

Uncertainty-aware data-driven predictive control in a stochastic setting[★]

V. Breschi^{*}, M. Fabris^{***}, S. Formentin^{**}, A. Chiuso^{***}

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

^{**} *Department of Electronics, Information and Bioengineering, Politecnico di Milano, via G. Ponzio 34/5, Milano, 20133, Italy (e-mail: simone.formentin@polimi.it).*

^{***} *Department of Information Engineering, University of Padova, via Gradeno 6/B, Padua, 35131, Italy. (e-mail: marco.fabris.1@unipd.it, alessandro.chiuso@unipd.it)*

Abstract: Data-Driven Predictive Control (DDPC) has been recently proposed as an effective alternative to traditional Model Predictive Control (MPC), in that the same constrained optimization problem can be addressed without the need to explicitly identify a full model of the plant. However, DDPC is built upon input/output trajectories. Therefore, the finite sample effect of stochastic data, due to, e.g., measurement noise, may have a detrimental impact on closed-loop performance. Exploiting a formal statistical analysis of the prediction error, in this paper we propose the first *systematic approach* to deal with uncertainty due to finite sample effects. To this end, we introduce two regularization strategies for which, differently from existing regularization-based DDPC techniques, we propose a tuning rationale allowing us to select the regularization hyper-parameters *before* closing the loop and *without* additional experiments. Simulation results confirm the potential of the proposed strategy when closing the loop.

Keywords: data-driven control, regularization, predictive control

1. INTRODUCTION

Among advanced control strategies, Model Predictive Control (MPC) is nowadays one of the most widely employed in practice, thanks to its intrinsic ability to handle constraints, time-varying dynamics and multiple (potentially conflicting) objectives, see e.g., Borrelli et al. (2017). Nonetheless, the ultimate performance attained in closed-loop with MPC critically depends on the predictive capabilities of the model featured within its optimization routine. As such, the ability of MPC to deliver the desired control performance might be jeopardized when such a mathematical description of the plant is not accurate enough. This well-known issue, that has led many research efforts towards the development of robust and adaptive MPC solutions (see, e.g., Ding (2017)), is particularly relevant when no mathematical model of the plant is available. In this case, system identification can come of help in allowing one to retrieve an accurate model of the system from data. Alternatively, in such a data-driven context, the unavoidable uncertainty of models can be dealt with by *skipping* an explicit modelling step, using data to directly map the control law.

One of the key ideas to make this shift possible is to think of past input/output records, traditionally used in system identification as training data to learn a parametric

dynamical model, as a *nonparametric* description of its dynamical behavior. This concept is at the core of subspace identification (see Moonen et al. (1989)), as well as behavioral theory in general (see Willems and Polderman (2013)), and of the (deterministic) result in Willems et al. (2005) in particular. More specifically, this last work shows that the future behaviour of a (deterministic) dynamical system can be expressed as the linear combination of a finite set of past trajectories, provided that the input satisfy certain persistency of excitation conditions. This deterministic result has paved the way for the recent developments of data-driven predictive control (DDPC), see, e.g., Coulson et al. (2019) and Berberich et al. (2020). This alternative predictive approach is proven to be equivalent to traditional MPC, if data are collected in a deterministic (noiseless) setting (see Krishnan and Pasqualetti (2021)) and, in special cases, to Subspace Predictive Control, see Favoreel et al. (1999); Fiedler and Lucia (2021); Breschi et al. (2023), while its performance rapidly deteriorates as stochastic data (e.g., noisy data) are used. To make DDPC less sensitive to noise in the data, different forms of regularization have then been embedded within the DDPC scheme, and they have been proven effective in handling noise (see Dorfler et al. (2022) for an overview of possible regularization strategies). Nonetheless, adding regularization terms to the predictive control cost implies that suitable regularization weights must be selected *a priori*, with a non-negligible impact on the final closed-loop performance. It follows that a proper tuning of such penalties could be performed only by means of a subsequent

[★] 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.

validation phase, which must be carried out in closed-loop. These experiments might be unsafe (and often unfeasible) for the plant, as one may even end up de-stabilizing the closed-loop.

