

Tuning of subspace predictive controls[★]

V. Breschi^{*}, T. Bou Hamdan^{**}, G. Mercère^{**},
S. Formentin^{***}

^{*} Department of Electrical Engineering, Eindhoven University of Technology, 5600 MB, Eindhoven, Netherlands (e-mail: v.breschi@tue.nl)

^{**} University of Poitiers, Laboratoire d'Informatique et d'Automatique pour les Systèmes, 40 avenue du Recteur Pineau, 86000 Poitiers, France (e-mail: {*name.surname*}@univ-poitiers.fr)

^{***} Politecnico di Milano, P.zza Leonardo da Vinci 32, 20133, Milan, Italy (e-mail: simone.formentin@polimi.it)

Abstract: Data-driven predictive control has recently gained increasing attention, as it makes it possible to design constrained controls directly from a set of data, without requiring an intermediate identification step. In this paper, we focus on a Subspace Predictive Control (SPC) scheme, with the aim of clarifying the sensitivity of the final closed-loop performance to its main hyperparameters, namely the length of the past horizon and the regularization penalties. Moreover, by delving deep into the structural properties of the control problem formulation, we provide a set of guidelines for the choice of such hyperparameters. The effectiveness of the resulting overall tuning strategy is assessed on two benchmark examples.

Copyright © 2023 The Authors. This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Keywords: Data-driven predictive control, Subspace identification, Constrained control

1. INTRODUCTION

Data-driven control solutions allow to directly exploit batches of data to design a controller for unknown (and potentially complex) systems (see *e.g.*, Formentin et al. (2014)). Specifically, *Data-Driven Predictive Control* (DDPC) techniques combine the advantages of *Model Predictive Control* (MPC) (see Rawlings (2000)), *i.e.*, the possibility to explicitly account for constraints, with the direct mapping of data onto control actions Coulson et al. (2019); Berberich et al. (2020). Among the others, *Subspace Predictive Control* (SPC) was first introduced in Favoreel et al. (1999) and exploits the tools of subspace identification (see Verhaegen and Verdult (2007); Mercère (2013)) to handle noisy data directly in the definition of the needed predictor (the connections between SPC and more recent DDPC have been investigated in Fiedler and Lucia (2020); Breschi et al. (2023)).

Along the footsteps of Favoreel et al. (1999), in this work we provide a comprehensive overview of the theory and procedures of subspace identification leveraged in the construction of SPC schemes. This allows us to pinpoint that, apart from the hyperparameters typical of MPC (*e.g.*, the penalties in the cost function, the prediction horizon etc.), the only additional degree of freedom in classical SPC is the *past horizon* used to construct the initialization for the predictive control problem, from input/output data only. Nonetheless, well-known results in subspace identification allows us to indicate a set of guidelines for assembling the initial conditions of the problem, *without the need for*

tuning the past horizon with closed-loop experiments. As briefly shown in Breschi et al. (2023), the applicability of these guidelines extends beyond the subspace-based control approach considered in this work, since also other existing data-driven predictive control strategies estimate the initial conditions in the same way (see *e.g.*, Coulson et al. (2019); Berberich et al. (2020)). We then shift to the augmented SPC scheme proposed in Fiedler and Lucia (2020), which comprises a set of slack variables used to handle noise in the online data. The derivation of the solution to this subspace predictive control problem in closed-form (in a constrained case) allows us to discuss the impact that the regularization penalties introduced along with the slacks potentially have on the resulting optimal control action. Our conclusions are supported by the results attained on two benchmark numerical examples.

The remainder of the paper is as follows. The considered stochastic setting and the control problem of interest are introduced in Section 2. In Section 3, we review the main tools of subspace identification, which we rely upon in Section 4 to shed light on some key hyperparameters of the SPC scheme augmented with slacks proposed in Fiedler and Lucia (2020). The new insights are numerically tested in Section 5. Some closing remarks end the paper.

Notation. Given a rectangular matrix $B \in \mathbb{R}^{m \times n}$, $B^\top \in \mathbb{R}^{n \times m}$ denotes its transpose and B^\dagger indicates its right inverse. If $Q \succ 0$ ($Q \succeq 0$), then the matrix $Q \in \mathbb{R}^{n \times n}$ is positive definite (positive semi-definite). Let $x \in \mathbb{R}^n$, the quadratic form $x^\top Q x$ is compactly denoted as $\|x\|_Q^2$, while $\|x\|_2$ denotes its 2-norm. For any vector $x_t \in \mathbb{R}^n$, with $t \in \mathbb{N}$, we define the finite past and future stack vectors as

[★] This project was partially supported by the Italian Ministry of University and Research under the PRIN'17 project "Data-driven learning of constrained control systems", contract no. 2017J89ARP.

$$x_{\rho,t}^- = \begin{bmatrix} x_{t-\rho} \\ \vdots \\ x_{t-2} \\ x_{t-1} \end{bmatrix} \in \mathbb{R}^{n\rho}, \quad x_{L,t}^+ = \begin{bmatrix} x_t \\ x_{t+1} \\ \vdots \\ x_{t+L-1} \end{bmatrix} \in \mathbb{R}^{nL}, \quad (1)$$

where $\rho > 0$ and $L > 0$ denotes the length of the considered past and future windows. Accordingly, we can define the associated Hankel matrices as:

$$X_{\rho,M}^- = [x_{\rho,t}^- \cdots x_{\rho,t+M-1}^-] \in \mathbb{R}^{n\rho \times M}, \quad (2a)$$

$$X_{L,M}^+ = [x_{L,t}^+ \cdots x_{L,t+M-1}^+] \in \mathbb{R}^{nL \times M}. \quad (2b)$$

