

# Predictive Safety Filter using System Level Synthesis

**Conference Paper****Author(s):**

[Leeman, Antoine](#) ; [Köhler, Johannes](#) ; Bennani, Samir; Zeilinger, Melanie N.

**Publication date:**

2023

**Permanent link:**

<https://doi.org/10.3929/ethz-b-000615512>

**Rights / license:**

[In Copyright - Non-Commercial Use Permitted](#)

**Originally published in:**

Proceedings of Machine Learning Research 211

**Funding acknowledgement:**

180545 - NCCR Automation (phase I) (SNF)

# Predictive safety filter using system level synthesis

Antoine P. Leeman<sup>†</sup>

Johannes Köhler<sup>†</sup>

Samir Bennani<sup>‡</sup>

Melanie N. Zeilinger<sup>†</sup>

ALEEMAN@ETHZ.CH\*

JKOEHL@ETHZ.CH

SAMIR.BENNANI@ESA.INT

MZEILINGER@ETHZ.CH

<sup>†</sup>*Institute for Dynamic Systems and Control, ETH Zurich, Switzerland*

<sup>‡</sup>*ESTEC, European Space Agency, the Netherlands*

**Editors:** N. Matni, M. Morari, G. J. Pappas

## Abstract

Safety filters provide modular techniques to augment potentially unsafe control inputs (e.g. from learning-based controllers or humans) with safety guarantees in the form of constraint satisfaction. In this paper, we present an improved model predictive safety filter (MPSF) formulation, which incorporates system level synthesis techniques in the design. The resulting SL-MPSF scheme ensures safety for linear systems subject to bounded disturbances in an enlarged safe set. It requires less severe and frequent modifications of potentially unsafe control inputs compared to existing MPSF formulations to certify safety. In addition, we propose an explicit variant of the SL-MPSF formulation, which maintains scalability, and reduces the required online computational effort - the main drawback of the MPSF. The benefits of the proposed system level safety filter formulations compared to state-of-the-art MPSF formulations are demonstrated using a numerical example.

**Keywords:** Safety Filter, Safety certification, Model Predictive Control, System Level Synthesis

## 1. Introduction

Learning-based controllers have demonstrated high performance for control of uncertain systems in complex environments, see, e.g. [Hwangbo et al. \(2019\)](#); [Ibarz et al. \(2021\)](#); [Mnih et al. \(2015\)](#) and references therein. However, their potential can still not be fully exploited in many industrial applications due to the lack of safety guarantees i.e., in terms of constraint satisfaction.

Safety filters are modular techniques that can be combined with any controller and modify/filter a potentially unsafe input as little as needed while ensuring that the state remains in some safe set. Methods based on Hamilton-Jacobi reachability analysis [Gillula and Tomlin \(2012\)](#); [Fisac et al. \(2019\)](#) provide a general approach for computing safe sets. However, their practical applicability is often limited due to scalability issues associated with solving nonlinear partial differential equations. Control barrier functions (CBFs) provide a popular alternative to define safe sets [Ames et al. \(2019\)](#), but their design commonly requires system-specific intuition related to the choice of the control Lyapunov function. Model predictive safety filters (MPSFs) [Wabersich and Zeilinger \(2018a, 2021\)](#) provide a flexible way to implicitly define a safe set and a safety filter as the solution to an optimisation problem, at the cost of increased online computational effort.

\* This work has been supported by the European Space Agency under OSIP 4000133352, the Swiss Space Center, and the Swiss National Science Foundation under NCCR Automation (grant agreement 51NF40 180545).

Overall, these three main approaches used to define safety filters offer different trade-offs between offline design complexity, online computational demand, and conservativeness in the design. Conservativeness is reflected by the safe set size and the amount of control intervention. Maximising the size of the safe set increases the operational range of the system and enables, e.g. more extensive (safe) exploration particularly in the learning setting. In addition, safety filter interventions affect the performance of the learning-based (or the human) controller, such as stability [Dai et al. \(2021\)](#), operational cost, or tracking error. This paper presents an MPSF approach that reduces conservativeness compared with existing MPSF by increasing the size of the safe set, and reducing the required safety filter interventions, and thereby enables higher performance for problems with larger model uncertainty.

The main concept of the MPSF is based on model predictive control (MPC): a safe backup trajectory is predicted, ensuring the system could be steered to a safe (invariant) set, while minimally modifying the potentially unsafe input  $u_{\mathcal{L}}$ . The safe set and safe backup strategy in MPSF schemes are thereby characterised by the solution to an optimisation problem. A key challenge is the computation of safe backup trajectories in the presence of disturbances and model uncertainties, which are particularly relevant in the learning setting. The first MPSF was presented in [Wabersich and Zeilinger \(2018a\)](#), and exploits the linear robust MPC (RMPC) from [Mayne et al. \(2005\)](#). This formulation was later extended to distributed systems [Muntwiler et al. \(2020\)](#), parametric uncertainties [Didier et al. \(2021\)](#), nonlinear systems [Wabersich and Zeilinger \(2021\)](#), and stochastic noise [Wabersich et al. \(2022\)](#). Notably, all of these MPSF formulations use an offline fixed auxiliary tube-controller to mitigate the effect of disturbances/uncertainties. While the tube-controller allows for a scalable and efficient offline design, this simplification can result in overly cautious MPSF schemes, limiting their practical performance.

**Contribution** In this paper, we propose a novel MPSF formulation focused on a robust design for linear systems, that addresses the main limitation of the fixed tube-controller, and hence significantly mitigates the conservativeness of existing MPSF schemes. To achieve this, we combine MPSF with system level synthesis (SLS) [Anderson et al. \(2019\)](#) to create a safety filter that ensures safety with reduced control interventions. In particular, SLS enables optimisation over affine feedback policies in RMPC [Chen et al. \(2021\)](#); [Sieber et al. \(2022\)](#), compare also disturbance feedback MPC [Goulart et al. \(2006\)](#) (Section 3.1). As a result, the proposed system level MPSF (SL-MPSF) overcomes the limitation of the fixed tube-controller in existing MPSF schemes, and ensures a larger safe set, and lower control intervention at the cost of increased computational complexity. The proposed flexible SL-MPSF hence simultaneously optimises the safe backup trajectory and tube-controller (Section 3.2). Notably, the presented SL-MPSF formulation uses a general terminal set to ensure recursive feasibility and does not require the finite-impulse-response constraint from [Sieber et al. \(2022\)](#), which constitutes an alternative RMPC formulation based on SLS.

The second contribution is an explicit safety filter that leverages the flexibility of SLS to design an explicit safe set and an explicit backup control law offline (Section 4). The resulting scheme is more conservative than SL-MPSF but requires no online optimisation. The explicit safe set is constructed by imposing the disturbance reachable set to safely return to the safe set within the prediction horizon. The safe set is defined using a simple shape and is hence easy to implement and scalable. Similar methods are used in [Parsi et al. \(2022\)](#) and ([Sieber et al., 2022](#), Section 4.2) to compute feedback policies and disturbance reachable sets offline, however, they still require online optimisation to ensure robust constraint satisfaction.