Instead of looking at the DDPC design problem from a behavioral perspective, in this paper we look at input/output trajectories from a subspace identification oriented perspective. In particular, building upon the so-called γ -DDPC formulation presented in Breschi et al. (2023), we propose a *systematic framework* to deal with uncertainty in designing data-driven predictive controllers within a stochastic setting. Specifically, by relying on the statistical analysis of the uncertainty in the data-driven predictions, we introduce two regularization schemes to limit mismatches between the true outputs and their prediction and, ultimately, improve closed-loop performance. Differently from existing regularized DDPC approaches, we additionally propose strategies to tune regularization parameters *without* requiring closed-loop experiments. The validity of these procedures for a proper tuning of the overall scheme is shown on a benchmark simulation example.

The remainder of the paper is structured as follows. In Section 2 we initially provide a summary of the main features needed to construct the γ -DDPC scheme proposed in Breschi et al. (2023). Section 3 is then devoted to the formalization of the problem, i.e., the design of uncertainty-aware regularization ingredients for γ -DDPC. The statistical analysis of the data-driven multi-step predictor employed in the considered predictive scheme is provided in Section 4. In light of these results, in Section 5 we propose two alternative tuning policies for the regularization penalties. Their effectiveness is illustrated through a numerical case study in Section 6. The paper is ended by some concluding remarks.

Notation. Given a signal $w(k) \in \mathbb{R}^s$, the associated (block) Hankel matrix $W_{[t_0, t_1], N} \in \mathbb{R}^{s(t_1 - t_0 + 1) \times N}$ is defined as:

$$W_{[t_0, t_1], N} := \frac{1}{\sqrt{N}} \begin{bmatrix} w(t_0) & w(t_0+1) & \cdots & w(t_0+N-1) \\ w(t_0+1) & w(t_0+2) & \cdots & w(t_0+N) \\ \vdots & \vdots & \ddots & \vdots \\ w(t_1) & w(t_1+1) & \cdots & w(t_1+N-1) \end{bmatrix}, \quad (1)$$

while we use the shorthand $W_{t_0} := W_{[t_0, t_0], N}$ to denote a single (block) row Hankel, namely:

$$W_{t_0} := \frac{1}{\sqrt{N}} [w(t_0) \ w(t_0+1) \ \cdots \ w(t_0+N-1)]. \quad (2)$$

2. BACKGROUND

Consider an *unknown* discrete-time, *linear time-invariant* (LTI) stochastic plant \mathcal{S} . Without loss of generality, let \mathcal{S} be described in the so-called minimal (i.e., reachable and observable) *innovation form*, namely

$$\begin{cases} x(t+1) = Ax(t) + Bu(t) + Ke(t) \\ y(t) = Cx(t) + Du(t) + e(t), \end{cases} \quad t \in \mathbb{Z} \quad (3)$$

where $x(t) \in \mathbb{R}^n$, $u(t) \in \mathbb{R}^m$ and $e(t) \in \mathbb{R}^p$ are the state, input and innovation process respectively, while $y(t) \in \mathbb{R}^p$ is the corresponding output signal.

Let us introduce the joint input/output process $z(t)$, given by

$$z(t) := \begin{bmatrix} u(t) \\ y(t) \end{bmatrix}. \quad (4)$$

Given a set of N_{data} input/output pairs and, thus, the sequence $\{z(j)\}_{j=1}^{N_{data}}$, let the associated Hankel matrix be

$$Z_P := Z_{[0, \rho-1], N}, \quad (5)$$

where $N := N_{data} - T - \rho + 1$, T is the “future horizon”, i.e., the prediction horizon when solving a predictive control problem, and ρ is the “past horizon”, shaping the number of past input/output samples used to reconstruct the state at time t , when it is not directly measurable. In addition, let us define the following input and output Hankel matrices:

$$U_F := U_{[\rho, \rho+T-1], N}, \quad Y_F := Y_{[\rho, \rho+T-1], N}. \quad (6)$$

Based on (3), the Hankel of future outputs Y_F can be written as a function of the previous matrices in the form

$$Y_F = \Gamma X_\rho + \mathcal{H}_d U_F + \mathcal{H}_s E_F, \quad (7a)$$

where E_F is the Hankel of future innovations, $\Gamma \in \mathbb{R}^{pT \times n}$ is the extended observability matrix associated with the system, i.e.,

$$\Gamma = \begin{bmatrix} C \\ CA \\ CA^2 \\ \vdots \\ CA^{T-1} \end{bmatrix}, \quad (7b)$$