Moreover, given a set of matrices (A, B, C, D) of appropriate dimensions, we define the following block matrices:

$$\Omega_\ell(A, B) = [A^{\ell-1}B \cdots AB \ B], \quad (3a)$$

$$\Gamma_\ell(A, C) = \begin{bmatrix} A \\ CA \\ \vdots \\ CA^{\ell-1} \end{bmatrix}, \quad (3b)$$

$$H_\ell(A, B, C, D) = \begin{bmatrix} D & 0 & \cdots & 0 \\ CB & D & \cdots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ CA^{\ell-2}B & \cdots & CB & D \end{bmatrix}, \quad (3c)$$

where $\ell > 0$.

2. SETTING AND GOAL

Consider a *linear time invariant* (LTI) system \mathcal{S} , with dynamics characterized by the following difference equations:

$$x_{t+1} = Ax_t + Bu_t + w_t, \quad (4a)$$

$$y_t = Cx_t + Du_t + v_t, \quad (4b)$$

where $x_t \in \mathbb{R}^n$, $u_t \in \mathbb{R}^m$ and $y_t \in \mathbb{R}^p$ are the state, input and output of the system, while $w_t \in \mathbb{R}^n$ and $v_t \in \mathbb{R}^p$ are realizations of the process and measurement noises at time $t \in \mathbb{N}$. Assume that the latter are statistically independent of the input (when the plant operates in open-loop), and that they are mutually independent, zero-mean white noises with joint covariance matrix:

$$\mathbb{E} \left\{ \begin{bmatrix} v_t \\ w_t \end{bmatrix} \begin{bmatrix} v_j^\top & w_j^\top \end{bmatrix} \right\} = \begin{bmatrix} R^v & S^\top \\ S & Q^w \end{bmatrix} \delta_{t,j} \succeq 0, \quad (4c)$$

where $R^v \succ 0$, $Q^w \succeq 0$ and $\delta_{t,j}$ is the Kronecker delta

$$\delta_{t,j} = \begin{cases} 1 & \text{if } t = j, \\ 0 & \text{if } t \neq j. \end{cases}$$

Additionally, suppose that \mathcal{S} is *minimal*, namely the pair (A, C) is observable and the pair $(A, [B, (Q^w)^{1/2}])$ is controllable. By considering the typical data-driven control setting, let us finally assume that the matrices (A, B, C, D) , along with (Q^w, R^v, S) are *unknown*. Nonetheless, suppose that we can collect a set of input/output pairs $\mathcal{D}_N = \{u_t, y_t\}_{t=0}^N$ from \mathcal{S} . Our goal is to steer the output of the system towards a reference target $r \in \mathbb{R}^p$, while guaranteeing that

$$u_t \in \mathcal{U}, \quad y_t \in \mathcal{Y}, \quad \forall t, \quad (5)$$

with $\mathcal{U} \subseteq \mathbb{R}^m$ and $\mathcal{Y} \in \mathbb{R}^p$ formalizing convex limitations on the actuators and desired bounds on the output.

3. SUBSPACE IDENTIFICATION TOOLS FOR SPC

We now review the main steps for the derivation of the predictor used in SPC schemes, see *e.g.*, Favoreel et al. (1999); Fiedler and Lucia (2020). In particular, given a sufficiently large set of input/output data \mathcal{D}_N , most linear regression-based subspace methods (see, *e.g.*, Bauer (2003)) initially involve the estimation of the state sequence $\{x_t\}_{t=0}^{N+1}$. This step is crucial to recast the state-space model of \mathcal{S} as:

$$\omega_t = \Theta\phi_t + \nu_t, \quad (6a)$$

with

$$\omega_t = \begin{bmatrix} x_{t+1} \\ y_t \end{bmatrix}, \quad \Theta = \begin{bmatrix} A & B \\ C & D \end{bmatrix}, \quad \phi_t = \begin{bmatrix} x_t \\ u_t \end{bmatrix}, \quad \nu_t = \begin{bmatrix} v_t \\ w_t \end{bmatrix}, \quad (6b)$$

thus ultimately allowing for the identification of \mathcal{S} by solving a linear regression problem. For a comprehensive summary of the main steps towards the derivation of the data-based output predictor, we thus initially discuss how this state estimate is derived from input/output data and, then, we go through some preliminary model-based results, finally retrieving the predictor used in SPC.

