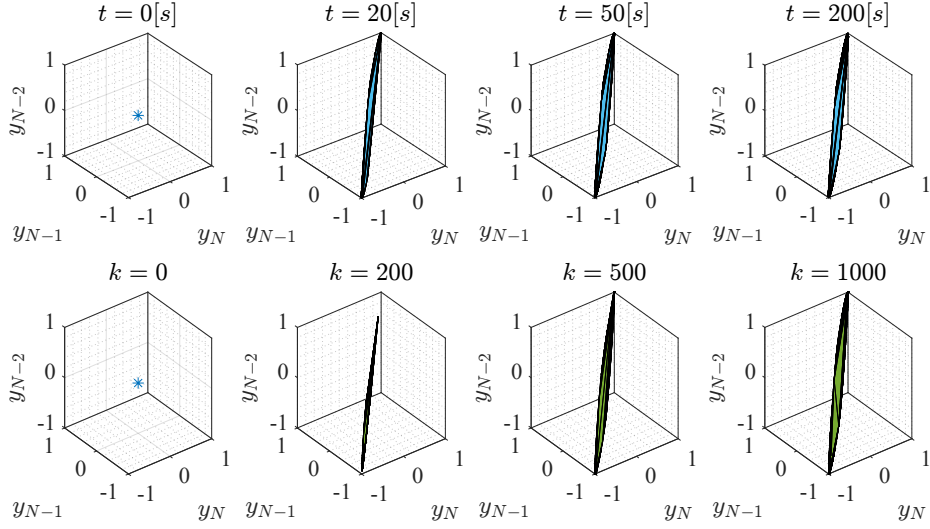


Figure 3: First case study: The sampled terminal safe set at four different times $t = 0[s]$, $t = 20[s]$, $t = 50[s]$, $t = 200[s]$ computed by the online algorithm (1), and the sampled terminal safe set (1) at four different iterations $k = 0$, $k = 200$, $k = 500$, $k = 1000$ computed by the offline algorithm (2).




Figure 4: Second case study: Input-output trajectories of system (13) under the policy of (12) for the exact and overestimated system's lags, $T_{ini} = 2$ and $T_{ini} = 3$, respectively.

Acknowledgments

We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC) [RGPIN-2023-03660].

References

- Aaron D Ames, Samuel Coogan, Magnus Egerstedt, Gennaro Notomista, Koushil Sreenath, and Paulo Tabuada. Control barrier functions: Theory and applications. In *2019 18th European control conference (ECC)*, pages 3420–3431. IEEE, 2019.
- Joel A E Andersson, Joris Gillis, Greg Horn, James B Rawlings, and Moritz Diehl. CasADi – A software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*, 11(1):1–36, 2019. doi: 10.1007/s12532-018-0139-4.
- Mohammad Bajelani and Klaske van Heusden. Data-driven safety filter: An input-output perspective. *arXiv preprint arXiv:2309.00189*, 2023.
- Somil Bansal, Mo Chen, Sylvia Herbert, and Claire J Tomlin. Hamilton-jacobi reachability: A brief overview and recent advances. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 2242–2253. IEEE, 2017.
- Julian Berberich, Johannes Köhler, Matthias A Müller, and Frank Allgöwer. Data-driven model predictive control with stability and robustness guarantees. *IEEE Transactions on Automatic Control*, 66(4):1702–1717, 2020a.
- Julian Berberich, Johannes Köhler, Matthias A Müller, and Frank Allgöwer. Robust constraint satisfaction in data-driven MPC. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pages 1260–1267. IEEE, 2020b.
- Julian Berberich, Johannes Köhler, Matthias A. Müller, and Frank Allgöwer. Data-driven tracking mpc for changing setpoints. *IFAC-PapersOnLine*, 53(2):6923–6930, 2020c. ISSN 2405-8963. 21st IFAC World Congress.
- Julian Berberich, Johannes Köhler, Matthias A Müller, and Frank Allgöwer. On the design of terminal ingredients for data-driven MPC. *IFAC-PapersOnLine*, 54(6):257–263, 2021.
- Lukas Brunke, Melissa Greeff, Adam W Hall, Zhaocong Yuan, Siqi Zhou, Jacopo Panerati, and Angela P Schoellig. Safe learning in robotics: From learning-based control to safe reinforcement learning. *Annual Review of Control, Robotics, and Autonomous Systems*, 5:411–444, 2022.
- Jeremy Coulson, John Lygeros, and Florian Dörfler. Data-enabled predictive control: In the shallows of the deepc. In *2019 18th European Control Conference (ECC)*, pages 307–312. IEEE, 2019.
- Florian Dörfler. Data-driven control: Part two of two: Hot take: Why not go with models? *IEEE Control Systems Magazine*, 43(6):27–31, 2023. doi: 10.1109/MCS.2023.3310302.
- Sylvia Herbert, Jason J Choi, Suvansh Sanjeev, Marsalis Gibson, Koushil Sreenath, and Claire J Tomlin. Scalable learning of safety guarantees for autonomous systems using hamilton-jacobi reachability. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5914–5920. IEEE, 2021.
- Lukas Hewing, Kim P Wabersich, Marcel Menner, and Melanie N Zeilinger. Learning-based model predictive control: Toward safe learning in control. *Annual Review of Control, Robotics, and Autonomous Systems*, 3:269–296, 2020.

