Available online at www.sciencedirect.com
ScienceDirect
IFAC PapersOnLine 54-6 (2021) 257–263

On the design of terminal ingredients for data-driven MPC

Julian Berberich * Johannes Köhler * Matthias A. Müller ** Frank Allgöwer *

* University of Stuttgart, Institute for Systems Theory and Automatic Control, 70550 Stuttgart, Germany (email:{julian.berberich, johannes.koehler, frank.allgower}@ist.uni-stuttgart.de).
** Leibniz University Hannover, Institute of Automatic Control, 30167 Hannover, Germany (e-mail:mueller@irt.uni-hannover.de).

Abstract: We present a model predictive control (MPC) scheme to control linear time-invariant systems using only measured input-output data and no model knowledge. The scheme includes a terminal cost and a terminal set constraint on an extended state containing past input-output values. We provide an explicit design procedure for the corresponding terminal ingredients that only uses measured input-output data. Further, we prove that the MPC scheme based on these terminal ingredients exponentially stabilizes the desired setpoint in closed loop. Finally, we illustrate the advantages over existing data-driven MPC approaches with a numerical example.

1. INTRODUCTION

Designing model predictive control (MPC) schemes based on measured data with only partial or no model knowledge is an active field of research using, e.g., model adaptation (Adetola and Guay, 2011; Aswani et al., 2013; Tanaskovic et al., 2014) or reinforcement learning based approaches (Berkenkamp et al., 2017; Zanon and Gros, 2020). A key motivation for such approaches is that obtaining accurate model knowledge can be difficult in practice, whereas large amounts of data are often available and can be exploited for control. In this paper, we investigate MPC of unknown linear time-invariant (LTI) systems using only measured input-output data and no model knowledge. Our approach relies on a result by Willems et al. (2005), which shows that all trajectories of an LTI system can be parametrized via one persistently exciting data trajectory. This fact has been used, e.g., for data-driven simulation by Markovsky and Rapisarda (2008) and, subsequently, for data-driven MPC by Yang and Li (2015); Coulson et al. (2019). While Coulson et al. (2020) derive open-loop robustness guarantees of the scheme, Berberich et al. (2021a) prove closed-loop stability and robustness properties, even if the measured data are affected by noise. Further theoretical results on robust constraint satisfaction based on noisy data and on a tracking MPC formulation are derived by Berberich et al. (2020b) and Berberich et al. (2020a), respectively. Notably, the existing works with closed-loop guarantees by Berberich et al. (2021a, 2020a,b) require terminal equality constraints, which may result in a small region of attraction and a small robustness margin.

In this paper, we propose a novel data-driven MPC scheme with terminal cost and terminal constraints, thereby improving robustness and increasing the region of attraction of the closed loop, in analogy to results from model-based MPC (Chen and Allgöwer (1998); Rawlings et al. (2020)). Similar to other works on data-driven MPC based on Willems et al. (2005), we assume that only input-output data of the unknown system are available and thus, we employ an extended state vector involving consecutive input-output measurements. We prove that the MPC scheme with suitable terminal ingredients involving this extended state exponentially stabilizes the closed loop. Further, using recent results on data-driven control by Berberich et al. (2020c), we show how the terminal ingredients can be constructed using only measured data. A key benefit of the presented approach is that closed-loop stability guarantees can be given in an end-to-end fashion, without prior system identification steps. Throughout this paper, we assume that the measured data are not affected by noise, but we conjecture that practical stability using noisy measurements can be proven based on recent robustness results for data-driven MPC in Berberich et al. (2021b).

The work by Dutta et al. (2014) is conceptually related to our results since it provides a terminal set constraint for (model-based) MPC with input-output models using an implicit characterization of the maximal invariant set. Moreover, Abbas et al. (2016) propose a linear matrix inequality (LMI) based procedure to construct stabilizing terminal ingredients for linear parameter-varying input-output models. Compared to these model-based results, the proposed MPC scheme and the computation of the terminal ingredients require only one input-output trajectory of an LTI system and no model knowledge.

The remainder of the paper is structured as follows. After introducing required preliminaries in Section 2, we present the data-driven MPC scheme, prove closed-loop stability, and provide a design procedure for terminal ingredients in

* This work was funded by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2075 - 390740016. The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting Julian Berberich, and the International Research Training Group Soft Tissue Robotics (GRK 2198/1 - 277536708).

10.1016/j.ifacol.2021.08.554

Section 3. The approach is applied to a numerical example in Section 4, and the paper is concluded in Section 5.

Notation We denote the set of integers in the interval $[a, b]$ by $\mathbb{I}_{[a,b]}$ and the set of nonnegative integers by $\mathbb{I}_{\geq 0}$. Moreover, we write $\|x\|_2$ and $\|A\|_2$ for the (induced) 2-norm of some vector x and matrix A. For matrices $P = P^\top$, $P_2 = P_2^\top$, we write $\lambda_{\min}(P)$ for the minimum eigenvalue of P and we define $\lambda_{\min}(P, P_2) := \min\{\lambda_{\min}(P), \lambda_{\min}(P_2)\}$, and similarly for $\lambda_{\max}(P)$ and $\lambda_{\max}(P, P_2)$. We write $P \succ 0$ ($P \succeq 0$) if P is positive (semi-) definite, and similarly $P \prec 0$ ($P \preceq 0$) if P is negative (semi-) definite. Further, we define $\|x\|_P^2 := x^\top P x$, we denote matrix entries which can be inferred from symmetry by $\star$, and we write I for an identity matrix of appropriate dimension. For some generic sequence $\{x_k\}_{k=0}^{N-1}$, we introduce $x_{[a,b]} := \begin{bmatrix} x_a^\top & \dots & x_b^\top \end{bmatrix}^\top$ and we abbreviate $x := x_{[0,N-1]}$. Finally, we define the Hankel matrix

$$H_L(x) := \begin{bmatrix} x_0 & x_1 & \dots & x_{N-L} \\ x_1 & x_2 & \dots & x_{N-L+1} \\ \vdots & \vdots & \vdots & \vdots \\ x_{L-1} & x_L & \dots & x_{N-1} \end{bmatrix}.$$

2. PRELIMINARIES

In this section, we introduce preliminaries regarding the data-driven system representation (Section 2.1) and the extended state-space description (Section 2.2). Further, we describe the problem setting in Section 2.3.