while $\mathcal{H}_d \in \mathbb{R}^{pT \times mT}$ and $\mathcal{H}_s \in \mathbb{R}^{pT \times pT}$ are the Toeplitz matrices formed with the Markov parameters of the system, namely

$$\mathcal{H}_d = \begin{bmatrix} D & 0 & 0 & \cdots & 0 \\ CB & D & 0 & \cdots & 0 \\ CAB & CB & D & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ CA^{T-2}B & CA^{T-3}B & CA^{T-4}B & \cdots & D \end{bmatrix}, \quad (7c)$$

$$\mathcal{H}_s = \begin{bmatrix} I & 0 & 0 & \cdots & 0 \\ CK & I & 0 & \cdots & 0 \\ CAK & CK & I & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ CA^{T-2}K & CA^{T-3}K & CA^{T-4}K & \cdots & I \end{bmatrix}. \quad (7d)$$

Let us additionally define \hat{Y}_F as the orthogonal projection of Y_F onto the row space of Z_P and U_F , i.e.,

$$\hat{Y}_F = \Gamma \hat{X}_\rho + \mathcal{H}_d U_F + \underbrace{\mathcal{H}_s \Pi_{Z_P, U_F}(E_F)}_{O_P(1/\sqrt{N})} \quad (8)$$

where the last term vanishes¹ (in probability) as $1/\sqrt{N}$.

When the matrices (A, B, C, D, K) are *unknown*, future outputs can still be predicted from the Hankel matrices in (5)-(6). Indeed, defining the future input and output vectors

$$u_f := \begin{bmatrix} u(t) \\ u(t+1) \\ \vdots \\ u(t+T-1) \end{bmatrix}, \quad y_f := \begin{bmatrix} y(t) \\ y(t+1) \\ \vdots \\ y(t+T-1) \end{bmatrix}, \quad (9)$$

¹ For a more formal statement on this, we refer the reader to standard literature on subspace identification.

Z_P, U_F and Y_F can be used in a deterministic setting to predict, given past data and future inputs, the future outputs y_f as follows:

$$\begin{bmatrix} z_{init} \\ u_f \\ y_f \end{bmatrix} = \begin{bmatrix} Z_P \\ U_F \\ Y_F \end{bmatrix} \alpha, \quad \text{where } z_{init} := \begin{bmatrix} z(t-\rho) \\ \vdots \\ z(t-2) \\ z(t-1) \end{bmatrix}, \quad (10)$$

with $\alpha \in \mathbb{R}^N$ being the variable typically optimized in DDPC problems stemming from Willems et al. (2005), see e.g., Berberich et al. (2020) and Coulson et al. (2019).

Following the same rationale of Breschi et al. (2023), we instead reformulate the previous relationship through the LQ decomposition of the joint input-output block Hankel matrix:

$$\begin{bmatrix} Z_P \\ U_F \\ Y_F \end{bmatrix} = \begin{bmatrix} L_{11} & 0 & 0 \\ L_{21} & L_{22} & 0 \\ L_{31} & L_{32} & L_{33} \end{bmatrix} \begin{bmatrix} Q_1 \\ Q_2 \\ Q_3 \end{bmatrix}, \quad (11)$$

where the matrices $\{L_{ii}\}_{i=1}^3$ are all non-singular and Q_i have orthonormal rows, i.e., $Q_i Q_i^\top = I$, for $i = 1, \dots, 3$, $Q_i Q_j^\top = 0$, $i \neq j$. Combining (10) with (11), we can further retrieve the following relationship:

$$\begin{bmatrix} z_{init} \\ u_f \\ y_f \end{bmatrix} = \begin{bmatrix} Z_P \\ U_F \\ Y_F \end{bmatrix} \alpha = \begin{bmatrix} L_{11} & 0 & 0 \\ L_{21} & L_{22} & 0 \\ L_{31} & L_{32} & L_{33} \end{bmatrix} \underbrace{\begin{bmatrix} Q_1 \\ Q_2 \\ Q_3 \end{bmatrix}}_{\gamma} \alpha. \quad (12)$$

This relation allows us to establish a connection between the standard optimization variable of DDPC strategies α , and the new parameters

$$\gamma = [\gamma_1^\top \ \gamma_2^\top \ \gamma_3^\top]^\top, \quad (13)$$

which is the starting point for the derivation of the γ -DDPC scheme proposed in Breschi et al. (2023), and that is at the core of this work.