Both safety filters proposed in this paper offer a different trade-off between conservativeness, scalability, offline design complexity and online computation demand compared with available techniques and thereby enlarge the scope of this promising safety approach. The benefits of both methods are illustrated by comparing them against state-of-the-art safety filters using a numerical example (Section 5).

**Notation** The set of non-negative integers is given by  $\mathbb{N}_{\geq 0}$ . For a vector  $x \in \mathbb{R}^n$ , we denote the  $p$ -norm with  $p \in \{1, 2, \infty\}$  by  $\|x\|_p$ . The infinity-norm ball is given by  $\mathcal{B}_{\infty}^n := \{w \in \mathbb{R}^n \mid \|w\|_{\infty} \leq 1\}$ . For two sets  $\mathcal{W}_1$  and  $\mathcal{W}_2$ , the Minkowski sum is defined as  $\mathcal{W}_1 \oplus \mathcal{W}_2 := \{w_1 + w_2 \mid w_1 \in \mathcal{W}_1, w_2 \in \mathcal{W}_2\}$ . We also define  $\mathcal{W}^k$  as the Cartesian product of  $\mathcal{W}$  with itself  $k$  times, i.e.,  $\mathcal{W}^k = \mathcal{W} \times \dots \times \mathcal{W}$  ( $k$  times). The identity matrix is denoted by  $\mathcal{I}_n \in \mathbb{R}^{n \times n}$  and the dimensions are omitted if they can be inferred from the context. Let  $0_{p,q} \in \mathbb{R}^{p \times q}$  and  $0_n \in \mathbb{R}^n$  be respectively a matrix and a vector of zeros. We denote stacked vectors by  $(a, b) = [a^\top \ b^\top]^\top$ . The block diagonal matrix consisting of matrices  $A_1, \dots, A_T$  is denoted by  $\mathcal{A} = \text{diag}(A_1, \dots, A_T)$ . Let  $\mathcal{L}^{N,p \times q}$  denote the set of all block lower-triangular matrices with the following structure

$$M = \begin{bmatrix} M^{0,0} & 0_{p,q} & \dots & 0_{p,q} \\ M^{1,1} & M^{1,0} & \dots & 0_{p,q} \\ \vdots & \vdots & \ddots & \vdots \\ M^{N,N} & M^{N,N-1} & \dots & M^{N,0} \end{bmatrix},$$

where  $M^{i,j} \in \mathbb{R}^{p \times q}$ . We denote the  $k^{\text{th}}$  block row of  $M$  as  $M^k := [M^{k,k} \ \dots \ M^{k,0}, 0_{p,q(N-k)}]$ .

## 2. Preliminaries

We consider a linear system

$$x(t+1) = Ax(t) + Bu(t) + B_w w(t), \quad t \in \mathbb{N}_{\geq 0}, \quad (1)$$

with state  $x(t) \in \mathbb{R}^n$ , control input  $u(t) \in \mathbb{R}^m$  and unknown disturbance  $w(t) \in B_w \mathcal{W} \subset \mathbb{R}^n$ , at discrete time instances  $t \in \mathbb{N}_{\geq 0}$ . For simplicity of exposition, we consider a box  $\mathcal{W} := \mathcal{B}_{\infty}^n$ , and assume that  $B_w$  is an invertible square matrix, compare also Remark 3 below. The system is safe if it satisfies the following state and input constraints

$$x(t) \in \mathcal{X} := \{x \mid A_{x,j} x \leq b_{x,j}, \quad j = 1, \dots, n_x\}, \quad u(t) \in \mathcal{U} := \{u \mid A_{u,j} u \leq b_{u,j}, \quad j = 1, \dots, n_u\},$$

for all  $t \in \mathbb{N}_{\geq 0}$ , with  $A_{x,j} \in \mathbb{R}^n$ ,  $b_{x,j} > 0$ ,  $A_{u,j} \in \mathbb{R}^n$ ,  $b_{u,j} > 0$ . We assume that  $\mathcal{U}$  is compact. In this paper, we provide a safety filter and safe set, which enable a safety certificate for arbitrary control signals  $u_{\mathcal{L}}(t) \in \mathbb{R}^m$ .

**Definition 1** A set  $\mathcal{S} \subseteq \mathcal{X}$  and a control law  $u_{\mathcal{S}}(x(t), u_{\mathcal{L}}(t), t)$ , with  $u_{\mathcal{S}} : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{N}_{\geq 0} \mapsto \mathcal{U}$  are called a safe set and a safety filter, respectively, if the application of the safety filter  $u = u_{\mathcal{S}}$  in (1), ensures robust constraint satisfaction, i.e.,  $x(t) \in \mathcal{X} \forall t \geq \bar{t} \forall w(t) \in \mathcal{B}_{\infty}^n$ , if  $x(\bar{t}) \in \mathcal{S}$ .

We search for a safety filter  $u_{\mathcal{S}}$  and a corresponding safe set  $\mathcal{S} \subseteq \mathcal{X}$ , such that state and input constraints satisfaction can be guaranteed for all future time steps and for all disturbances. While safety is the primary objective, the safety filter should also interfere as little as possible, i.e., whenever possible the safety filter should yield  $u_{\mathcal{S}} = u_{\mathcal{L}}$ . The main benefit of a modular safety filter

according to Definition 1 is that it can certify any potentially unsafe control input  $u_{\mathcal{L}}$ , including the important case of learning-based controllers or humans.

In the nominal case (i.e., without disturbance  $w(t)$ ), the MPSF formulation proposed in [Wabersich and Zeilinger \(2018a\)](#) consists in solving the following optimisation problem at each time step

$$\min_{\mathbf{z}, \mathbf{v}} \|\mathbf{v}_0 - u_{\mathcal{L}}(t)\|_2^2 \quad \text{s.t.} \quad \mathbf{z}_{k+1} = A\mathbf{z}_k + B\mathbf{v}_k, \quad x(t) = \mathbf{z}_0, \quad k = 0, \dots, N-1, \quad (2a)$$

$$\mathbf{z}_k \in \mathcal{X}, \quad \mathbf{v}_k \in \mathcal{U}, \quad \mathbf{z}_N \in \mathcal{X}_f, \quad k = 0, \dots, N-1, \quad (2b)$$

where  $\mathbf{z}_k, \mathbf{v}_k$  denote the nominal state and input, respectively, and  $\mathcal{X}_f \subseteq \mathcal{X} \subseteq \mathbb{R}^n$  is a positively invariant set as commonly required in MPC, see e.g. [Kouvaritakis and Cannon \(2016\)](#). In order to reduce the effects of model uncertainties, MPSFs [Wabersich and Zeilinger \(2018a, 2021\)](#) employ robust MPC techniques based on a so-called tube-controller, i.e.,  $u(t) = v(t) + K(x(t) - z(t))$ , where  $K$  is fixed offline. In the following, we develop a robust counterpart of (2) based on SLS (Section 3.1), where we not only optimise the (nominal) backup trajectory  $\mathbf{z}$  and  $\mathbf{v}$ , but also a tube-controller and disturbance reachable set for the error between the uncertain and nominal system. The resulting feedback policy guarantees robust constraint satisfaction and reduces conservativeness compared to other MPSFs (Section 3.2). While offering a flexible approach, the required online optimisation can be computationally demanding. To address this limitation, we introduce an explicit system level safety filter (Section 4), which removes most of the offline computational complexity. Notably, both methods ensure robust constraint satisfaction.