Based on our assumptions on \mathcal{S} and the process and measurement noises, it can be proven that (4) is equivalent to the innovation form, namely

$$x_{t+1} = Ax_t + Bu_t + Ke_t, \quad (7a)$$

$$y_t = Cx_t + Du_t + e_t, \quad (7b)$$

where $\mathcal{K} \in \mathbb{R}^{n \times p}$ is the steady-state Kalman gain, $e_t \in \mathbb{R}^p$ is the so-called innovation vector. Because of the properties of \mathcal{S} , all eigenvalues of the matrix $\tilde{A} = A - \mathcal{K}C$ lay inside the unit circle. By exploiting the output equation (7b), this alternative representation of the system can be manipulated as follows:

$$x_{t+1} = \underbrace{(A - \mathcal{K}C)}_{\tilde{A}} x_t + \underbrace{(B - \mathcal{K}D)}_{\tilde{B}} u_t + \mathcal{K}y_t. \quad (8)$$

Accordingly, the state at each time instant can be straightforwardly recast as the infinite sum of past inputs and outputs

$$x_t = \sum_{j=1}^{\infty} \tilde{A}^{j-1} \tilde{B} u_{t-j} + \sum_{j=1}^{\infty} \tilde{A}^{j-1} \mathcal{K} y_{t-1}, \quad (9)$$

which can be further decomposed as

$$\begin{aligned} x_t &= \sum_{j=1}^{\rho} \tilde{A}^{j-1} \tilde{B} u_{t-j} + \sum_{j=1}^{\rho} \tilde{A}^{j-1} \mathcal{K} y_{t-1} + \delta_t \\ &= \Omega_\rho(\tilde{A}, \tilde{B}) u_{\rho,t}^- + \Omega_\rho(\tilde{A}, \mathcal{K}) y_{\rho,t}^- + \delta_t \\ &= \tilde{A}^\rho x_{t-\rho} + \Omega_\rho(\tilde{A}, \tilde{B}) u_{\rho,t}^- + \Omega_\rho(\tilde{A}, \mathcal{K}) y_{\rho,t}^-. \end{aligned} \quad (10a)$$

with

$$\delta_t = \sum_{j=\rho+1}^{\infty} \tilde{A}^{j-1} \tilde{B} u_{t-j} + \sum_{j=\rho+1}^{\infty} \tilde{A}^{j-1} \mathcal{K} y_{t-1}. \quad (10b)$$

Provided that ρ is *sufficiently large*, the state at time t can thus be consistently approximated by truncation as (see Bauer (2003))

$$\bar{x}_t = \Omega_\rho(\tilde{A}, \mathcal{K}) y_{\rho,t}^- + \Omega_\rho(\tilde{A}, \tilde{B}) u_{\rho,t}^-. \quad (11)$$

According to (4), we can now introduce the L -steps ahead predictor of the system's output, *i.e.*,

$$y_{L,t}^+ = \Gamma_L(A, C)x_t + H_L(A, B, C, D)u_{L,t}^+ + H_L(A, \mathcal{K}, C, I_p)e_{L,t}^+. \quad (12)$$

By defining the vector $z_{\rho,t}^- \in \mathbb{R}^{(m+p)\rho}$ stacking the past ρ outputs and the inputs at time t , namely

$$z_{\rho,t}^- = \begin{bmatrix} y_{\rho,t}^- \\ u_{\rho,t}^- \end{bmatrix}. \quad (13)$$

and replacing it into (10a) into (12), we then obtain

$$y_{L,t}^+ = \underbrace{\Gamma_L(A, C)\Omega_\rho(\tilde{A}, \mathcal{B})}_{\mathcal{L}_{L,\rho}} z_{\rho,t}^- + \underbrace{H_L(A, B, C, D)}_{H_L^u} u_{L,t}^+ + \eta_{L,t}^+, \quad (14a)$$

with $\mathcal{B} = [\mathcal{K} \tilde{B}]$ and

$$\eta_{L,t}^+ = \underbrace{H_L(A, \mathcal{K}, C, I_p)}_{H_L^e} e_{L,t}^+ + \Gamma_L(A, C)\tilde{A}^\rho x_{t-\rho}. \quad (14b)$$

Assume that the future horizon L and the hyperparameter ρ are chosen to be greater or equal than the order of the system n and sufficiently large to guarantee that \tilde{A} is *nilpotent* of degree ρ , namely $\tilde{A}^i = 0$ for all $i \geq \rho$. Based on this hypothesis and the ones on the process and measurement noises, then: (i) $\tilde{A}^\rho x_{t-\rho}$ is negligible: (ii) the noise term $H_L^e e_{L,t}^+$ is uncorrelated with both $z_{\rho,t}^-$ and $u_{L,t}^+$ and (iii) the following properties hold:

$$\text{rank}(\Gamma_{L,\rho}) = n, \quad (15a)$$

$$\text{col}(\Gamma_{L,\rho}) = \text{col}(\Gamma_L(A, C)). \quad (15b)$$

We have now all the elements to translate (14a) into its purely data-based counterpart as:

$$Y_{L,M}^+ = \mathcal{L}_{L,\rho} Z_{\rho,M}^- + H_L^u U_{L,M}^+ + N_{L,M}^+, \quad (16)$$

where $Y_{L,M}^+$, $U_{L,M}^+$ and $Z_{\rho,M}^- \in \mathbb{R}^{(m+p)L \times M}$ are Hankel matrices constructed from the available data, while $\mathcal{L}_{L,\rho}$ and H_L^u are *unknown*. Since $\eta_{L,t}^+$ is uncorrelated with $z_{\rho,t}^-$ and $u_{L,t}^+$ and this property can be extended to the Hankel matrix $N_{L,M}^+$, the unknown matrices in (16) can be estimated by solving the following least squares problem

$$\underset{\mathcal{L}_{L,\rho}, H_L^u}{\text{minimize}} \left\| Y_{L,M}^+ - [\mathcal{L}_{L,\rho} \ H_L^u] \begin{bmatrix} Z_{\rho,M}^- \\ U_{L,M}^+ \end{bmatrix} \right\|_F^2. \quad (17)$$