2.1 Data-driven system parametrization

We consider discrete-time LTI systems of the form

$$\begin{aligned} x_{k+1} &= Ax_k + Bu_k, \\ y_k &= Cx_k + Du_k \end{aligned} \tag{1}$$

with state $x_k \in \mathbb{R}^n$, input $u_k \in \mathbb{R}^m$, and output $y_k \in \mathbb{R}^p$, all at time $k \in \mathbb{I}_{\geq 0}$, where the matrices (A, B, C, D) are unknown. Throughout this paper, we assume that (A, B) is controllable and (A, C) is observable. Controllability is required both for applying the main result of Willems et al. (2005) and for the presented design of terminal ingredients, whereas observability is w.l.o.g. since unobservable modes do not affect the output cost or constraints.

Definition 1. We say that $\{u_k\}_{k=0}^{N-1}$ with $u_k \in \mathbb{R}^m$ is persistently exciting of order L if $\text{rank}(H_L(u)) = mL$.

The following main result of Willems et al. (2005) states that a given input-output trajectory $\{u_k^d, y_k^d\}_{k=0}^{N-1}$ can be used to parametrize all system trajectories if the input u^d is persistently exciting (compare Berberich and Allgöwer (2020) for a description in the state-space framework).

Theorem 2. (Willems et al. (2005)) Suppose $\{u_k^d, y_k^d\}_{k=0}^{N-1}$ is a trajectory of (1), where u^d is persistently exciting of order $L+n$. Then, $\{\bar{u}_k, \bar{y}_k\}_{k=0}^{L-1}$ is a trajectory of (1) if and only if there exists $\alpha \in \mathbb{R}^{N-L+1}$ such that

$$\begin{bmatrix} H_L(u^d) \\ H_L(y^d) \end{bmatrix} \alpha = \begin{bmatrix} \bar{u} \\ \bar{y} \end{bmatrix}. \tag{2}$$

Theorem 2 provides a direct non-parametric system description without identifying a model of the system. Thus, it can be used to predict system trajectories, which we exploit for the MPC scheme presented in this paper.

2.2 Extended state representation

We define the lag of the system (1) as follows.

Definition 3. The lag $\underline{l}$ of (1) is the smallest $l \in \mathbb{I}_{[1,n]}$ such that the following observability matrix has rank n:

$$\Phi_l := \begin{bmatrix} C^\top & (CA)^\top & \dots & (CA^{l-1})^\top \end{bmatrix}^\top. \tag{3}$$

For some integer l, we define the extended state ξ_t as

$$\xi_t := \begin{bmatrix} u_{[t-l,t-1]} \\ y_{[t-l,t-1]} \end{bmatrix}, \tag{4}$$

where $t \geq l$. We abbreviate the dimension of ξ_t by $n_\xi := (m+p)l$ and we denote the discrete-time LTI dynamics corresponding to the extended state by

$$\begin{aligned} \xi_{k+1} &= \tilde{A}\xi_k + \tilde{B}u_k, \\ y_k &= \tilde{C}\xi_k + \tilde{D}u_k \end{aligned} \tag{5}$$

for suitable matrices $\tilde{A}$, $\tilde{B}$, $\tilde{C}$, $\tilde{D}$. To be precise, the matrices in (5) take the form shown in (6), where the matrices F_i, G_i, $i \in \mathbb{I}_{[1,l]}$, D are unknown. It is straightforward to show that (1) and (5) have an equivalent input-output behavior if $l \geq \underline{l}$ (compare Goodwin and Sin (2014); Koch et al. (2020)). Throughout this paper, we assume that an upper bound l on the lag $\underline{l}$ is available. The proposed design method (Section 3.3) further requires $pl = n$, which is only valid if $\underline{l}$ is known exactly. Finally, we define the matrix T_y such that $y_{t-1} = T_y \xi_t$, i.e., $T_y = [0 \ \dots \ 0 \ I]$.

2.3 Problem setting

Throughout this paper, the matrices A, B, C, D, or equivalently, F_i, G_i, D, are unknown and one input-output trajectory $\{u_k^d, y_k^d\}_{k=0}^{N-1}$ of (1) is available. We consider pointwise-in-time constraints on the input and output, i.e., $u_t \in \mathbb{U} \subseteq \mathbb{R}^m$ and $y_t \in \mathbb{Y} \subseteq \mathbb{R}^p$ for $t \in \mathbb{I}_{\geq 0}$. Our objective is stabilization of a given input-output setpoint (u^s, y^s) which is an equilibrium in the following sense.

Definition 4. We say that $(u^s, y^s) \in \mathbb{R}^{m+p}$ is an equilibrium of (1), if the sequence $\{\bar{u}_k, \bar{y}_k\}_{k=0}^{\underline{l}}$ with $(\bar{u}_k, \bar{y}_k) = (u^s, y^s)$ for all $k \in \mathbb{I}_{[0,\underline{l}]}$ is a trajectory of (1).

We assume $(u^s, y^s) \in \text{int}(\mathbb{U} \times \mathbb{Y})$ and we denote the corresponding steady-state of (5) by ξ^s. While assuming knowledge of whether an input-output setpoint $(u^s, y^s) \neq 0$ is an equilibrium of the system (1) can be restrictive since A, B, C, D are unknown, this condition can be relaxed by introducing artificial setpoints in the MPC scheme which are optimized online (cf. Limón et al. (2008)).

Our contribution can be summarized as follows. We propose a data-driven MPC approach to control (1) with closed-loop stability guarantees. The MPC scheme uses (2) to predict future trajectories based on one persistently exciting input-output trajectory $\{u_k^d, y_k^d\}_{k=0}^{N-1}$. In contrast to earlier works, e.g., by Coulson et al. (2019); Berberich et al. (2021a), we ensure closed-loop stability via terminal ingredients, i.e., via an appropriate terminal set constraint and a terminal cost. This follows standard arguments from model-based MPC (Chen and Allgöwer, 1998; Rawlings

$$\begin{bmatrix} u_{k-l+1} \\ \vdots \\ u_{k-1} \\ u_k \\ \hline y_{k-l+1} \\ \vdots \\ y_{k-1} \\ y_k \end{bmatrix} = \left[\begin{array}{cccc|cccc} 0 & I & \dots & 0 & 0 & \dots & \dots & 0 \\ \vdots & \ddots & \ddots & \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & \ddots & \ddots & I & \vdots & \ddots & \ddots & \vdots \\ 0 & \dots & \dots & 0 & 0 & \dots & \dots & 0 \\ \hline 0 & \dots & \dots & 0 & 0 & I & \dots & 0 \\ \vdots & \ddots & \ddots & \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & \dots & \dots & 0 & 0 & \dots & \dots & I \\ G_l & \dots & \dots & G_1 & F_l & \dots & \dots & F_1 \end{array}\right] \begin{bmatrix} u_{k-l} \\ \vdots \\ u_{k-2} \\ u_{k-1} \\ \hline y_{k-l} \\ \vdots \\ y_{k-2} \\ y_{k-1} \end{bmatrix} + \begin{bmatrix} 0 \\ \vdots \\ 0 \\ I \\ \hline 0 \\ \vdots \\ 0 \\ D \end{bmatrix} u_k \quad (6)$$

et al., 2020). The key difficulty is that, since we do not have access to the state x_t or the matrices A, B, C, D in (1), the terminal ingredients need to be designed for the extended state ξ_t in (5) using only measured data. Based on recent results on data-driven control by Berberich et al. (2020c), we compute terminal ingredients which can be used to implement the MPC scheme. Finally, we illustrate the advantages w.r.t. earlier works with a numerical example.