- Ting-Wei Hsu, Jason J Choi, Divyang Amin, Claire Tomlin, Shaun C McWherter, and Michael Piedmonte. Towards flight envelope protection for the nasa tiltwing evtol flight mode transition using hamilton–jacobi reachability. *Journal of the American Helicopter Society*, 2023.
- Johannes Köhler, Kim P Wabersich, Julian Berberich, and Melanie N Zeilinger. State space models vs. multi-step predictors in predictive control: Are state space models complicating safe data-driven designs? In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 491–498. IEEE, 2022.
- Ivan Markovskiy and Florian Dörfler. Behavioral systems theory in data-driven analysis, signal processing, and control. *Annual Reviews in Control*, 52:42–64, 2021.
- Ivan Markovskiy, Jan C Willems, Sabine Van Huffel, and Bart De Moor. *Exact and approximate modeling of linear systems: A behavioral approach*. SIAM, 2006.
- Tamas G Molnar and Aaron D Ames. Safety-critical control with bounded inputs via reduced order models. *arXiv preprint arXiv:2303.03247*, 2023.
- Ugo Rosolia and Francesco Borrelli. Learning model predictive control for iterative tasks: A computationally efficient approach for linear system. *IFAC-PapersOnLine*, 50(1):3142–3147, 2017a.
- Ugo Rosolia and Francesco Borrelli. Learning model predictive control for iterative tasks. a data-driven control framework. *IEEE Transactions on Automatic Control*, 63(7):1883–1896, 2017b.
- Ben Tearle, Kim P Wabersich, Andrea Carron, and Melanie N Zeilinger. A predictive safety filter for learning-based racing control. *IEEE Robotics and Automation Letters*, 6(4):7635–7642, 2021.
- Kim P Wabersich and Melanie N Zeilinger. Linear model predictive safety certification for learning-based control. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 7130–7135. IEEE, 2018.
- Kim P Wabersich, Andrew J Taylor, Jason J Choi, Koushil Sreenath, Claire J Tomlin, Aaron D Ames, and Melanie N Zeilinger. Data-driven safety filters: Hamilton-jacobi reachability, control barrier functions, and predictive methods for uncertain systems. *IEEE Control Systems Magazine*, 43(5):137–177, 2023.
- Kim Peter Wabersich and Melanie N Zeilinger. A predictive safety filter for learning-based control of constrained nonlinear dynamical systems. *Automatica*, 129:109597, 2021.
- Li Wang, Evangelos A. Theodorou, and Magnus Egerstedt. Safe learning of quadrotor dynamics using barrier certificates. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2460–2465, 2018. doi: 10.1109/ICRA.2018.8460471.
- Jan C Willems, Paolo Rapisarda, Ivan Markovskiy, and Bart LM De Moor. A note on persistency of excitation. *Systems & Control Letters*, 54(4):325–329, 2005.