Whenever the following rank condition holds:

$$\text{rank} \left(\lim_{M \rightarrow \infty} \begin{bmatrix} Z_{\rho,M}^- \\ U_{L,M}^+ \end{bmatrix} \begin{bmatrix} Z_{\rho,M}^- \\ U_{L,M}^+ \end{bmatrix}^\top \right) = m(\rho + L) + p\rho, \quad (18)$$

the unique solution of the optimization problem in (17) can be solved in closed form, resulting in the following estimates:

$$[\hat{\mathcal{L}}_{L,\rho} \ \hat{H}_L^u] = Y_{L,M}^+ \begin{bmatrix} Z_{\rho,M}^- \\ U_{L,M}^+ \end{bmatrix}^\dagger, \quad (19a)$$

with

$$\lim_{M \rightarrow \infty} \frac{1}{M} N_{L,M}^+ [(U_{L,M}^+)^\top (Z_{\rho,M}^-)^\top] = 0. \quad (19b)$$

As $N_{L,M}^+$, $U_{L,M}^+$ and $Z_{\rho,M}^-$ satisfy (19b), the estimates $\hat{\mathcal{L}}_{L,\rho}$ and \hat{H}_L^u are asymptotically unbiased. While this feature is directly exploited in SPC, the latter is neglected in the behavioral predictors used in Coulson et al. (2019); Berberich et al. (2020). Note that the rank condition in (18) at the core of these results can be satisfied under mild excitation conditions on the input sequence used throughout data collection (see (Verhaegen and Verdult, 2007, Lemma 9.9)).

4. THE ROLE OF THE HYPERPARAMETERS

The estimates in (19a) in combination with (14a) allows to formulate the SPC problem as follows:

$$\underset{\tilde{u}_{L,t}^+, \sigma}{\text{minimize}} \sum_{k=0}^{L-1} \|\tilde{y}_{k,t} - r\|_Q^2 + \|\tilde{u}_{k,t}\|_R^2 + \lambda_u \|\sigma_u\|^2 + \lambda_y \|\sigma_y\|^2 \quad (20a)$$

$$\text{s.t.} \quad \tilde{y}_{L,t}^+ = Y_{L,M}^+ \begin{bmatrix} Z_{\rho,M}^- \\ U_{L,M}^+ \end{bmatrix}^\dagger \begin{bmatrix} z_{\rho,t}^- + \sigma \\ \tilde{u}_{L,t}^+ \end{bmatrix}, \quad (20b)$$

$$\tilde{u}_{k,t} \in \mathcal{U}, \quad \tilde{y}_{k,t} \in \mathcal{Y}, \quad k = 0, \dots, L-1, \quad (20c)$$

where $r \in \mathbb{R}^p$ is the constant reference we aim the system to track, $L \in \mathbb{N}$ is the prediction and $Q \succeq 0$ and $R \succ 0$ are the penalties associated with the tracking error and the control effort, respectively. These hyperparameters are assumed to be fixed a priori and based on the control objectives. Note that, as for the classical SPC formulation in Favoreel et al. (1999), $z_{\rho,t}^-$ is obtained by stacking closed-loop past inputs and outputs as in (13), while we retain the deterministic part of the output predictor (14a) only. As such, once the controller is deployed, the predictor is *simulated* in open-loop (so, without noise) over the horizon L to retrieve the optimal control sequence. With respect to the SPC scheme presented in Favoreel et al. (1999), here we look at its augmentation with additional slack variables introduced in Fiedler and Lucia (2020). Indeed, (20b) features a slack, *i.e.*, $\sigma = [\sigma_y^\top \ \sigma_u^\top]^\top \in \mathbb{R}^{\rho(m+p)}$, introduced to cope with noise potentially affecting $z_{\rho,t}^-$, while $\lambda_u, \lambda_y \geq 0$ in (20a) are two hyperparameters to be tuned. As in standard MPC, the optimal control action, to be fed to the unknown system \mathcal{S} at each time instant t , is still found by solving (20) to find the optimal input sequence $\{\tilde{u}_{k,t}^*\}_{k=0}^{L-1}$, then retaining its first element only and setting $u_t = \tilde{u}_0^*$. Since the prediction horizon L and the weights Q and R in (20) are fixed based on the predictive controller goals, the only hyperparameters left to be tuned are $\rho > 0$, which characterizes the approximation in (11), and $\lambda_u, \lambda_y > 0$, weighting the additional slacks.