3. DATA-DRIVEN MPC WITH TERMINAL INGREDIENTS

In this section, we present a data-driven MPC scheme with terminal ingredients and closed-loop stability guarantees. After defining the MPC scheme in Section 3.1, we prove the desired closed-loop properties in Section 3.2. Further, in Section 3.3, we show how the corresponding terminal ingredients can be computed using only measured data.

3.1 Proposed MPC scheme

Given an input-output data trajectory $\{u_k^d, y_k^d\}_{k=0}^{N-1}$ and initial conditions $\{u_k, y_k\}_{k=t-l}^{t-1}$, the proposed MPC scheme relies on the following open-loop optimal control problem

$$\min_{\substack{\alpha(t)\\ \bar{u}(t),\bar{y}(t)}} \sum_{k=0}^{L-1} \|\bar{u}_k(t) - u^s\|_R^2 + \|\bar{y}_k(t) - y^s\|_Q^2 \quad (7a)$$
$$+ \|\bar{\xi}_L(t) - \xi^s\|_P^2$$
$$\text{s.t.} \quad \begin{bmatrix} \bar{u}_{[-l,L-1]}(t) \\ \bar{y}_{[-l,L-1]}(t) \end{bmatrix} = \begin{bmatrix} H_{L+l}(u^d) \\ H_{L+l}(y^d) \end{bmatrix} \alpha(t), \quad (7b)$$
$$\begin{bmatrix} \bar{u}_{[-l,-1]}(t) \\ \bar{y}_{[-l,-1]}(t) \end{bmatrix} = \begin{bmatrix} u_{[t-l,t-1]} \\ y_{[t-l,t-1]} \end{bmatrix}, \quad (7c)$$
$$\bar{\xi}_L(t) \in \Xi_f,\ \bar{\xi}_L(t) = \begin{bmatrix} \bar{u}_{[L-l,L-1]}(t) \\ \bar{y}_{[L-l,L-1]}(t) \end{bmatrix}, \quad (7d)$$
$$\bar{u}_k(t) \in \mathbb{U},\ \ \bar{y}_k(t) \in \mathbb{Y},\ \ k \in \mathbb{I}_{[0,L-1]}. \quad (7e)$$

In the cost (7a), we penalize the distance of the predicted variables w.r.t. the desired setpoint (u^s, y^s), where $Q \succ 0$, $R \succ 0$ are design parameters. The constraint (7b) parametrizes all input-output trajectories $\bar{u}(t)$, $\bar{y}(t)$ of (5) using Theorem 2. Moreover, the predicted trajectory is initialized using past l input-output measurements in (7c), and the scheme includes pointwise-in-time constraints on the predicted input and output in (7e). Finally, the matrix $P \succ 0$ and the set Ξ_f characterize a suitable terminal cost $\|\bar{\xi}_L(t) - \xi^s\|_P^2$ and terminal constraint $\bar{\xi}_L(t) \in \Xi_f$ on the predicted extended state at time L. We assume that these terminal ingredients satisfy the following assumption.

Assumption 5. There exist matrices $P = P^\top \succ 0$, $K \in \mathbb{R}^{m\times n_\xi}$, and a set $\Xi_f \subseteq \mathbb{U}^l \times \mathbb{Y}^l$ with $\xi^s \in \text{int}(\Xi_f)$ such that for all $\xi \in \Xi_f$, $u = u^s + K(\xi - \xi^s)$, and $y = (\tilde{C} + \tilde{D}K)\xi$, we have

i) $\tilde{A}\xi + \tilde{B}u \in \Xi_f$,
ii) $u \in \mathbb{U}$ and $y \in \mathbb{Y}$,
iii) the following inequality holds:

$$\|(\tilde{A} + \tilde{B}K)\xi\|_P^2 \le \|\xi\|_P^2 - \|\xi\|_{K^\top RK} - \|y\|_Q^2. \quad (8)$$

Assumption 5 is a standard condition in model-based MPC to ensure closed-loop stability, compare Chen and Allgöwer (1998); Rawlings et al. (2020). In Section 3.3, we show how a matrix P and a set Ξ_f as in Assumption 5 can be constructed using only measured data and no model knowledge. If $\mathbb{U}$, $\mathbb{Y}$, and Ξ_f are convex polytopes (ellipsoids), then Problem (7) is a convex (quadratically constrained) quadratic program which can be solved efficiently. Throughout this paper, $\bar{u}^*(t)$, $\bar{y}^*(t)$, $\alpha^*(t)$ denote the optimal solution of Problem (7) at time $t \in \mathbb{I}_{\ge 0}$, whereas u_t, y_t denote the closed-loop input and output at time $t \in \mathbb{I}_{\ge 0}$. Further, we write $J_L^*(\xi_t)$ for the optimal cost of Problem (7) with initial condition $\xi_t = \begin{bmatrix} u_{[t-l,t-1]}^\top, y_{[t-l,t-1]}^\top \end{bmatrix}^\top$. Algorithm 6 summarizes the MPC scheme based on repeatedly solving Problem (7).

Algorithm 6. Data-Driven MPC Scheme

(1) At time t, take the past l measurements $u_{[t-l,t-1]}$, $y_{[t-l,t-1]}$ and solve (7).
(2) Apply the input $u_t = \bar{u}_0^*(t)$.
(3) Set $t = t + 1$ and go back to (1).

3.2 Closed-loop guarantees

We make the following common assumption.

Assumption 7. The optimal value function is quadratically upper bounded, i.e., there exists $c_u > 0$ such that $J_L^*(\xi) \le c_u\|\xi\|_2^2$ for all ξ such that Problem (7) is feasible.