3. PROBLEM SETTING

Consider now the predictive control problem designed for the outputs of the system to track a given reference $y_r(t)$, while satisfying the constraints encoded into the sets \mathcal{U} , \mathcal{Y} . This control problem can be cast as

$$\underset{\{u(k)\}_t^{t+T-1}}{\text{minimize}} \frac{1}{2} \sum_{k=t}^{t+T-1} \|\hat{y}(k) - y_r(k)\|_Q^2 + \|u(k)\|_R^2 \quad (14a)$$

$$\text{s.t. } \hat{x}(k+1) = A\hat{x}(k) + Bu(k), \quad k \in [t, t+T), \quad (14b)$$

$$\hat{y}(k) = C\hat{x}(k) + Du(k), \quad k \in [t, t+T), \quad (14c)$$

$$\hat{x}(t) = x_{init}, \quad (14d)$$

$$u(k) \in \mathcal{U}, \quad \hat{y}(k) \in \mathcal{Y}, \quad k \in [t, t+T), \quad (14e)$$

where $k \in \mathbb{Z}$, x_{init} is the state at time t , \hat{x} and \hat{y} are the *model-based estimates* of the deterministic components of the states and outputs, while the penalties $Q \in \mathbb{R}^{p \times p}$ and $R \in \mathbb{R}^{m \times m}$, with $Q \succeq 0$ and $R \succ 0$, are selected to trade-off between tracking performance and control effort.

Let us now assume that the system matrices (A, B, C, D, K) are *unknown*, while we have access to a sequence of input/output data $\mathcal{D}_{Ndata} = \{u(j), y(j)\}_{j=1}^{Ndata}$. Within this context, a data-driven predictive controller with the same objectives and constraints of (14) can be formulated as follows

$$\min_{\gamma_2, \gamma_3} \frac{1}{2} \sum_{k=t}^{t+T-1} \ell(u(k), \hat{y}(k), y_r(k)) + \Psi(\gamma_2, \gamma_3) \quad (15a)$$

$$\text{s.t. } \begin{bmatrix} u_f \\ \hat{y}_f \end{bmatrix} = \begin{bmatrix} L_{21} & L_{22} & 0 \\ L_{31} & L_{32} & L_{33} \end{bmatrix} \begin{bmatrix} \gamma_1^* \\ \gamma_2 \\ \gamma_3 \end{bmatrix}, \quad (15b)$$

$$u(k) \in \mathcal{U}, \quad \hat{y}(k) \in \mathcal{Y}, \quad k \in [t, t+T), \quad (15c)$$

with

$$\ell(u(k), \hat{y}(k), y_r(k)) = \|\hat{y}(k) - y_r(k)\|_Q^2 + \|u(k)\|_R^2, \quad (16)$$

$$\gamma_1^* = L_{11}^{-1} z_{init} \quad (17)$$

where z_{init} is defined as in (9). Note that, this formulation matches most² of the regularized DDPC schemes proposed in Dorfler et al. (2022), each based on a specific choice of³ $\Psi(\gamma_2, \gamma_3)$. Nonetheless, by decoupling α in the three components γ_1, γ_2 and γ_3 , this formulation turns out to be more convenient when discussing (and tuning) regularization.

The regularization term $\Psi(\gamma_2, \gamma_3)$ in (15a) is rather critical to design. Indeed, even in the presence of a small amount of noise, closed-loop performance might dramatically change depending on the chosen $\Psi(\gamma_2, \gamma_3)$. In the most extreme cases, one may go from not controlling the system at all (i.e., let it evolve in open loop) to overfitting noise, see e.g., Dorfler et al. (2022). Based on these considerations, our ultimate goal is *to provide a systematic approach for the design of the last term in (15a) within our stochastic framework, while avoiding the need for additional experiments and closed-loop tuning tests.*

Note that some observations on the problem in (15) have already been made in Breschi et al. (2023), where it is already argued that, for large N , the optimal choice is to set $\gamma_3 = 0$ and remove any regularization from γ_2 , so that problem (15) becomes:

$$\min_{\gamma_2} \frac{1}{2} \sum_{k=t}^{t+T-1} \|\hat{y}_0(k) - y_r(k)\|_Q^2 + \|u(k)\|_R^2 \quad (18a)$$

$$\text{s.t. } \begin{bmatrix} u_f \\ \hat{y}_{0,f} \end{bmatrix} = \begin{bmatrix} L_{21} & L_{22} \\ L_{31} & L_{32} \end{bmatrix} \begin{bmatrix} \gamma_1^* \\ \gamma_2 \end{bmatrix}, \quad (18b)$$

$$u(k) \in \mathcal{U}, \quad \hat{y}_0(k) \in \mathcal{Y}, \quad k \in [t, t+T), \quad (18c)$$