As it can be seen in (11), ρ dictates the approximation of the state as a combination of past inputs and outputs. Let us indicate as ξ_t the error in the state reconstruction due to the above approximation, namely,

$$\xi_t = x_t - \bar{x}_t. \quad (21)$$

As shown in Bauer (2003), ξ_t is orthogonal to $z_{\rho,t}^-$ and it satisfies the following:

$$\|\xi_t\|_2 \leq \|\tilde{A}^\rho\|_F \|x_{t-\rho}\|_2. \quad (22)$$

Based on the features of \tilde{A} (see Section 3), a direct consequence of this inequality is that the approximation error ξ_t becomes small if ρ is large *enough*, vanishing in the ideal case of an infinite past window (*i.e.*, $\rho \rightarrow \infty$). At the same time, the variance of the predictor's parameters tends to increase for growing values of ρ . Although here (14a) is not exploited for identification purposes, this effect is still undesirable. Indeed, it implies a lack of robustness of the predictive controller to different realizations of noise on the batch data used to construct (20b). Another downside of augmenting ρ lays in the increasing complexity of the optimization problem to be solved, with higher dimension matrices that need to be

inverted when $\rho \rightarrow \infty$. A possible approach to trade-off between reducing the approximation error ξ_t , the parameters' variance and the computational costs of data-based control is to select ρ that minimizes specific Akaike information criteria. Note that, this procedure can be carried out *offline*, by exploiting the available batch of data \mathcal{D}_N and without endangering the safety of the system.

Let us now recast the predictor in (20b) as follows:

$$\tilde{y}_L^+ = \mathcal{H}_\rho z_{\rho,t}^- + \underbrace{\begin{bmatrix} \mathcal{H}_f & \mathcal{H}_{\rho,y} & \mathcal{H}_{\rho,u} \end{bmatrix}}_{\mathcal{H}_\alpha} \underbrace{\begin{bmatrix} \tilde{u}_{L,t}^+ \\ \sigma_y \\ \sigma_u \end{bmatrix}}_{\alpha_t}, \quad (23a)$$

where

$$\mathcal{H} = Y_{L,M}^+ \begin{bmatrix} Z_{\rho,M}^- \\ U_{L,M}^+ \end{bmatrix}^\dagger \quad (23b)$$

and

$$\mathcal{H}_\rho = [\mathcal{H}]_{1:\tilde{n}}, \quad \mathcal{H}_f = [\mathcal{H}]_{\tilde{n}+1:mL+\tilde{n}}, \quad (23c)$$

$$\mathcal{H}_{\rho,y} = [\mathcal{H}]_{1:p\rho}, \quad \mathcal{H}_{\rho,u} = [\mathcal{H}]_{p\rho+1:\tilde{n}}, \quad (23d)$$

with $\tilde{n} = (m+p)\rho$. We now aim at shedding a light on the impact that the additional slack variables, operating as a sort of *filter* with respect to the noise affecting measurements in closed-loop, have on the solution of the final optimization problem. To this end, we aim at finding a closed-form expression for the solution of the following predictive control problem:

$$\underset{\alpha_t}{\text{minimize}} \quad \|\alpha_t\|_{W_\alpha}^2 + 2(\mathcal{H}_\rho z_{\rho,t}^- - r)^\top c_\alpha \alpha_t \quad (24a)$$

$$\text{s.t.} \quad [I \ 0 \ 0] \alpha_t \in \mathcal{U}^L, \quad (24b)$$

$$\mathcal{H}_\rho z_{\rho,t}^- + \mathcal{H}_\alpha \alpha_t \in \mathcal{Y}^L, \quad (24c)$$

where the new weights are

$$W_\alpha = \mathcal{H}_\alpha^\top \mathcal{Q} \mathcal{H}_\alpha + \begin{bmatrix} \mathcal{R} & 0 & 0 \\ 0 & \lambda_y I & 0 \\ 0 & 0 & \lambda_u I \end{bmatrix}, \quad (24d)$$

$$c_\alpha = \mathcal{Q} \mathcal{H}_\alpha, \quad (24e)$$

with $\mathcal{Q} = \text{diag}([Q, \dots, Q])$ and $\mathcal{R} = \text{diag}([R, \dots, R])$. Since $R \succ 0$, it is straightforward to prove that \mathcal{R} and, thus, W_α is positive definite even when $\lambda_u = \lambda_y = 0$, provided that also σ is set to zero (*i.e.*, the classical SPC scheme of Favoreel et al. (1999) is recovered). Indeed, the first term in (24d) leads to a mutual dependence of the optimal input sequence and slack variables.

For the sake of a better understanding, let us assume the convex constraints in (24b)-(24c) to be polytopic¹, so as to recast them as follows:

$$\mathcal{C}_\alpha \alpha_t + \mathcal{C}_z z_{\rho,t}^- \leq \mathcal{G}. \quad (25)$$

Given the characteristics of (24), we can thus find its unique solution via the *Karush-Kuhn-Tucker* (KKT) conditions, *i.e.*, by solving the following set of equalities:

$$W_\alpha \alpha_t + c_\alpha^\top (\mathcal{H}_\rho z_{\rho,t}^- - r) + \mathcal{C}_\alpha^\top \beta = 0, \quad (26a)$$

$$\beta^\top (\mathcal{C}_\alpha \alpha_t + \mathcal{C}_z z_{\rho,t}^- - \mathcal{G}) = 0, \quad (26b)$$

$$\mathcal{C}_\alpha \alpha_t + \mathcal{C}_z z_{\rho,t}^- - \mathcal{G} \leq 0, \quad (26c)$$

$$\beta \geq 0, \quad (26d)$$

where β is the vector of scaled Lagrange multipliers associated with the polytopic constraints. To find a closed-form expression for α_t , let us focus on a single set of active