Assumption 7 is, e.g., satisfied if the sets $\mathbb{U}$, $\mathbb{Y}$, Ξ_f are polytopes (Bemporad et al., 2002) or if $\mathbb{U}$ and $\mathbb{Y}$ are compact, compare (Rawlings et al., 2020, Proposition 2.16). More generally, Assumption 7 holds with a non-quadratic upper bound, thus leading only to *asymptotic* stability guarantees for the closed loop, if $\mathbb{U}$ is compact (Rawlings et al., 2020, Proposition 2.16).

Theorem 8. Suppose Assumptions 5 and 7 hold and the input u^d is persistently exciting of order $L + l + n$. If Problem (7) is feasible at initial time $t = 0$, then

i) it is feasible at any $t \in \mathbb{I}_{\ge 0}$,
ii) the closed-loop trajectory satisfies the constraints, i.e., $u_t \in \mathbb{U}$ and $y_t \in \mathbb{Y}$ for all $t \in \mathbb{I}_{\ge 0}$,

iii) the equilibrium ξ^s is exponentially stable for the resulting closed loop.

Proof. W.l.o.g., we assume $u^s = 0$, $y^s = 0$, $\xi^s = 0$. At time $t+1$, we construct a feasible candidate solution $(\bar{u}'(t+1), \bar{y}'(t+1))$ of (7) by shifting the previously optimal solution $\bar{u}^*(t)$, $\bar{y}^*(t)$ of (7) and appending the input $\bar{u}'_{L-1}(t+1) = K \begin{bmatrix} \bar{u}^*_{[L-l,L-1]}(t) \\ \bar{y}^*_{[L-l,L-1]}(t) \end{bmatrix}$, with K as in Assumption 5. Since $(\bar{u}'(t+1), \bar{y}'(t+1))$ is a trajectory of (1), there exists $\alpha'(t+1)$ satisfying (7b) by Theorem 2, i.e., all constraints of (7) are satisfied. Denoting the cost of (7) for this candidate solution by $J'_L(\xi_{t+1})$, we have

$$\begin{aligned} &J^*_L(\xi_{t+1}) - J^*_L(\xi_t) \leq J'_L(\xi_{t+1}) - J^*_L(\xi_t) \qquad (9)\\ =&\|\bar{u}'_{L-1}(t+1)\|_R^2 + \|\bar{y}'_{L-1}(t+1)\|_Q^2 - \|\bar{u}^*_0(t)\|_R^2 - \|\bar{y}^*_0(t)\|_Q^2 \\ &+ \|\bar{\xi}'_L(t+1)\|_P^2 - \|\bar{\xi}^*_L(t)\|_P^2 \overset{(8)}{\leq} -\|u_t\|_R^2 - \|y_t\|_Q^2. \end{aligned}$$

Since (A, C) is observable, the pair $(\tilde{A}, \tilde{C})$ is detectable and hence, there exists an input-output-to-state stability (IOSS) Lyapunov function $W(\xi) = \xi^\top P_W \xi$ satisfying

$$W(\tilde{A}\xi + \tilde{B}u) - W(\xi) \leq -\frac{1}{2}\|\xi\|_2^2 + c_1\|u\|_2^2 + c_2\|y\|_2^2 \quad (10)$$

for suitable $c_1, c_2 > 0$, $P_W \succ 0$, and all $u \in \mathbb{R}^m$, $\xi \in \mathbb{R}^{n_\xi}$, $y = \tilde{C}\xi + \tilde{D}u$, compare Cai and Teel (2008). We consider as a Lyapunov function candidate the weighted sum of J^*_L and W, i.e., $V(\xi_t) = J^*_L(\xi_t) + \gamma W(\xi_t)$ with $\gamma = \frac{\lambda_{\min}(Q,R)}{\max\{c_1,c_2\}} > 0$. Combining (9) and (10), this implies $V(\xi_{t+1}) - V(\xi_t) \leq -\frac{\gamma}{2}\|\xi_t\|_2^2$. The Lyapunov function candidate $V(\xi_t)$ is clearly quadratically lower bounded and it is quadratically upper bounded due to Assumption 7 such that ξ^s is exponentially stable in closed loop due to standard Lyapunov arguments (Rawlings et al., 2020). □

Theorem 8 shows that the proposed MPC scheme based on Problem (7) exponentially stabilizes the desired setpoint ξ^s and hence, the closed-loop input/output exponentially converges to u^s/y^s. Since the cost of Problem (7) is only positive semidefinite in the minimal state x, the proof relies on detectability to construct a Lyapunov function based on IOSS arguments, compare also (Rawlings et al., 2020, Theorem 2.24), Berberich et al. (2021a); Grimm et al. (2005). Note that Theorem 8 requires persistence of excitation of order $L+l+n$ for the input since the length of the constructed input-output trajectory is $L+l$ due to the initial conditions in (7c).

Remark 9. We note that, if the terminal set constraint in (7d) is dropped, i.e., $\Xi_f = \mathbb{R}^{n_\xi}$, closed-loop stability as shown in Theorem 8 holds locally with some region of attraction as long as the terminal cost matrix P satisfies the conditions in Assumption 5, compare Limón et al. (2006). For the special case of open-loop stable systems with no output constraints, i.e., $\mathbb{Y} = \mathbb{R}^p$, Assumption 5 can always be satisfied with $K = 0$ and without the terminal set constraint in (7d), i.e., $\Xi_f = \mathbb{R}^{n_\xi}$, compare Rawlings and Muske (1993). In this case, Problem (7) is globally feasible and Theorem 8 ensures global exponential stability.