Under this choice ($\gamma_3 = 0$) and with no regularization on γ_2 , the predicted output $\hat{y}_{0,f} = L_{31}\gamma_1^* + L_{32}\gamma_2$ can further be written (see Breschi et al. (2023)) in the form:

$$\hat{y}_{0,f} = [L_{31} \ L_{32}] \begin{bmatrix} \gamma_1^* \\ \gamma_2 \end{bmatrix} = \hat{Y}_F \alpha \quad (19)$$

where the last equation exploits the fact that the projected future output \hat{Y}_F can be written in terms of the LQ decomposition (12) as:

$$\hat{Y}_F = [L_{31} \ L_{32}] \begin{bmatrix} Q_1 \\ Q_2 \end{bmatrix}.$$

² Some of the schemes in Dorfler et al. (2022) require also a regularization on γ_1 , which instead in this work is fixed based on (17). See Breschi et al. (2023) for further discussion about this issue.

³ Based on the relationship in (12), $\Psi(\gamma_2, \gamma_3)$ is indeed equivalent to a regularization on α .

4. FINITE SAMPLE UNCERTAINTY OF DATA-DRIVEN PREDICTORS

Exploiting the relation (8), the predicted output $\hat{y}_{0,f}$ in (19) is subject to $O_p(1/\sqrt{N})$ perturbations that are due to the projection residuals of the future innovations E_f onto the joint past Z_p and future input U_f spaces.

More precisely, denoting with \hat{y}_f^* the ‘‘true’’ output predictor corresponding to the given initial conditions and inputs, $\hat{y}_{0,f}$ satisfies the relation

$$\underbrace{\hat{y}_{0,f}}_{=Y_f\alpha} = \underbrace{\hat{y}_f^*}_{=[\Gamma\hat{X}_p+\mathcal{H}_dU_F]\alpha} + \mathcal{H}_s \underbrace{\Pi_{Z_p,U_f}(E_f)}_{\tilde{e}_f} \alpha,$$

Defining the prediction error $\tilde{y}_f := \hat{y}_f^* - \hat{y}_{0,f}$, from the previous equation we see that $\tilde{y}_f = -\mathcal{H}_s \tilde{e}_f$. If \tilde{y}_f were equal to zero, then the optimal control problem (18) would coincide with the *oracle* model based predictive control problem, i.e., the optimal MPC using the true model of the system (14).

For future use, let us observe that \tilde{e}_f can be written in the form:

$$\tilde{e}_f := \Pi_{Z_p,U_f}(E_f)\alpha = E_f \begin{bmatrix} Q_1^\top & Q_2^\top \end{bmatrix} \begin{bmatrix} Q_1 \\ Q_2 \end{bmatrix} \alpha.$$

Denoting with $e_f(t)$ the t -th column of E_f and with $q(t)$ the t -th column of

$$\sqrt{N} \begin{bmatrix} Q_1 \\ Q_2 \end{bmatrix} := \begin{bmatrix} L_{11} & 0 \\ L_{21} & L_{22} \end{bmatrix}^{-1} \sqrt{N} \begin{bmatrix} Z_p \\ U_f \end{bmatrix},$$

we can rewrite $\sqrt{N}\tilde{e}_f$ in the form:

$$\sqrt{N}\tilde{e}_f = \frac{1}{\sqrt{N}} \sum_{t=1}^N e_f(t)q(t)^\top \underbrace{\begin{bmatrix} \gamma_1 \\ \gamma_2 \end{bmatrix}}_{:=\gamma_{12}}.$$

The following proposition characterizes the statistical properties of \tilde{e}_f , and is the core result that will be used in the next section to design data-driven tuning strategies for regularization in DDPC.