constraints, so as to split β into $\bar{\beta}$ and $\tilde{\beta}$, respectively denoting the Lagrange multipliers associated with inactive and active constraints. While $\bar{\beta} = 0$ for (26b) to be satisfied, $\tilde{\beta}$ is such that the following have to hold:

$$W_\alpha \alpha_t + c_\alpha^\top (\mathcal{H}_\rho z_{\rho,t}^- - r) + \tilde{\mathcal{C}}_\alpha \tilde{\beta} = 0, \quad (27)$$

$$\tilde{\mathcal{C}}_\alpha \alpha_t + \tilde{\mathcal{C}}_z z_{\rho,t}^- - \tilde{\mathcal{G}} = 0, \quad (28)$$

where $\tilde{\mathcal{C}}_\alpha$, $\tilde{\mathcal{C}}_z$ and $\tilde{\mathcal{G}}$ are the rows of \mathcal{C}_α , \mathcal{C}_z and \mathcal{G} associated to the considered set of active constraints, respectively. By combining these relations, the optimal α_t^* and the associated $\tilde{\beta}$ are given by:

$$\tilde{\beta} = -\tilde{\mathcal{C}}_{\alpha,W} \left[\tilde{\mathcal{C}}_{z,W} z_{\rho,t}^- - \tilde{\mathcal{G}}_W \right], \quad (29a)$$

$$\alpha_t^* = -W_\alpha^{-1} \left[\tilde{\mathcal{C}}_\alpha^\top \tilde{\beta} + c_\alpha^\top \mathcal{H}_\rho z_{\rho,t}^- + c_\alpha^\top r \right], \quad (29b)$$

where

$$\tilde{\mathcal{C}}_{\alpha,W} = (\tilde{\mathcal{C}}_\alpha W_\alpha^{-1} \tilde{\mathcal{C}}_\alpha^\top)^{-1},$$

$$\tilde{\mathcal{C}}_{z,W} = \tilde{\mathcal{C}}_\alpha W_\alpha^{-1} c_\alpha^\top \mathcal{H}_\rho - \tilde{\mathcal{C}}_z,$$

$$\tilde{\mathcal{G}}_W = \tilde{\mathcal{C}}_\alpha W_\alpha^{-1} c_\alpha^\top r - \tilde{\mathcal{G}}.$$

This closed-form solution only holds in the region identified by the following inequalities²

$$\tilde{\mathcal{C}}_\alpha \alpha_t^* + \tilde{\mathcal{C}}_z z_{\rho,t}^- - \tilde{\mathcal{G}} < 0, \quad (29c)$$

$$-\tilde{\beta} < 0. \quad (29d)$$

This explicit expression of the *local* optimal solution highlights the dependence on the inverse of W_α , which is shaped by λ_u and λ_y (see (24d)). At the same time, these parameters regulate the strength of the shrinkage on the slack variables. As such, high values of λ_u and λ_y tend to steer σ_u and σ_y toward zero, while heavily impacting on the value of W_α and, thus, influencing the resulting optimal control action. On the contrary, smaller λ_u and λ_y are likely to lead to slack variables that are not negligible, that, in turn, shape the input fed to the system, while having a reduced impact on the weight of (24a). Trading off between these two behaviors thus becomes crucial for the effect of the slack to eventually improve the robustness with respect to additive noise acting in closed-loop, while not excessively biasing the resulting optimal sequence.

By looking at the effect of the additional slack, it is clear that for ρ sufficiently high (potentially $\rho \rightarrow \infty$) the true state of the system can be perfectly reconstructed from *noisy data*. In such a case, the slack variables are thus not needed. Although the selection of an infinite past horizon is not feasible in practice, a proper choice of ρ , *e.g.*, through Akaike's criteria, would once again make the introduction of the slack potentially unnecessary.

It is worth to further point out that currently the effect of the choice of λ_u and λ_y can be assessed only online. As such, the calibration of these parameters has to be performed once the predictive controller has been deployed, possibly undermining the safety of \mathcal{S} (*e.g.*, in case the controller destabilizes the system).

5. NUMERICAL CASE STUDIES

To assess the validity of the previous conjectures, we now analyze the impact of ρ and the slacks on the SPC solutions

¹ This assumption is commonly verified in linear predictive control, see Rawlings (2000).

² By following this procedure for all possible combinations of active constraints, one can retrieve the explicit solution.




Fig. 1. Benchmark example: oracle (black dashed line) *vs* average input/output trajectories (red line) and their standard deviation (shaded area) over 20 Monte Carlo predictors and simulations, $\lambda = 0$ and $\rho = 23$.

for two benchmark case studies. To quantitatively assess the capability of the SPC+slack controller, we consider the following performance indexes:

$$\mathcal{J} = \sum_{t=0}^{T-1} \|y_t - r\|_Q^2 + \|u_t\|_R^2, \quad (30a)$$

$$\mathcal{J}_u = \sum_{t=0}^{T-1} \|u_t\|_2^2, \quad (30b)$$

reflecting the overall cost of the optimal input sequence and the corresponding effort. These indexes are compared with their *oracle* counterparts, *i.e.*, those obtained by solving an MPC problem with the exact model of the system and within a noise-free setting, respectively denoted as \mathcal{J}^o and \mathcal{J}_u^o . To slightly simplify our analysis, in all case studies we impose $\lambda_u = \lambda_y = \lambda$ in (20a). Moreover, we always assume the output of the system to be unconstrained, namely we set $\mathcal{Y} = \mathbb{R}^p$.