3.3 Data-driven design of terminal ingredients

In this section, we show how terminal ingredients satisfying Assumption 5 can be computed using only measured data and no model knowledge, based on recent results on data-driven controller design by Berberich et al. (2020c). To this end, we write (5) as a linear fractional transformation

$$\begin{bmatrix} \xi_{k+1} \\ \hline z_k \end{bmatrix} = \left[\begin{array}{c|cc} A' & B_w & B' \\ \hline \begin{bmatrix} I \\ 0 \end{bmatrix} & 0 & \begin{bmatrix} 0 \\ I \end{bmatrix} \end{array}\right] \begin{bmatrix} \xi_k \\ \hline w_k \\ u_k \end{bmatrix}, \qquad (11)$$
$$w_k = \Delta z_k,$$

where $w \mapsto z$ represents an additional uncertainty channel capturing all unknown elements of (5) (compare Zhou et al. (1996)). More precisely, A', B', $B_w = [0 \ I]^\top$ are known matrices according to the structure in (6), and

$$\Delta = [G_l \ \dots \ G_1 \ F_l \ \dots \ F_1 \ D]$$

contains the unknown system parameters. We factorize $T_y^\top Q T_y = Q_r^\top Q_r$, $R = R_r^\top R_r$ with T_y as in Section 2.2, and we define the data-dependent matrices

$$\Xi := \begin{bmatrix} \xi_l^d \ \xi_{l+1}^d \ \dots \ \xi_{N-1}^d \end{bmatrix}, \ \Xi_+ := \begin{bmatrix} \xi_{l+1}^d \ \xi_{l+2}^d \ \dots \ \xi_N^d \end{bmatrix},$$
$$U := \begin{bmatrix} u_l^d \ u_{l+1}^d \ \dots \ u_{N-1}^d \end{bmatrix}, \ Z := \begin{bmatrix} \Xi \\ U \end{bmatrix},$$

where $\{\xi_k^d\}_{k=l}^{N-1}$ is the extended state trajectory corresponding to the available input-output measurements $\{u_k^d, y_k^d\}_{k=0}^{N-1}$. Using that $B_w^\top$ is the Moore-Penrose inverse of B_w, (Berberich et al., 2020c, Proposition 1) implies the following data-dependent bound on the "uncertainty" Δ:

$$\begin{bmatrix} \Delta^\top \\ I \end{bmatrix}^\top \underbrace{\begin{bmatrix} -ZZ^\top & Z\Xi_+^\top B_w \\ B_w^\top \Xi_+ Z^\top & -B_w^\top \Xi_+ \Xi_+^\top B_w \end{bmatrix}}_{P_\Delta^w :=} \begin{bmatrix} \Delta^\top \\ I \end{bmatrix} \succeq 0. \qquad (12)$$

Moreover, we abbreviate

$$\bar{P}_\Delta^w := \begin{bmatrix} 0 & I \\ B_w^\top & 0 \end{bmatrix}^\top P_\Delta^w \begin{bmatrix} 0 & I \\ B_w^\top & 0 \end{bmatrix}.$$

Proposition 10. Suppose there exist $\mathcal{X} \succ 0$, $\Gamma \succ 0$, M, $\tau \geq 0$, $\gamma > 0$ such that trace$(\Gamma) < \gamma^2$, $\begin{bmatrix} \Gamma & I \\ I & \mathcal{X} \end{bmatrix} \succ 0$, and

$$\begin{bmatrix} \left(\tau \bar{P}_\Delta^w - \begin{bmatrix} \mathcal{X} & 0 \\ 0 & 0 \end{bmatrix}\right) & \begin{bmatrix} A'\mathcal{X} + B'M \\ \mathcal{X} \\ M \end{bmatrix} & 0 \\ \star & -\mathcal{X} & \begin{bmatrix} Q_r\mathcal{X} \\ R_r M \end{bmatrix}^\top \\ \star & \star & -I \end{bmatrix} \prec 0, \qquad (13)$$

define $P := \mathcal{X}^{-1} - T_y^\top Q T_y$, $K := M\mathcal{X}^{-1}$, and choose $\beta > 0$ such that

$$\Xi_f := \{\xi \in \mathbb{R}^{n_\xi} \mid \|\xi - \xi^s\|_P^2 \leq \beta\} \subseteq \mathbb{U}^l \times \mathbb{Y}^l. \qquad (14)$$

Then, Assumption 5 holds with P, K, and Ξ_f.

Proof. By applying the Schur complement to the right-lower block of (13) and subsequently left- and right-multiplying diag$(I, \mathcal{X}^{-1})$, we obtain

$$\begin{bmatrix} \tau \bar{P}_\Delta^w - \begin{bmatrix} \mathcal{X} & 0 \\ 0 & 0 \end{bmatrix} & \begin{bmatrix} A' + B'K \\ I \\ K \end{bmatrix} \\ \star & T_y^\top Q T_y + K^\top R K - \mathcal{X}^{-1} \end{bmatrix} \prec 0. \qquad (15)$$

Applying the Schur complement to the right-lower block of (15) and re-arranging terms leads to $\mathcal{Q} := \mathcal{X}^{-1} - T_y^\top Q T_y - K^\top R K \succ 0$ as well as

$$\begin{bmatrix} I & 0 \\ (A'+B'K)^\top \begin{bmatrix} I \\ K \end{bmatrix}^\top \\ \hline 0 & I \\ B_w^\top & 0 \end{bmatrix}^\top \begin{bmatrix} -\mathcal{X} & 0 & 0 \\ 0 & \mathcal{Q}^{-1} & \\ \hline 0 & & \tau P_\Delta^w \end{bmatrix} \begin{bmatrix} \star \\ \star \\ \hline \star \\ \star \end{bmatrix} \prec 0.$$

Using the full-block S-procedure (compare Scherer (2001)) and $\tilde{A} = A' + B_w\Delta \begin{bmatrix} I \\ 0 \end{bmatrix}$, $\tilde{B} = B' + B_w\Delta \begin{bmatrix} 0 \\ I \end{bmatrix}$, where Δ satisfies (12), the latter inequality implies

$$-\mathcal{X} + (\tilde{A}+\tilde{B}K)\mathcal{Q}^{-1}(\tilde{A}+\tilde{B}K)^\top \prec 0. \tag{16}$$

Applying the Schur complement twice, this in turn implies

$$(\tilde{A}+\tilde{B}K)^\top \mathcal{X}^{-1}(\star) - \mathcal{X}^{-1} + T_y^\top Q T_y + K^\top R K \prec 0. \tag{17}$$

Let now $\xi^+ = (\tilde{A}+\tilde{B}K)\xi$, $y = (\tilde{C}+\tilde{D}K)\xi$ for some $\xi \in \mathbb{R}^{n_\xi}$. We then have

$$\begin{aligned} \|\xi^+\|_P^2 &\overset{(17)}{\leq} \|\xi\|_{\mathcal{X}^{-1}} - \|\xi\|^2_{T_y^\top Q T_y + K^\top R K} - \|\xi^+\|^2_{T_y^\top Q T_y} \\ &= \|\xi\|_P - \|y\|_Q^2 - \|\xi\|^2_{K^\top R K}, \end{aligned}$$

which proves Part (iii) of Assumption 5. Part (i) is a simple consequence of Part (iii). Finally, Part (ii) follows from invariance of $\Xi_f \subseteq \mathbb{U}^l \times \mathbb{Y}^l$. □