### 3. System Level Model Predictive Safety Filter

In the following, we first present the parametrisation of the disturbance reachable sets under affine tube-controllers, before introducing the SL-MPSF formulation.

#### 3.1. Parameterisation of affine controllers

The employed parameterisation of affine feedback policies for LTI systems relies on SLS techniques [Goulart et al. \(2006\)](#); [Anderson et al. \(2019\)](#). We define the prediction state  $\mathbf{x}_k \in \mathbb{R}^n$  and input  $\mathbf{u}_k \in \mathbb{R}^m$  for system (1) as

$$\mathbf{x}_{k+1} = A\mathbf{x}_k + B\mathbf{u}_k + B_w\mathbf{w}_k, \quad k = 0, \dots, N-1, \quad (3)$$

where  $N \in \mathbb{N}_{\geq 1}$  is the prediction horizon and  $\mathbf{w}_k \in \mathcal{B}_{\infty}^n = \mathcal{W}$ . We define the error (deviation) between the uncertain and nominal system states and inputs (2a) with  $\Delta\mathbf{x}_k := \mathbf{x}_k - \mathbf{z}_k$  and  $\Delta\mathbf{u}_k := \mathbf{u}_k - \mathbf{v}_k$ . In the following, we consider initial conditions of the form  $\Delta\mathbf{x}_0 \in \mathcal{X}_0 := P_{\text{init}}\mathcal{B}_{\infty}^n$ , with  $P_{\text{init}} \in \mathbb{R}^{n \times n}$ . This general initial set formulation will become essential for the explicit safe set (Section 4), while the online optimisation-based SL-MPSF (Section 3.2) only considers the special case of fixed initial conditions with  $P_{\text{init}} = 0_{n,n}$  and  $\Delta\mathbf{x}_0 = 0_n$ . We can write the error dynamics compactly as

$$\Delta\mathbf{x} = \mathcal{Z}A\Delta\mathbf{x} + \mathcal{Z}B\Delta\mathbf{u} + \mathcal{E}\boldsymbol{\delta}, \quad (4)$$

with  $\Delta\mathbf{x} = (\Delta\mathbf{x}_0, \dots, \Delta\mathbf{x}_N)$ ,  $\Delta\mathbf{u} = (\Delta\mathbf{u}_0, \dots, \Delta\mathbf{u}_N)$ ,  $\mathbf{w} := (\mathbf{w}_0, \dots, \mathbf{w}_{N-1})$ ,  $\boldsymbol{\delta} := (\Delta\mathbf{x}_0, \mathbf{w}) \in \mathcal{X}_0 \times \mathcal{W}^N$ ,  $A := \text{diag}(A, \dots, A, 0_{n,n})$ ,  $B := \text{diag}(B, \dots, B, 0_{n,m})$ ,  $\mathcal{E} := \text{diag}(\mathcal{I}, B_w, \dots, B_w) \in \mathcal{L}^{N, n \times n}$ , and the block downshift operator  $\mathcal{Z} \in \mathcal{L}^{N, n \times n}$ , i.e., a matrix with identity matrices along

its first block sub-diagonal and zeros elsewhere. We introduce the causal linear feedback controller  $\Delta \mathbf{u} = \mathcal{K} \Delta \mathbf{x}$ , with  $\mathcal{K} \in \mathcal{L}^{N, m \times n}$ , i.e.,

$$\mathbf{u} = \mathbf{v} + \mathcal{K}(\mathbf{x} - \mathbf{z}), \quad (5)$$

yielding  $\mathbf{u}_k = \mathbf{v}_k + \sum_{j=0}^k K^{k,j}(\mathbf{x}_{k-j} - \mathbf{z}_{k-j})$ . Using this feedback, we can write the closed-loop error dynamics as

$$\begin{bmatrix} \mathbf{x} \\ \mathbf{u} \end{bmatrix} - \begin{bmatrix} \mathbf{z} \\ \mathbf{v} \end{bmatrix} = \begin{bmatrix} \Delta \mathbf{x} \\ \Delta \mathbf{u} \end{bmatrix} = \begin{bmatrix} \mathcal{Z}(\mathcal{A} + \mathcal{B}\mathcal{K})\Delta \mathbf{x} + \mathcal{E}\delta \\ \mathcal{K}\Delta \mathbf{x} \end{bmatrix} =: \begin{bmatrix} \Phi_{\mathbf{x}} \\ \Phi_{\mathbf{u}} \end{bmatrix} \delta =: \left[ \begin{array}{c|c} \Phi_{\mathbf{x},0} & \Phi_{\mathbf{x},w} \\ \hline \Phi_{\mathbf{u},0} & \Phi_{\mathbf{u},w} \end{array} \right] \begin{bmatrix} \Delta \mathbf{x}_0 \\ \mathbf{w} \end{bmatrix}, \quad (6)$$

where  $\Phi_{\mathbf{x}} \in \mathcal{L}^{N, n \times n}$ ,  $\Phi_{\mathbf{u}} \in \mathcal{L}^{N, m \times n}$ . The matrices  $\Phi_{\mathbf{x}}$  and  $\Phi_{\mathbf{u}}$  are called the system responses from the disturbances to the closed-loop error state and input, respectively. The following proposition shows that the closed-loop responses under arbitrary affine feedbacks lie on a linear subspace.

**Proposition 2** (*Chen et al., 2021, adapted from Theorem 1*) Consider some disturbance sequence  $\delta \in \mathcal{X}_0 \times \mathcal{W}^N$ .