Benchmark example. Consider the system

$$x_{t+1} = \begin{bmatrix} 0.7326 & -0.0861 \\ 0.1722 & 0.9909 \end{bmatrix} x_t + \begin{bmatrix} 0.0609 \\ 0.0064 \end{bmatrix} u_t + \mathcal{K}e_t \quad (31a)$$

$$y_t = [0 \quad 1.4142] x_t + e_t, \quad (31b)$$

taken from Bemporad et al. (2002), to which we have added both process and measurement noise, with $e \sim \mathcal{N}(0, 0.01)$ and $\mathcal{K} = [0.9089 \ 0.4838]^T$, chosen at random, yet guaranteeing that all eigenvalues of $A - \mathcal{K}\mathcal{C}$ are inside the unit circle. Our control objective in this case is steering the output of the system to zero, namely $r = 0$, over a closed-loop test of $T = 50$ steps, while limiting the required control effort. To attain this goal, we impose $Q = 1$ and $R = 10^{-3}$ in (20a), and we fix the prediction horizon at $L = 10$, which is here equal to M in (16). Meanwhile, the inputs to the system need to be such that

$$-2 \leq u_t \leq 2, \quad \forall t \quad (32)$$

which dictates the set of feasible inputs \mathcal{U} in (20c).

In defining the predictor in (20b), we exploit a set comprising $N = 1000$ input/output samples collected throughout a single noisy open-loop experiment, performed by feeding the system with a random input sequence, uniformly distributed in $[-5, 5]$. Nonetheless, to evaluate the robustness of the approach with respect to different realization of this batch dataset, we carry out all tests over 20 instances³ of \mathcal{D}_N . Fig. 1 shows the performance in closed-loop with on-

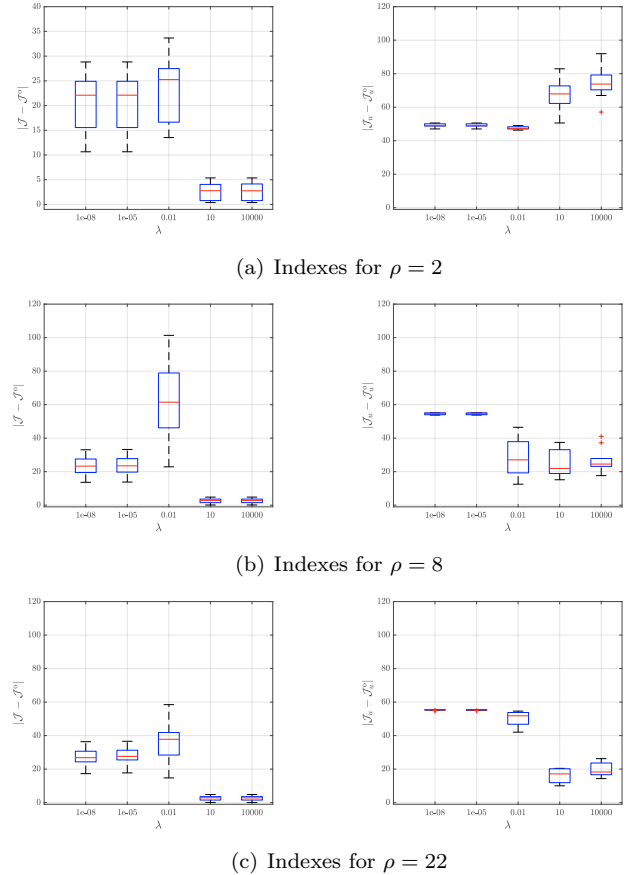


Fig. 2. Benchmark example: oracle indexes and average \mathcal{J} , \mathcal{J}_u over 20 Monte Carlo tests *vs* λ , for increasing ρ .

line noise, when ρ is fixed to 23, *i.e.*, the past horizon minimizing Akaike's *final prediction error* (FPE) over a the set of *autoregressive models with exogenous inputs* (ARX) estimated with the 20 datasets. Despite a variability in the input, due to the unpredictability of future noises, overall the SPC+slack scheme performs comparably to the oracle. This corroborates the discussion in Section 4, showing the advantages of incorporating subspace identification insights within the control design problem. In addition, the results in Fig. 2 highlight the existing intertwining between the choice of λ and ρ . In particular, as soon as ρ is sufficiently high, noise on the online data tends to be correctly handled by the predictive control scheme, without the need for additional slacks.

Flexible transmission. Consider now the stable benchmark single-input, single-output, 5th order, linear time-invariant system of Dörfler et al. (2021). To gather the data used to build the Monte Carlo predictors to be used in (20b), we fed the system with 20 input sequences uniformly generated at random within the interval $[-5, 5]$ of length $N = 125$. Meanwhile, by considering the innovation form (see (7)), for each of the 20 dataset gathered in the data collection phase and when closing the loop, the evolution of the state and the measured output of the system are affected by a different zero mean white noise e_t , with standard deviation 0.01. The gain \mathcal{K} in (7a) has instead been chosen randomly, guaranteeing that $A - \mathcal{K}\mathcal{C}$ has all eigenvalues inside the unit circle. The SPC+ slack scheme is designed by considering the constant output reference

³ We consider 20 realizations of input and noise sequences.