Proposition 1. Under the assumption that the innovation process $e(t)$ in (3) is, conditionally on the joint input-output past data $\{y(s), u(s), s < t\}$, a martingale difference sequence with constant conditional variance, i.e.,

$$\begin{aligned} \mathbf{E}[e(t)|y(s), u(s), s < t] &= 0, \\ \text{Var}[e(t)|y(s), u(s), s < t] &= \text{Var}[e(t)] = \sigma^2, \end{aligned}$$

then

$$\mathbf{E}[\sqrt{N}\tilde{e}_f] \xrightarrow{N \rightarrow \infty} 0,$$

and

$$\text{Var}[\sqrt{N}\tilde{e}_f] \xrightarrow{N \rightarrow \infty} \sum_{k=-T}^T \sigma^2 J_k \frac{(N-|k|)}{N} \gamma_{12}^\top \Sigma_q^\top(k) \gamma_{12}, \quad (20)$$

where $\Sigma_q(k)$ is the covariance matrix $\mathbf{E}[q(t+k)q^\top(t)]$ and J_i is the shift matrix such that $[J_i]_{h,k} = 0$ for $k-h \neq i$ and $[J_i]_{h,k} = 1$ for $k-h = i$ (i.e., with zeros everywhere except for ones on the superdiagonal ($i > 0$) or subdiagonal ($i < 0$)).

The covariance matrices $\Sigma_q(k)$ in (20) can be estimated from data. However, it is easy to prove that, asymptoti-

cally in N , $\Sigma_q(0) = I$. Therefore, exploiting the fact that $\text{Trace}[J_i] = 0$ for $i \neq 0$, we have that

$$\text{Trace}[\text{Var}[\sqrt{N}\tilde{e}_f]] \xrightarrow{N \rightarrow \infty} T\sigma^2 \|\gamma_{12}\|^2. \quad (21)$$

This relation will be extremely useful in the next section because it connects the average scalar variance of \tilde{e}_f to the optimization variable γ . In particular, if the control horizon T is large enough, one expects that $\frac{\|\tilde{e}_f\|^2}{T}$ can be seen as a sample estimate of the average variance $\frac{\text{Trace}[\text{Var}[\tilde{e}_f]]}{T}$ of the components of the vector \tilde{e}_f , so that:

$$\frac{\|\tilde{e}_f\|^2}{T} \simeq \frac{\text{Trace}[\text{Var}[\tilde{e}_f]]}{T} \simeq \frac{\sigma^2 \|\gamma_{12}\|^2}{N}.$$

We conclude this section discussing how the connection between the prediction error \tilde{y}_f and the vector \tilde{e}_f , given by $\tilde{y}_f = -\mathcal{H}_s \tilde{e}_f$ can be estimated in a model-free, data driven fashion. The result is formalized in the following lemma.

Lemma 2. Consider the LQ decomposition in (12). Then

$$\lim_{N \rightarrow \infty} L_{33} L_{33}^\top = \sigma^2 \mathcal{H}_s \mathcal{H}_s^\top. \quad (22)$$

The previous lemma shows that the matrix L_{33} can be seen as a (nonparametric/model-free) estimate of $\sigma \mathcal{H}_s$. Exploiting this fact we obtain the fundamental result of this section that will be used later on in the design of our tuning strategy.

Proposition 3. The prediction error \tilde{y}_f can be written as

$$\tilde{y}_f = -\mathcal{H}_s \tilde{e}_f \simeq -L_{33} \frac{\tilde{e}_f}{\sigma} = L_{33} \tilde{\gamma}_3,$$

where $\tilde{\gamma}_3 := -\frac{\tilde{e}_f}{\sigma}$ satisfies

$$\|\tilde{\gamma}_3\|^2 \simeq \text{Trace}[\text{Var}[\tilde{\gamma}_3]] \simeq \frac{T\|\gamma_{12}\|^2}{N}.$$

5. REGULARIZATION DESIGN

In Section 4, we have seen that (i) the predictor \hat{y}_f used in the data-driven predictive control design scheme (18) is affected by an uncertainty \tilde{y}_f , (ii) how the latter can be (statistically) characterized, and (iii) how it can be expressed in a data-driven (model-free) fashion. Based on these considerations, in the next two subsections, we introduce two regularization strategies to mitigate the effect of the prediction error \tilde{y}_f on the control performance.