- a) Any trajectory  $\Delta \mathbf{x}$ ,  $\Delta \mathbf{u}$  satisfying the dynamics (4), (5) also satisfy (6) with some  $\Phi_{\mathbf{x}} \in \mathcal{L}^{N, n \times n}$ ,  $\Phi_{\mathbf{u}} \in \mathcal{L}^{N, m \times n}$  lying on the subspace

$$[\mathcal{I} - \mathcal{Z}\mathcal{A} \quad -\mathcal{Z}\mathcal{B}] \begin{bmatrix} \Phi_{\mathbf{x}} \\ \Phi_{\mathbf{u}} \end{bmatrix} = \mathcal{E}. \quad (7)$$

- b) Let  $\Phi_{\mathbf{x}}$  and  $\Phi_{\mathbf{u}}$  be arbitrary matrices satisfying (7). Then the corresponding  $\Delta \mathbf{x}$  and  $\Delta \mathbf{u}$  computed with (6) also satisfy (4), (5) with  $\mathcal{K} = \Phi_{\mathbf{u}}\Phi_{\mathbf{x}}^{-1} \in \mathcal{L}^{N, m \times n}$ .

**Remark 3** We consider square matrices  $B_w$ ,  $P_{\text{init}}$  to provide a simple exposition<sup>1</sup>. However, the results can be directly extended to more general disturbances and initial condition sets of the form  $w(k) \in \mathcal{B}_{\infty}^{n_w}$ ,  $\mathcal{X}_0 = P_0 \mathcal{B}_{\infty}^{n_0}$ , compare, e.g. *Herold et al. (2022)*.

For any nominal trajectory  $\mathbf{z}$ ,  $\mathbf{v}$  satisfying (2a), and any error feedback  $\Delta \mathbf{u} = \mathcal{K} \Delta \mathbf{x}$  with  $\Phi_{\mathbf{x}}$ ,  $\Phi_{\mathbf{u}}$  satisfying (7), the disturbance reachable sets of system (6) are given exactly by

$$\begin{aligned} \mathbf{x}_k &\in \mathcal{R}_{\mathbf{x}}(\mathbf{z}_k, \Phi_{\mathbf{x}}^k) := \{\mathbf{z}_k\} \oplus \Phi_{\mathbf{x},0}^k \mathcal{X}_0 \oplus \Phi_{\mathbf{x},w}^k \mathcal{B}_{\infty}^{Nn}, \\ \mathbf{u}_k &\in \mathcal{R}_{\mathbf{u}}(\mathbf{v}_k, \Phi_{\mathbf{u}}^k) := \{\mathbf{v}_k\} \oplus \Phi_{\mathbf{u},0}^k \mathcal{X}_0 \oplus \Phi_{\mathbf{u},w}^k \mathcal{B}_{\infty}^{Nn}. \end{aligned} \quad (8)$$

Hence, state and input constraint satisfaction can be robustly ensured by constraining the disturbance reachable sets to lie within the respective constraint sets, i.e.,  $\mathcal{R}_{\mathbf{x}}(\mathbf{z}_k, \Phi_{\mathbf{x}}^k) \subseteq \mathcal{X}$  and  $\mathcal{R}_{\mathbf{u}}(\mathbf{v}_k, \Phi_{\mathbf{u}}^k) \subseteq \mathcal{U}$  for all  $k = 0, \dots, N-1$ .

**Proposition 4** There exists an affine feedback law of the form (5) such that for any  $\delta \in \mathcal{X}_0 \times \mathcal{W}^N$ ,

$$\mathbf{x}_k \in \mathcal{X} := \{x | A_{x,j} x \leq b_{x,j}, j = 1, \dots, n_x\}, \mathbf{u}_k \in \mathcal{U} := \{u | A_{u,j} u \leq b_{u,j}, j = 1, \dots, n_u\}, \quad (9)$$

with  $\mathbf{x}_k$  and  $\mathbf{u}_k$  satisfying (1), if and only if there exist matrices  $\Phi_{\mathbf{x}}$  and  $\Phi_{\mathbf{u}}$  and a nominal trajectory  $\mathbf{z}$ ,  $\mathbf{v}$  satisfying (7), (2a), and, for  $k = 0, \dots, N-1$ ,  $j_x = 1, \dots, n_x$ ,  $j_u = 1, \dots, n_u$ :

$$\begin{aligned} A_{x,j_x} \mathbf{z}_k + \|A_{x,j_x} \Phi_{\mathbf{x},0}^k P_{\text{init}}\|_1 + \|A_{x,j_x} \Phi_{\mathbf{x},w}^k\|_1 &\leq b_{x,j_x}, \\ A_{u,j_u} \mathbf{v}_k + \|A_{u,j_u} \Phi_{\mathbf{u},0}^k P_{\text{init}}\|_1 + \|A_{u,j_u} \Phi_{\mathbf{u},w}^k\|_1 &\leq b_{u,j_u}. \end{aligned} \quad (10)$$

1. For more general disturbance and initial condition sets, Proposition 2 requires  $\mathcal{E}$  to be right-invertible, see e.g. *Herold et al. (2022)*.

**Proof** As per Proposition 2, the state and input trajectory satisfy (6). Hence, for each  $k = 0, \dots, N-1$ ,  $j = 1, \dots, n_x$ , we have

$$\begin{aligned} \max_{\substack{\mathbf{w} \in \mathcal{W}^N \\ \Delta \mathbf{x}_0 \in \mathcal{X}_0}} A_{x,j} \mathbf{x}_k &\stackrel{(6)}{=} A_{x,j} \mathbf{z}_k + \max_{\Delta \mathbf{x}_0 \in \mathcal{X}_0} A_{x,j} \Phi_{x,0}^k \Delta \mathbf{x}_0 + \max_{\mathbf{w} \in \mathcal{W}^N} A_{x,j} \Phi_{x,w}^k \mathbf{w} \\ &= A_{x,j} \mathbf{z}_k + \|A_{x,j} \Phi_{x,0}^k P_{\text{init}}\|_1 + \|A_{x,j} \Phi_{x,w}^k\|_1 \leq b_{x,j}, \end{aligned} \quad (11)$$

where the second equality is given by the definition of the 1-norm. The same derivation applies to the input constraints.  $\blacksquare$

As standard in MPC, we ensure recursive feasibility by using a suitable terminal set [Kouvaritakis and Cannon \(2016\)](#).

**Assumption 1** *There exists a terminal set  $\mathcal{X}_f := \{x \in \mathbb{R}^n \mid A_f x \leq b_f\}$  and a terminal feedback  $K_f \in \mathbb{R}^{m \times n}$  such that  $(A + BK_f)\mathcal{X}_f \oplus B_w \mathcal{W} \subseteq \mathcal{X}_f$  and  $(\mathcal{X}_f \times K_f \mathcal{X}_f) \subseteq (\mathcal{X} \times \mathcal{U})$ .*