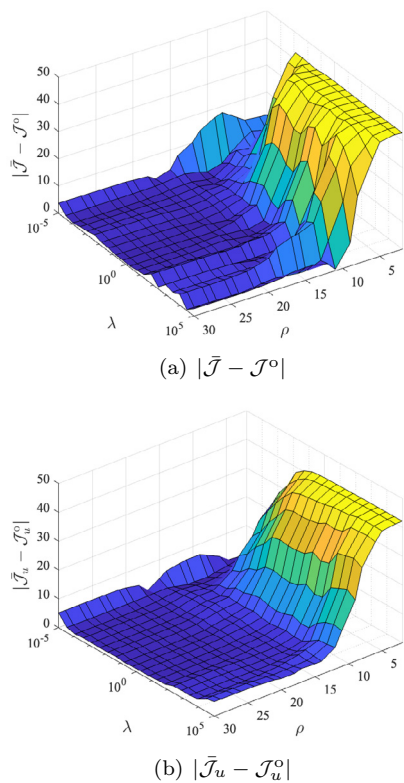


Fig. 3. Flexible transmission: absolute difference between the oracle indexes and the average \mathcal{J} and \mathcal{J}_u over 20 Monte Carlo predictors and simulations vs λ and ρ .




Fig. 4. Flexible transmission: oracle (black dashed line) vs average input/output trajectories (red line) and their standard deviation (shaded area) over 20 Monte Carlo predictors and simulations, $\lambda = 1.8$ and $\rho = 9$.

$r = 1$, a prediction horizon $L = 10$, $Q = 1$ and $R = 10^{-2}$, while the input is constrained to lay within the interval $[-1, 1]$. The SPC+slack scheme is deployed over an horizon of length $T = 50$ time steps for different values of the past horizon ρ and λ . As shown in Fig. 3, the average difference between the SPC+slack indicators in (30) and their oracle counterpart tends to consistently decrease for increasing ρ , while reaching values close to zero for regularization penalties λ that are not excessively high, nor extremely small. These results thus support the conclusions drawn in Section 4. Meanwhile, the effectiveness of the SPC+slack scheme is clearly visible when comparing the closed-loop input/output trajectories attained with this scheme and the oracle ones, as shown in Fig. 3. The past horizon is fixed to the average ρ minimizing Akaike's *final prediction error* (FPE) criterion over the 20 Monte Carlo predictors

in Fig. 4. Note that, in this case, the beneficial effect of slacks can be linked not only to their filtering effect on the incoming data, but also to the size of the dataset, which limits the accuracy of the predictor without slack.

6. CONCLUSIONS

This work provides a deep insight into data-driven predictive control based on subspace identification tools. In particular, we thoroughly elaborate upon the choice of some hyperparameters, which are crucial to attain satisfactory closed-loop performance. Our experimental outcomes show the advantages of exploiting such an insight for tuning the hyperparameters and the possible drawbacks of introducing additional variables and regularization terms.

Future research will be devoted to extending this analysis to other data-driven predictive control schemes.

REFERENCES

- Bauer, D. (2003). Subspace algorithms. In *Proceedings of the IFAC Symposium on System Identification*. Rotterdam, The Netherlands.
- Bemporad, A., Morari, M., Dua, V., and Pistikopoulos, E. (2002). The explicit linear quadratic regulator for constrained systems. *Automatica*, 38(1), 3–20.
- Berberich, J., Köhler, J., Müller, M., and Allgöwer, F. (2020). Data-driven model predictive control with stability and robustness guarantees. *IEEE Transactions on Automatic Control*, 66(4), 1702–1717.
- Breschi, V., Chiuso, A., and Formentin, S. (2023). Data-driven predictive control in a stochastic setting: a unified framework. *Automatica*, 152, 110961.
- Coulson, J., Lygeros, J., and Dörfler, F. (2019). Data-enabled predictive control: In the shallows of the deep. In *2019 18th European Control Conference (ECC)*, 307–312. IEEE.
- Dörfler, F., Coulson, J., and Markovskiy, I. (2021). Bridging direct & indirect data-driven control formulations via regularizations and relaxations. *arXiv preprint arXiv:2101.01273*.
- Favoreel, W., De Moor, B., and Gevers, M. (1999). Spc: Subspace predictive control. 32(2), 4004–4009. 14th IFAC World Congress.
- Fiedler, F. and Lucia, S. (2020). On the relationship between data-enabled predictive control and subspace predictive control. doi:10.48550/ARXIV.2011.13868. URL <https://arxiv.org/abs/2011.13868>.
- Formentin, S., Van Heusden, K., and Karimi, A. (2014). A comparison of model-based and data-driven controller tuning. *International Journal of Adaptive Control and Signal Processing*, 28(10), 882–897.
- Mercère, G. (2013). Regression techniques for subspace-based black-box state-space system identification: an overview. doi:10.48550/ARXIV.1305.7121.
- Rawlings, J. (2000). Tutorial overview of model predictive control. *IEEE Control Systems Magazine*, 20(3), 38–52. doi:10.1109/37.845037.
- Verhaegen, M. and Verdult, V. (2007). *Filtering and system identification: a least squares approach*. Cambridge university press.