Proposition 10 provides a procedure to compute terminal ingredients satisfying Assumption 5, using no model knowledge but only input-output data. The proof relies on the full-block S-procedure (Scherer (2001)), which is used to robustify the model-based matrix inequality (8) against the unknown parameters Δ satisfying the known quadratic bound (12) derived by Berberich et al. (2020c), compare also van Waarde et al. (2020) for a similar approach to data-driven $\mathcal{H}_2$-state-feedback design. The design procedure in Proposition 10 is an LMI feasibility problem which can be solved efficiently in practice. Although any feasible solution $\mathcal{X} \succ 0$, M, $\tau \geq 0$ satisfying (13) leads to terminal ingredients satisfying Assumption 5, the additional conditions $\operatorname{trace}(\Gamma) < \gamma^2$, $\begin{bmatrix} \Gamma & I \\ I & \mathcal{X} \end{bmatrix} \succ 0$ for some $\Gamma \succ 0$, $\gamma > 0$ lead to a smaller terminal cost. To be precise, the optimal solution minimizing γ subject to the conditions in Proposition 10 (provided it exists) is equivalent to a linear quadratic regulator for the system (5), compare Berberich et al. (2020c) for details. Hence, this optimal solution provides a standard choice for terminal ingredients, compare Chen and Allgöwer (1998), leading to a good closed-loop performance. For a reliable numerical implementation, we perform a bisection algorithm over $\gamma > 0$ subject to the conditions in Proposition 10.

Remark 11. The proposed MPC scheme as well as the theoretical analysis in Section 3.2 and the design of the terminal ingredients explained above all require that the measured data are noise-free. However, an extension of the presented results to noisy data is straightforward, where the noise may enter both the offline data used to build the Hankel matrices in (7b) and to compute the terminal ingredients via Proposition 10 as well as the online data used to specify initial conditions in (7c). In particular, it is shown in Berberich et al. (2021b) that (offline and online) output measurement noise in data-driven MPC can be viewed as an input disturbance for the corresponding nominal (noise-free) data-driven MPC scheme. Since model-based MPC with terminal ingredients possesses inherent robustness properties (Yu et al., 2014), we conjecture that the closed loop is practically exponentially stable w.r.t. the noise and disturbance level, provided that Problem (7) is slightly modified (by including a slack variable in (7b) and a regularization on $\alpha(t)$ in the cost, cf. Berberich et al. (2021b) for details). As a noteworthy advantage, this is possible based on a one-step MPC scheme as in Algorithm 6, whereas the robustness guarantees from Berberich et al. (2021a,b) require the application of a multi-step MPC scheme due to the presence of terminal equality constraints. Finally, the (offline) design by Berberich et al. (2020c) which forms the basis for Proposition 10 also applies in the presence of noise and hence, our design of terminal ingredients can be carried out in this case. More precisely, replacing the definition of P_Δ^w in (12) by

$$P_\Delta^w = \begin{bmatrix} -ZZ^\top & Z\Xi_+^\top B_w \\ B_w^\top \Xi_+ Z^\top & \bar{d}I - B_w^\top \Xi_+ \Xi_+^\top B_w \end{bmatrix}$$

for some $\bar{d} > 0$, Proposition 10 leads to terminal ingredients which satisfy Assumption 5 robustly for all systems (5) which are consistent with the measured data, when assuming that data generated offline are perturbed by process noise $\{d_k\}_{k=0}^{N-1}$ and measurement noise $\{\varepsilon_k\}_{k=0}^{N}$ satisfying a bound $\sum_{k=0}^{N-1} \|d_k\|_2^2 + \|\varepsilon_{k+1} - \tilde{A}\varepsilon_k\|_2^2 \leq \bar{d}$ (compare Berberich et al. (2020c) for details). To summarize, all results in this paper can be extended to the realistic scenario where the measured data are affected by noise and the system is subject to disturbances.

Remark 12. In case of noise-free data, an *indirect* approach consisting of system identification and model-based MPC is a simple alternative to our data-driven MPC scheme. In particular, with persistently exciting data, the system matrices can be computed exactly and thus, a model-based design of terminal ingredients and implementation of MPC can be carried out analogous to Proposition 10 and Algorithm 6, respectively. The resulting computational complexity is comparable to that of the *direct* approach proposed in this paper. On the other hand, as discussed in more detail in Remark 11, our data-driven MPC can be extended to noisy data, in which case it may be preferable over the indirect approach due to the challenges in obtaining tight estimation bounds from data (Matni et al., 2019). A theoretical comparison of indirect and direct data-driven MPC is an interesting direction for future research, cf. Dörfler et al. (2021); Krishnan and Pasqualetti (2021) for existing open-loop results.

In the remainder of this section, we investigate when a crucial necessary condition for feasibility of the conditions in Proposition 10 is fulfilled. Since (13) is a strict LMI, its feasibility requires $\tau\bar{P}_\Delta^w - \begin{bmatrix} \mathcal{X} & 0 \\ 0 & 0 \end{bmatrix} \prec 0$, which implies $-ZZ^\top \prec 0$. This in turn requires that the data matrix $Z = \begin{bmatrix} \Xi^\top & U^\top \end{bmatrix}^\top$ has full row rank. The latter condition is related to persistence of excitation (Definition 1). In particular, (Willems et al., 2005, Corollary 2 (ii)) implies that Z has full row rank if $(\tilde{A}, \tilde{B})$ is controllable and u^d is persistently exciting of a sufficiently high order. Thus, one possibility to ensure that (13) is feasible via a sufficiently rich input signal is to require that $(\tilde{A}, \tilde{B})$ is controllable, which need not be the case in general, even if (A, B) is controllable. In order to determine when this sufficient condition for Z having full row rank holds, the following result provides an equivalent condition for controllability of $(\tilde{A}, \tilde{B})$ based on the observability matrix Φ_l in (3).

Lemma 13. The pair $(\tilde{A}, \tilde{B})$ is controllable if and only if Φ_l is square.