### 3.2. System level predictive safety filter

In this section, we present the SL-MPSF using the parameterisation in Section 3.1. In particular, at each time step  $t$ , for a given control input  $u_{\mathcal{L}}(t)$  and a measured state  $x(t)$ , the SL-MPSF is defined via the following optimisation problem:

$$\min_{\Phi, \mathbf{z}, \mathbf{v}} \quad \|\mathbf{v}_0 - u_{\mathcal{L}}(t)\|_2^2, \quad (12a)$$

$$\text{s.t.} \quad [\mathcal{I} - \mathcal{Z}\mathcal{A} \quad -\mathcal{Z}\mathcal{B}] \begin{bmatrix} \Phi_x \\ \Phi_u \end{bmatrix} = \mathcal{E}, \quad x(t) = \mathbf{z}_0, \quad (12b)$$

$$\mathbf{z}_{k+1} = \mathbf{A}\mathbf{z}_k + \mathbf{B}\mathbf{v}_k, \quad k = 0, \dots, N-1, \quad (12c)$$

$$A_{x,j} \mathbf{z}_k + \|A_{x,j} \Phi_{x,w}^k\|_1 \leq b_{x,j}, \quad j = 1, \dots, n_x, \quad k = 0, \dots, N-1, \quad (12d)$$

$$A_{u,j} \mathbf{v}_k + \|A_{u,j} \Phi_{u,w}^k\|_1 \leq b_{u,j}, \quad j = 1, \dots, n_u, \quad k = 0, \dots, N-1, \quad (12e)$$

$$A_{f,j} \mathbf{z}_N + \|A_{f,j} \Phi_{x,w}^N\|_1 \leq b_{f,j}, \quad j = 1, \dots, n_f. \quad (12f)$$

The constraints (12d), (12e), and (12f) follow from Proposition 4, where we consider the special case  $P_{\text{init}} = 0_{n,n}$  with  $x(t) = \mathbf{x}_0 = \mathbf{z}_0$ . The terminal set constraint (12f) ensures (robust) recursive feasibility. The constraint (12b) parameterises the tube-controller according to Proposition 2. The constraint (12c) defines the nominal dynamics used to compute a safe backup trajectory. The cost (12a) is chosen such that the applied input  $u_{\mathcal{L}}(t)$  is modified as little as necessary, i.e., the solution of (12) is  $u_{\mathcal{L}}(t) = \mathbf{v}_0^*$  whenever feasible. Problem (12) is a quadratic program (QP) and can hence be solved efficiently. We denote a minimiser of Problem (12) by  $\mathbf{v}^*(x(t), u_{\mathcal{L}}(t))$ . In the closed-loop, we apply the first element of the optimal input sequence, i.e., the safety filter control law is given by

$$u_{\mathcal{S}}(x(t), u_{\mathcal{L}}(t), t) = \mathbf{v}_0^*(x(t), u_{\mathcal{L}}(t)). \quad (13)$$

We denote the set of feasible states  $x(t)$  for Problem (12) by  $\mathcal{X}_N \subseteq \mathcal{X}$ .

**Theorem 5** *Let Assumption 1 hold. Then, the set  $\mathcal{X}_N$  is a safe set according to Definition 1 for system (1) with the safety filter (13).*

**Proof** Suppose Problem (12) is feasible for some state  $x(t) \in \mathcal{X}_N$ . Then, for any  $u_{\mathcal{L}}(t) \in \mathbb{R}^m$  and any  $w(t) \in \mathcal{W}$ , one can show that Problem (12) is feasible for  $x(t+1)$ . Hence, the feasible set  $\mathcal{X}_N$  is robust positive invariant (RPI), i.e.,  $x(t+1) = Ax(t) + B\mathbf{v}_0^*(x(t), u_{\mathcal{L}}(t)) + w(t) \in \mathcal{X}_N$ . Following standard MPC arguments, this can be shown by constructing a feasible candidate solution that shifts the previously optimal solution and appends the terminal control law  $K_f$  from Assumption 1. In fact, this proof can be found in Goulart et al. (2006), where the result was shown based on the equivalent disturbance feedback formulation (cf. (Sieber et al., 2021, Theorem 2)). Notably, this recursive feasibility property is completely independent of the considered cost function and hence the input  $u_{\mathcal{L}}(t)$ . ■

Compared to standard RMPC formulations Mayne et al. (2005); Chisci et al. (2001), the proposed SL-MPSF formulation (12) reduces the conservativeness by optimising over affine feedbacks and using a tight reachability analysis (Proposition 4). As a result, the SL-MPSF scheme is less conservative than state-of-the-art MPSF schemes, e.g. Wabersich and Zeilinger (2018a). As online optimisation can be computationally expensive, the next section introduces an alternative approach to synthesise a safe set and safety filter, which only requires the solution to a single LP offline. This drastically reduces the online computation complexity, albeit at the expense of potentially more frequent and stronger control interventions of the safety filter.

#### 4. Explicit system level safe set

The safety filter described in Section 3 requires the solution of a QP at each time step. To eliminate the need for an embedded QP-solver, we present an explicit safe set that can be generated through the solution of a single LP, solved offline. The optimisation problem solved offline relies on the parameterisation introduced in Section 3.1, and hence is comparable to Problem (12).

The key idea is that we maximise the size of the set of initial conditions  $\mathcal{S}_e$ , such that the system is guaranteed to return to that same set  $\mathcal{S}_e$  within a given number of time steps, while always satisfying the constraints. Consequently, besides the constraint satisfaction given by Propositions 2, 4, this is realised by adding the following two constraints

$$\mathbf{x}_0 \in \mathcal{S}_e, \mathcal{R}_x(\mathbf{z}_N, \Phi_x^N) \subseteq \mathcal{S}_e. \quad (14)$$

Previous approaches to design an (explicit) safe sets (cf., e.g. Wabersich and Zeilinger (2018b)), are based on a (robust) positively invariant set and offline controller design. In contrast, using system level disturbance reachable sets, the proposed safe set does not require any controller tuning and satisfies a weaker periodic invariance condition Gondhalekar and Jones (2011), and hence allows us to optimise over a simple shape of the safe set, e.g. unit infinity-norm balls. In the following, we consider for simplicity the safe set  $\mathcal{S}_e := \{\mathbf{z}_0\} \oplus \alpha \mathcal{B}_{\infty}^n$ , with  $\alpha > 0$ , but hyperboxes, i.e.,  $\mathcal{S}_e = \{\mathbf{z}_0\} \oplus \text{diag}(\alpha) \mathcal{B}_{\infty}^n$ ,  $\text{diag}(\alpha) \in \mathbb{R}^{n \times n}$  with  $\alpha \in \mathbb{R}^n$  could be considered directly. Note that Inequalities (10) in Proposition 4 have a non-convex bilinearity if we optimise over the matrix  $P_{\text{init}}$ . Hence, we use the change of variables  $\tilde{\Phi}_x^{k,k} := \Phi_x^{k,k} P_{\text{init}}$ , and  $\tilde{\Phi}_u^{k,k} := \Phi_u^{k,k} P_{\text{init}}$  in (10) for lossless convexification. Due to the change of variables, the matrix  $\mathcal{E}$  in (7) is replaced by  $\tilde{\mathcal{E}}(\alpha) := \text{diag}(\alpha \mathcal{I}, B_w, \dots, B_w) \in \mathcal{L}^{N, n \times n}$ .



The following optimisation problem yields the proposed explicit safe set with an explicit safe backup control law:

$$\min_{\Phi, \mathbf{z}, \mathbf{v}, \alpha} -\alpha, \quad \text{s.t.} \quad [\mathcal{I} - \mathcal{Z}\mathcal{A} \quad -\mathcal{Z}\mathcal{B}] \begin{bmatrix} \Phi_{\mathbf{x}} \\ \Phi_{\mathbf{u}} \end{bmatrix} = \tilde{\mathcal{E}}(\alpha), \quad (15a)$$

$$(12b), (12d), (12e), \quad (15b)$$

$$|\mathcal{I}_{n,j}(\mathbf{z}_N - \mathbf{z}_0)| + \|\mathcal{I}_{n,j}\Phi_{\mathbf{x}}^N\|_1 \leq \alpha, \quad j = 1, \dots, n, \quad (15c)$$

where  $\mathcal{I}_{n,j}$  is the  $j^{\text{th}}$  row of  $\mathcal{I}_n$ . The constraints (15b) and (15c) are similar from Problem (12f). We denote the optimal solution of the optimisation problem (15) as  $\Phi^*$ ,  $\mathbf{z}^*$ ,  $\mathbf{v}^*$  and  $\alpha^*$ . The explicit safe set is given by  $\mathcal{S}_e^* := \{\mathbf{z}_0^*\} \oplus \alpha^* \mathcal{B}_{\infty}^n$  and  $\mathcal{K}^* = (\Phi_{\mathbf{x}}^*)^{-1} \Phi_{\mathbf{u}}^*$  characterises a safe affine backup controller, similar to (5)<sup>2</sup>. The objective function of (15) maximises the volume of the safe set. The parameterisation of the controller yields the constraints (15a) as per Proposition 2. Note that the nominal trajectory  $(\mathbf{z}, \mathbf{v})$  could also be chosen as zeros in (15), however, especially for asymmetrical constraints, this may introduce conservativeness. The constraints (15c) guarantee that the system will robustly return to the safe set within  $N$  time steps under the safe backup controller (cf. (14)). The constraints (15b) guarantee robust constraint satisfaction for any realisation of the disturbance and for any initial conditions given by (14). Algorithm 1 ensures safety by verifying that the learned-input  $u_{\mathcal{L}}(t)$  keeps the system within the safe set  $\mathcal{S}_e^*$ , which is computed using (15), before applying it. We note that the condition (16) is a set containment condition, evaluated by only a few arithmetic

**Algorithm 1:** Explicit system level safety filter

Compute  $\mathcal{S}_e^*$ ,  $\mathcal{K}^*$ ,  $\mathbf{z}^*$ , and  $\mathbf{v}^*$  using Problem (15), Initialise  $x(0) \in \mathcal{S}_e^*$

**for**  $t = 0, 1, 2, \dots$  **do**

**if**

$$u_{\mathcal{L}}(t) \in \mathcal{U} \wedge \{Ax(t) + Bu_{\mathcal{L}}(t)\} \oplus B_w \mathcal{W} \subseteq \mathcal{S}_e^* \quad (16)$$

**then**

            Apply  $u_{\mathcal{S}_e}(t) = u_{\mathcal{L}}(t)$ ,  $j \leftarrow 0$

**else**

            Apply

$$u_{\mathcal{S}_e}(t) = \mathbf{v}_j^* + \sum_{i=0}^j \mathcal{K}^{*j,i} (x(t-i:t) - \mathbf{z}_i^*), \quad j \leftarrow \text{modulo}(j+1, N) \quad (17)$$

**end**

**end**

operations thanks to the simple shape of the set  $\mathcal{S}_e^*$ .

**Theorem 6** *The control law resulting from Algorithm 1 is a safety filter and the set  $\mathcal{S}_e^*$  is the corresponding safe set according to Definition 1.*

**Proof** First, if the input  $u_{\mathcal{L}}(t)$  is safe, i.e.,  $u_{\mathcal{L}}(t) \in \mathcal{U} \wedge \{Ax(t) + Bu_{\mathcal{L}}(t)\} \oplus B_w \mathcal{W} \subseteq \mathcal{S}_e^*$ , then, the state is guaranteed to stay within the set  $\mathcal{S}_e^*$ , by definition, and hence within the state constraints. Hence, the safe backup controller can always be applied subsequently to the application of a learned input.

2. The constraint (15c) guarantees that  $\alpha^* > 0$  for  $B_w$  invertible and hence  $\tilde{\mathcal{E}}(\alpha)$  is invertible.

We now look at the case of a (potentially) unsafe input  $u_{\mathcal{L}}(t)$ . We denote the last time that the safety condition in (16) was fulfilled by  $\bar{t} \in \mathbb{N}_{\geq 0}$ . The safe backup controller constructed in (15) is a direct application of Propositions 2 and 4. Hence, starting at  $x(\bar{t}) \in \mathcal{S}_e^*$ , it results that by applying the affine feedback (17) for  $N$  time steps to the system (1), the constraints

$$x(\bar{t} + k) \in \mathcal{X}, \quad u(\bar{t} + k) \in \mathcal{U}, \quad k = 0, \dots, N - 1,$$

are robustly satisfied for any realisation of  $\delta \in \mathcal{X}_0 \times \mathcal{W}^N$ , i.e., for any  $x(\bar{t}) \in \alpha^* \mathcal{B}_{\infty}^n$  and any sequence  $(w(\bar{t}), \dots, w(\bar{t} + N - 1)) \in \mathcal{W}^N$ . Along the same lines, the constraint (15c) guarantees that the system returns to the safe set within  $N$  time steps, i.e.,  $x(\bar{t} + N) \in \mathcal{S}_e^*$ . Hence the same backup controller is still feasible by re-setting  $\bar{t} = t$ .

It results that the backup controller associated to the safe set  $\mathcal{S}_e^*$  is explicitly described by (17), and Algorithm 1 implements the associated safety filter as per Definition 1.  $\blacksquare$

**Remark 7** *Based on the solution of (15), we could also construct an explicit robust control invariant (RCI) set  $\mathcal{C}^*$  as*

$$\mathcal{C}^* := \text{Conv} \left( \bigcup_{k=0}^{N-1} \{z_k^*\} \oplus \Phi_x^{*k} \mathcal{B}_{\infty}^{Nn} \right), \quad (18)$$

where Conv denotes the convex hull. The set  $\mathcal{C}^*$  defines a safe set as per Definition 1. Similar to CBFs Cheng et al. (2019), the safety filter is implicitly defined by solving the following QP

$$\min_{u \in \mathcal{U}} \|u - u_{\mathcal{L}}(t)\|_2^2, \quad \text{s.t. } Ax(t) + Bu + B_w w \in \mathcal{C}^* \quad \forall w \in \mathcal{B}_{\infty}^n. \quad (19a)$$