Proof. We show that an arbitrary $\xi_k \in \mathbb{R}^{n_\xi}$ is reachable from any initial condition if and only if Φ_l is square. The system dynamics (1) imply for any $k \geq n_\xi$

$$\xi_k = \begin{bmatrix} u_{[k-l,k-1]} \\ \Phi_l x_{k-l} + \Gamma_l u_{[k-l,k-1]} \end{bmatrix} \tag{18}$$

with a suitably defined matrix Γ_l. By observability, i.e., full column rank of Φ_l, we have $pl \geq p\underline{l} \geq n$ and hence, $k-l \geq n_\xi - l \geq n$. Thus, using controllability of (A, B), the system can be steered to an arbitrary state x_{k-l} from any initial condition x_0 by appropriately selecting the input $u_{[0,k-l-1]}$. Therefore, the state ξ_k is reachable if and only if there exist an input $u_{[k-l,k-1]}$ and a state x_{k-l} such that (18) holds. The first block row of (18) is trivially satisfied by definition of ξ_k. Using $u_{[k-l,k-1]} = [I \ 0]\,\xi_k$, the second block row of (18) takes the form

$$[-\Gamma_l \ \ I]\,\xi_k = \Phi_l x_{k-l}. \tag{19}$$

Clearly, there exists a state x_{k-l} for any given ξ_k such that (19) holds if and only if $[-\Gamma_l \ \ I]\,\xi_k$ lies in the image of Φ_l for any $\xi_k \in \mathbb{R}^{n_\xi}$. Since $[-\Gamma_l \ \ I]$ has full row rank, this means that Φ_l has full row rank. Since Φ_l has full column rank by the assumption that (A, C) is observable, full row rank of Φ_l is equivalent to Φ_l being square. □

Lemma 13 shows that, under the present assumptions, the extended dynamics (5) are controllable if and only if the matrix Φ_l is square. This means that full row rank of Z, which is a necessary condition for the presented design of terminal ingredients and thus the proposed MPC scheme, can be guaranteed if Φ_l is square. The latter condition requires $l = \underline{l}$, i.e., the lag of the system is known and chosen for defining the extended state ξ, as well as $pl = n$. For single-output systems with $p = 1$, this in turn requires that $l = n$, i.e., the system order is known exactly. For multiple-output systems on the other hand, the condition $pl = n$ does in general not hold even if both the system order and the lag are known exactly. Addressing this issue and designing data-driven controllers for the extended system (5) in case it is not controllable is an interesting issue for future research, independently of its application to data-driven MPC in this paper.

4. NUMERICAL EXAMPLE

We apply our MPC scheme to a linearized version of the four-tank system considered by Raff et al. (2006), see Berberich et al. (2021a) for the numerical values of A, B, C, D. We consider the prediction horizon $L = 15$, cost matrices $Q = I$, $R = 5 \cdot 10^{-3} I$, an input constraint set $\mathbb{U} = [-2, 2]^2$, but no output constraints, i.e., $\mathbb{Y} = \mathbb{R}^p$, and the input-output setpoint $(u^s, y^s) = \left(\begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 0.65 \\ 0.77 \end{bmatrix}\right)$. Moreover, we regularize $\alpha(t)$ in the cost of Problem (7) by penalizing $10^{-4} \cdot \|\alpha(t)\|_2^2$ for better numerical stability and robustness (compare Coulson et al. (2020); Berberich et al. (2021a)). We generate data of length $N = 100$ of the above system by applying an input sampled uniformly from $u_k \in [-1, 1]^2$. First, we note that the observability matrix Φ_l of this system is square, i.e., the extended system is controllable (compare Lemma 13), for $l = \underline{l} = 2$, i.e., if l is chosen as the lag of the system. Based on this knowledge, we apply Proposition 10 to compute terminal ingredients satisfying Assumption 5. To be precise, using Yalmip (Löfberg (2004)) with the solver MOSEK (MOSEK ApS (2015)), we obtain a feasible solution of the conditions in Proposition 10 with $K = 0$ and $\gamma = 52$. We then apply the MPC scheme based on Problem (7) with $\Xi_f = \mathbb{R}^{n_\xi}$. Since the above system is open-loop stable, the closed loop under this MPC scheme is globally exponentially stable (compare Remark 9). The closed-loop input and output trajectories are displayed in Figure 1. It can be seen that the MPC scheme tracks the desired setpoint reliably. On the other hand, an MPC scheme using no terminal ingredients (i.e., setting $P = 0$ in (7)) leads to an *unstable* closed loop since the prediction horizon is chosen too short. Moreover, an MPC scheme based on terminal equality constraints as proposed by Berberich et al. (2021a) is not initially feasible for the initial condition $x_0 = [0.1 \ 0.1 \ 0.2 \ 0.2]^\top$ with any prediction horizon $L \leq 24$ due to the input constraints. This means that x_0 does not lie in the (bounded) region of attraction of the scheme by Berberich et al. (2021a), in contrast to the global stability guarantees of the MPC scheme designed above. Finally, we note that it is also possible to design a non-zero K based on Proposition 10 resulting in a less conservative terminal penalty. However, in this case, non-trivial terminal set constraints need to be included which reduces the region of attraction. To conclude, our results show that enhancing data-driven MPC with non-trivial terminal ingredients improves both the theoretical properties as well as the practical performance.

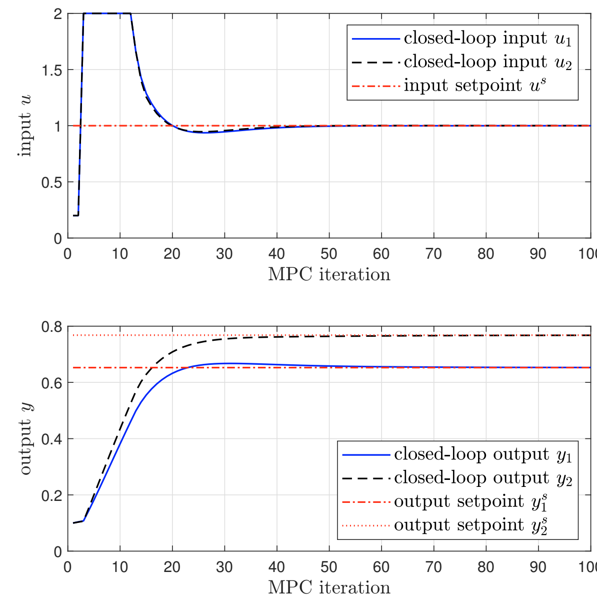

Fig. 1. Closed loop under the data-driven MPC scheme.

5. CONCLUSION

We presented a data-driven MPC scheme with terminal ingredients. The scheme uses only measured input-output data for prediction and thus, the terminal ingredients involve an extended state consisting of consecutive input-output values. We proved closed-loop stability and showed that the required terminal ingredients can be computed based on data without explicit model knowledge. If compared to existing data-driven MPC schemes with closed-loop guarantees (Berberich et al. (2021a)), the presented approach has superior closed-loop robustness and performance properties both in theory and practice.