## 5. Numerical example

In this section, we demonstrate the benefits of the proposed system level safety filters compared to state-of-the-art MPSF formulations specifically highlighting improvements in the size of the safe set and the maximum amount of control intervention. We consider the illustrative example of controlling a double integrator system which relates to many practical problems, e.g. a spacecraft in rotation around one of its axis Mammarella et al. (2018). The dynamics are given by  $x(t + 1) = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} x(t) + \begin{bmatrix} 0.5 \\ 1 \end{bmatrix} u(t) + \begin{bmatrix} 0.3 & 0 \\ 0 & 0.3 \end{bmatrix} w(k)$ , with  $|u(t)| \leq 3$ ,  $\|x(t)\|_{\infty} \leq 5$ ,  $\|w(t)\|_{\infty} \leq 1$ ,  $t \in \mathbb{N}_{\geq 0}$ . Consider the linear feedback  $K_f$  computed via LQR design, with  $Q = \mathcal{I}_2$ , and  $R = 10^2$ . For this system and a prediction horizon of  $N = 10$ , we compare the MPSF safe set  $\mathcal{S}_{\text{MPSF}}$  Wabersich and Zeilinger (2018a), the maximal RPI set  $\Omega_{\text{max}}$ , the SL-MPSF safe set  $\mathcal{S}_{\text{SL-MPSF}}$  (Section 3.2), the SL-based explicit safe set  $\mathcal{S}_e$  (Section 4) and the maximum RCI set  $\Xi_{\text{max}}$ . Although  $\Xi_{\text{max}}$  is not practically relevant because of its poor scalability, it serves as a basis for comparison as it is the maximum theoretically possible safe set. Using the MPT3 toolbox Herceg et al. (2013), we find approximations of a minimal RPI set  $\Omega_{\text{min}}$ , maximal RPI set  $\Omega_{\text{max}}$ , and maximal positively invariant set  $\Pi_{\text{max}}$  based on the closed-loop dynamics  $A + BK_f$ . We use the terminal set  $\mathcal{X}_f = \Omega_{\text{max}}$  for the proposed SL-MPSF. For the MPSF approach, we use  $\Omega_{\text{min}}$  for constraint tightening and  $\Pi_{\text{max}}$  as the terminal set (see e.g. Mayne et al. (2005) for further details).

In Figure 1(right), the safe sets for each method are plotted and their size is compared against the largest possible safe set  $\Xi_{\text{max}}$ , with larger sets being preferable. The SL-MPSF's safe set is

nearly maximal and significantly larger than the one associated with the MPSF, demonstrating the effectiveness of our proposed method. The explicit safe set covers regions not covered by either  $\Omega_{\max}$  or  $\mathcal{S}_{\text{MPSF}}$ . However, in general, the maximal RPI set is not contained within the explicit safe set because of the simple scalar parameterisation of the safe set.

In Figure 1(left-middle), we show the maximum intervention for any input  $u_{\mathcal{L}}(t) \in \mathcal{U}$  for each point in the safe set given by SL-MPSF and MPSF, where lower values are preferable.




Figure 1: Comparison of the maximal control intervention for MPSF (left) and SL-MPSF (middle). Size comparison for safe set from MPSF  $\mathcal{S}_{\text{MPSF}}$ , explicit  $\mathcal{S}_e$ , SL-MPSF  $\mathcal{S}_{\text{SL-MPSF}}$ , maximal RPI set  $\Omega_{\max}$ , maximal RCI set  $\Xi_{\max}$ , and constraint set  $\mathcal{X}$  (right).

The results highlight that SL-MPSF is not only able to ensure safety for a larger set of states, but it also requires a less aggressive modification of potentially unsafe inputs. This allows the learning-based controller to maintain high performance with minimal interference. SL-MPSF is hence the less conservative filter to explore the state space with minimal control interventions when the maximal RCI cannot be computed.

**Computation times** We uniformly sample  $10^4$  initial conditions in  $\mathcal{X}$ . The MPSF solve time is  $4.8\text{ms} \pm 1.8\text{ms}$ , while SL-MPSF solve time is  $65.7\text{ms} \pm 20.0\text{ms}$ , i.e., a factor  $\approx 10$  slower. In contrast, the evaluation of the explicit safety filter (see Eq. (16)) takes  $0.031\text{ms} \pm 0.028\text{ms}$ , i.e., a factor  $\approx 10^3$  faster compared to SL-MPSF. However, this code has not been optimised for speed.<sup>3</sup>

## 6. Conclusion

The paper has presented new predictive safety filters formulation based on SLS techniques, to augment any control policy (including the important case of learning-based controllers) with safety guarantees. Two distinct safety filters have been proposed leveraging the flexibility of SLS. The first method extends MPSF by enlarging the safe set through online optimisation over the robustness-ensuring tube-controller. This approach allows to verify safety for a larger set of states with less control intervention, thereby reducing the impact on the performance of the potentially unsafe controller. The second method presents an efficient explicit safety filter formulation that does not require solving optimisation problems online or tuning a controller offline. We expect that the proposed approaches can be extended to the nonlinear setting using the tools developed in Leeman et al. (2023).

3. The code is available online with permanent link <https://doi.org/10.3929/ethz-b-000611665>. We use YALMIP Löfberg (2004) and MOSEK ApS (2021) on a Intel® Core™ i7-8565U CPU @1.80GHz with 16.0GB of RAM.