REFERENCES

Abbas, H. S., Tóth, R., Meskin, N., Mohammadpour, J., Hanema, J., 2016. A robust MPC for input-output LPV models. Trans. Automat. Control 61 (12), 4183–4188.

Adetola, V., Guay, M., 2011. Robust adaptive MPC for constrained uncertain nonlinear systems. Int. J. Adaptive Control and Signal Processing 25 (2), 155–167.

Aswani, A., Gonzalez, H., Sastry, S. S., Tomlin, C., 2013. Provably safe and robust learning-based model predictive control. Automatica 49 (5), 1216–1226.

Bemporad, A., Morari, M., Dua, V., Pistikopoulos, E. N., 2002. The explicit linear quadratic regulator for constrained systems. Automatica 38 (1), 3–20.

Berberich, J., Allgöwer, F., 2020. A trajectory-based framework for data-driven system analysis and control. In: Proc. European Control Conf. pp. 1365–1370.

Berberich, J., Köhler, J., Müller, M. A., Allgöwer, F., 2020a. Data-driven tracking MPC for changing setpoints. IFAC-PapersOnLine 53 (2), 6923–6930.

Berberich, J., Köhler, J., Müller, M. A., Allgöwer, F., 2020b. Robust constraint satisfaction in data-driven MPC. In: Proc. Conf. Decision and Control. pp. 1260–1267.

Berberich, J., Köhler, J., Müller, M. A., Allgöwer, F., 2021a. Data-driven model predictive control with stability and robustness guarantees. Trans. Automat. Control 66 (4), 1702–1717.

Berberich, J., Köhler, J., Müller, M. A., Allgöwer, F., 2021b. Linear tracking MPC for nonlinear systems part II: the data-driven case. arXiv:2105.08567.

Berberich, J., Scherer, C. W., Allgöwer, F., 2020c. Combining prior knowledge and data for robust controller design. arXiv:2009.05253.

Berkenkamp, F., Turchetta, M., Schoellig, A., Krause, A., 2017. Safe model-based reinforcement learning with stability guarantees. In: NeurIPS. pp. 908–918.

Cai, C., Teel, A. R., 2008. Input–output-to-state stability for discrete-time systems. Automatica 44 (2), 326–336.

Chen, H., Allgöwer, F., 1998. A quasi-infinite horizon nonlinear model predictive control scheme with guaranteed stability. Automatica 34 (10), 1205–1217.

Coulson, J., Lygeros, J., Dörfler, F., 2019. Data-enabled predictive control: in the shallows of the DeePC. In: Proc. European Control Conf. pp. 307–312.

Coulson, J., Lygeros, J., Dörfler, F., 2020. Distributionally robust chance constrained data-enabled predictive control. arXiv:2006.01702.

Dörfler, F., Coulson, J., Markovsky, I., 2021. Bridging direct & indirect data-driven control formulations via regularizations and relaxations. arXiv:2101.01273.

Dutta, A., Hartley, E., Maciejowski, J., De Keyser, R., 2014. Certification of a class of industrial predictive controllers without terminal conditions. In: Proc. Conf. Decision and Control. pp. 6695–6700.

Goodwin, G. C., Sin, K. S., 2014. Adaptive filtering prediction and control. Courier Corporation.

Grimm, G., Messina, M. J., Tuna, S. E., Teel, A. R., 2005. Model predictive control: for want of a local control Lyapunov function, all is not lost. Trans. Automat. Control 50 (5), 546–558.

Koch, A., Berberich, J., Allgöwer, F., 2020. Provably robust verification of dissipativity properties from data. arXiv:2006.05974.

Krishnan, V., Pasqualetti, F., 2021. On direct vs indirect data-driven predictive control. arXiv:2103.14936.

Limón, D., Alamo, T., Salas, F., Camacho, E. F., 2006. On the stability of constrained MPC without terminal constraint. Trans. Automat. Control 51 (5), 832–836.

Limón, D., Alvarado, I., Alamo, T., Camacho, E. F., 2008. MPC for tracking piecewise constant references for constrained linear systems. Automatica 44 (9), 2382–2387.

Löfberg, J., 2004. Yalmip: a toolbox for modeling and optimization in matlab. In: Proc. Int. Conf. Robotics and Automation. pp. 284–289.

Markovsky, I., Rapisarda, P., 2008. Data-driven simulation and control. Int. J. Control 81 (12), 1946–1959.

Matni, N., Proutiere, A., Rantzer, A., Tu, S., 2019. From self-tuning regulators to reinforcement learning and back again. In: Proc. Conf. Decision and Control. pp. 3724–3740.

MOSEK ApS, 2015. The MOSEK optimization toolbox for MATLAB manual, version 7.1 (revision 28).

Raff, T., Huber, S., Nagy, Z. K., Allgöwer, F., 2006. Nonlinear model predictive control of a four tank system: An experimental stability study. In: Proc. Int. Conf. Control Applications. pp. 237–242.

Rawlings, J. B., Mayne, D. Q., Diehl, M. M., 2020. Model Predictive Control: Theory, Computation, and Design, 2nd Edition. Nob Hill Pub, 3rd printing.

Rawlings, J. B., Muske, K. R., 1993. The stability of constrained receding horizon control. Trans. Automat. Control 38 (10), 1512–1516.

Scherer, C. W., 2001. LPV control and full block multipliers. Automatica 37 (3), 361–375.

Tanaskovic, M., Fagiano, L., Smith, R., Morari, M., 2014. Adaptive receding horizon control for constrained MIMO systems. Automatica 50 (12), 3019–3029.

van Waarde, H. J., Kanat Camlibel, M., Mesbahi, M., 2020. From noisy data to feedback controllers: nonconservative design via a matrix S-lemma. Trans. Automat. Control.

Willems, J. C., Rapisarda, P., Markovsky, I., De Moor, B., 2005. A note on persistency of excitation. Systems & Control Letters 54, 325–329.

Yang, H., Li, S., 2015. A data-driven predictive controller design based on reduced hankel matrix. In: Proc. Asian Control Conf. pp. 1–7.

Yu, S., Reble, M., Chen, H., Allgöwer, F., 2014. Inherent robustness properties of quasi-infinite horizon nonlinear model predictive control. Automatica 50 (9), 2269–2280.

Zanon, M., Gros, S., 2020. Safe reinforcement learning using robust MPC. Trans. Automat. Control.

Zhou, K., Doyle, J. C., Glover, K., 1996. Robust and optimal control. Prentice-Hall, Englewood Cliffs, N.J.