## References

- Aaron D. Ames, Samuel Coogan, Magnus Egerstedt, Gennaro Notomista, Koushil Sreenath, and Paulo Tabuada. Control Barrier Functions: Theory and Applications. In *Proc. 18th European Control Conference (ECC)*, pages 3420–3431, 2019. doi: 10.23919/ECC.2019.8796030.
- James Anderson, John C. Doyle, Steven H. Low, and Nikolai Matni. System level synthesis. *Annual Reviews in Control*, 47:364–393, 2019. ISSN 13675788. doi: 10.1016/j.arcontrol.2019.03.006.
- Mosek ApS. *The MOSEK optimization software*, 2021. URL <https://www.mosek.com>.
- Shaoru Chen, Nikolai Matni, Manfred Morari, and Victor M. Preciado. System Level Synthesis-based Robust Model Predictive Control through Convex Inner Approximation. *arXiv preprint arXiv:2111.05509*, 2021.
- Richard Cheng, Gábor Orosz, Richard M. Murray, and Joel W. Burdick. End-to-End Safe Reinforcement Learning through Barrier Functions for Safety-Critical Continuous Control Tasks. In *AAAI Conference on Artificial Intelligence*, volume 33, pages 3387–3395, 2019. doi: 10.1609/aaai.v33i01.33013387.
- Luigi Chisci, J. Anthony Rossiter, and Giovanni Zappa. Systems with persistent disturbances: Predictive control with restricted constraints. *Automatica*, 37(7):1019–1028, 2001. ISSN 00051098. doi: 10.1016/S0005-1098(01)00051-6.
- Hongkai Dai, Benoit Landry, Lujie Yang, Marco Pavone, and Russ Tedrake. Lyapunov-stable neural-network control. *arXiv preprint arXiv:2109.14152*, 2021.
- Alexandre Didier, Kim P. Wabersich, and Melanie N. Zeilinger. Adaptive Model Predictive Safety Certification for Learning-based Control. In *Proc. 60th IEEE Conference on Decision and Control (CDC)*, pages 809–815. IEEE, 2021. ISBN 9781665436595. doi: 10.1109/CDC45484.2021.9682832.
- Jaime F. Fisac, Anayo K. Akametalu, Melanie N. Zeilinger, Shahab Kaynama, Jeremy Gillula, and Claire J. Tomlin. A General Safety Framework for Learning-Based Control in Uncertain Robotic Systems. *IEEE Transactions on Automatic Control*, 64(7):2737–2752, 2019. ISSN 15582523. doi: 10.1109/TAC.2018.2876389.
- Jeremy H. Gillula and Claire J. Tomlin. Guaranteed Safe Online Learning via Reachability: tracking a ground target using a quadrotor. In *Proc. IEEE International Conference on Robotics and Automation (ICRA)*, pages 2723–2730, 2012. doi: 10.1109/ICRA.2012.6225136.
- Ravi Gondhalekar and Colin N. Jones. MPC of constrained discrete-time linear periodic systems — A framework for asynchronous control: Strong feasibility, stability and optimality via periodic invariance. *Automatica*, 47(2):326–333, 2011. ISSN 0005-1098. doi: <https://doi.org/10.1016/j.automatica.2010.10.021>.
- Paul J. Goulart, Eric C. Kerrigan, and Jan M. Maciejowski. Optimization over state feedback policies for robust control with constraints. *Automatica*, 42(4):523–533, 2006. ISSN 00051098. doi: 10.1016/j.automatica.2005.08.023.

- M. Herceg, M. Kvasnica, Colin N. Jones, and M. Morari. Multi-Parametric Toolbox 3.0. In *Proc. European Control Conference (ECC)*, pages 502–510, Zürich, Switzerland, July 17–19 2013. <http://control.ee.ethz.ch/~mpt>.
- Thilo Herold, Felix Berkel, Adrian Trachte, and Thomas Specker. A Computationally Efficient System Level Parametrization for Robust Model Predictive Control. In *Proc. European Control Conference (ECC)*, pages 826–832. EUCA, 2022. ISBN 9783907144077. doi: 10.23919/ECC55457.2022.9838178.
- Jemin Hwangbo, Joonho Lee, Alexey Dosovitskiy, Dario Bellicoso, Vassilios Tsounis, Vladlen Koltun, and Marco Hutter. Learning agile and dynamic motor skills for legged robots. *Science Robotics*, 4(26):eaau5872, 2019. doi: 10.1126/scirobotics.aau5872.
- Julian Ibarz, Jie Tan, Chelsea Finn, Mrinal Kalakrishnan, Peter Pastor, and Sergey Levine. How to train your robot with deep reinforcement learning: lessons we have learned. *The International Journal of Robotics Research*, 40(4-5):698–721, 2021. doi: 10.1177/0278364920987859.
- Basil Kouvaritakis and Mark Cannon. *Model Predictive Control: Classical, Robust, and Stochastic*. Springer, 2016. ISBN 9783319248516. doi: 10.1109/mcs.2016.2602738.
- Antoine P. Leeman, Johannes Köhler, Andrea Zanelli, Samir Bennani, and Melanie N Zeilinger. Robust nonlinear optimal control via system level synthesis. *arXiv preprint arXiv:2301.04943*, 2023.
- Johan Löfberg. Yalmip : A toolbox for modeling and optimization in matlab. In *Proc. CACSD Conference*, Taipei, Taiwan, 2004.
- Martina Mammarella, Elisa Capello, Hyeongjun Park, Giorgio Guglieri, and Marcello Romano. Tube-based robust model predictive control for spacecraft proximity operations in the presence of persistent disturbance. *Aerospace Science and Technology*, 77:585–594, 2018. ISSN 12709638. doi: 10.1016/j.ast.2018.04.009.
- David Q. Mayne, María M. Seron, and SV Raković. Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica*, 41(2):219–224, 2005. ISSN 00051098. doi: 10.1016/j.automatica.2004.08.019.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Belle-mare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, and others. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- Simon Muntwiler, Kim P. Wabersich, Andrea Carron, and Melanie N. Zeilinger. Distributed Model Predictive Safety Certification for Learning-based Control. *IFAC-PapersOnLine*, 53(2):5258–5265, 2020. ISSN 2405-8963. doi: <https://doi.org/10.1016/j.ifacol.2020.12.1205>.
- Anilkumar Parsi, Panagiotis Anagnostaras, Andrea Iannelli, and Roy S. Smith. Computationally efficient robust MPC using optimized constraint tightening. In *Proc. IEEE 61st Conference on Decision and Control (CDC)*, pages 1770–1775. IEEE, 2022.

- Jerome Sieber, Samir Bennani, and Melanie N. Zeilinger. A System Level Approach to Tube-based Model Predictive Control. *IEEE Control Systems Letters*, 6:776–781, 2021. ISSN 24751456. doi: 10.1109/LCSYS.2021.3086190.
- Jerome Sieber, Andrea Zanelli, Samir Bennani, and Melanie N. Zeilinger. System Level Disturbance Reachable Sets and their Application to Tube-based MPC. *European Journal of Control*, 2022. ISSN 0947-3580. doi: <https://doi.org/10.1016/j.ejcon.2022.100680>.
- Kim P. Wabersich and Melanie N. Zeilinger. Linear model predictive safety certification for learning-based control. In *Proc. Conference on Decision and Control (CDC)*, pages 809–815. IEEE, 2018a. ISBN 9781665436595. doi: 10.1109/CDC45484.2021.9682832.
- Kim P. Wabersich and Melanie N. Zeilinger. Scalable synthesis of safety certificates from data with application to learning-based control. In *Proc. European Control Conference (ECC)*, pages 1691–1697. European Control Association (EUCA), 2018b. ISBN 9783952426982. doi: 10.23919/ECC.2018.8550288.
- Kim P. Wabersich and Melanie N. Zeilinger. A predictive safety filter for learning-based control of constrained nonlinear dynamical systems. *Automatica*, 129:109597, 2021. ISSN 00051098. doi: 10.1016/j.automatica.2021.109597.
- Kim P. Wabersich, Lukas Hewing, Andrea Carron, and Melanie N. Zeilinger. Probabilistic Model Predictive Safety Certification for Learning-Based Control. *IEEE Transactions on Automatic Control*, 67(1):176–188, 2022. doi: 10.1109/TAC.2021.3049335.