5.1 Regularizing γ_2

First of all let us observe that the average scalar variance $\text{Trace}[\text{Var}[\gamma_3]]$ of the vector $\gamma_3 := L_{33}^{-1} \tilde{y}_f$ scales linearly with squared norm of the optimization parameter $\|\gamma_{12}\|^2 = \|\gamma_1\|^2 + \|\gamma_2\|^2$. The vector γ_1 is fixed in the optimization problem (15) to the value γ_1^* that matches the initial conditions, whereas γ_2 is optimized to achieve the control goal. Should $\|\gamma_2\|$ grow, then also the variance of the prediction error would increase, possibly jeopardizing the closed loop performance. Thus, it is desirable to regularize the control problem (15) by adding a term of the form $\Psi(\gamma_2, \gamma_3) = \beta_2 \|\gamma_2\|^2$ while constraining $\gamma_3 = 0$, i.e.,

6. NUMERICAL EXAMPLES

$$\min_{\gamma_2} \frac{1}{2} \sum_{k=t}^{t+T-1} \ell(u(k), \hat{y}_0(k), y_r(k)) + \beta_2 \|\gamma_2\|^2 \quad (23a)$$

$$\text{s.t.} \quad \begin{bmatrix} u_f \\ \hat{y}_{0,f} \end{bmatrix} = \begin{bmatrix} L_{21} & L_{22} \\ L_{31} & L_{32} \end{bmatrix} \begin{bmatrix} \gamma_1^* \\ \gamma_2 \end{bmatrix}, \quad (23b)$$

$$u(k) \in \mathcal{U}, \hat{y}_0(k) \in \mathcal{Y}, k \in [t, t+T), \quad (23c)$$

where $\ell(u(k), \hat{y}_0(k), y_r(k))$ is defined as in (16). This scheme is designed so as to keep the norm of γ_2 (and thus the variance of the prediction error) small. In particular, we would like to avoid situations in which the control input is (erroneously) exploiting prediction errors to make $\hat{y}_{0,f}$ (too) close to the reference trajectory $y_r(t)$ (indeed, there are no reasons to fit below noise level). Thus, it makes sense to chose β_2 large enough so that

$$\|L_{33}^{-1}(\hat{y}_{0,f} - y_r)\|^2 \simeq \|L_{33}^{-1}\tilde{y}_f\|^2 \simeq T \frac{\|\gamma_1^*\|^2 + \|\gamma_2^*(\beta_2)\|^2}{N}, \quad (24)$$

where $\gamma_2^*(\beta_2)$ denotes the optimal parameter γ_2 that solve (23) as a function of β_2 . As it will be shown in Section 6, this reduces to a linear search problem that can be solved *before* actually closing the loop.

5.2 Slack on output prediction

As an alternative, Proposition 3 suggests that the “true” predictor \hat{y}_f^* can be written as

$$\hat{y}_f^* = \hat{y}_{0,f} + L_{33}\tilde{\gamma}_3,$$

where $\tilde{\gamma}_3$ depends on the “true” but unknown noise. This suggests that the effect of noise can be compensated for by adding a slack of the form $\xi_f := L_{33}\tilde{\gamma}_3$ with $\tilde{\gamma}_3$ to be optimized. We do so by defining

$$\hat{y}_f := \hat{y}_{0,f} + \xi_f = L_{11}\gamma_1^* + L_{22}\gamma_2 + L_{33}\gamma_3,$$

where γ_2 and γ_3 are optimization variables. In this case, introducing some slack has the effect of avoiding an unnecessarily large control effort, thus inducing an implicit regularization on γ_2 . The “size” of the slack variable can be controlled by regularizing the norm of γ_3 , adding a regularization term $\Psi(\beta_2, \beta_3) := \beta_3\|\gamma_3\|^2$ in (15). The resulting optimization problem (changing sign to the optimization variable γ_3 , which of course does not affect the result) takes the form

$$\min_{\gamma_2, \gamma_3} \frac{1}{2} \sum_{k=t}^{t+T-1} \ell(u(k), \hat{y}(k), y_r(k)) + \beta_3\|\gamma_3\|^2 \quad (25a)$$

$$\text{s.t.} \quad \begin{bmatrix} u_f \\ \hat{y}_f \end{bmatrix} = \begin{bmatrix} L_{21} & L_{22} & 0 \\ L_{31} & L_{32} & L_{33} \end{bmatrix} \begin{bmatrix} \gamma_1^* \\ \gamma_2 \\ \gamma_3 \end{bmatrix}, \quad (25b)$$

$$u(k) \in \mathcal{U}, \hat{y}(k) \in \mathcal{Y}, k \in [t, t+T), \quad (25c)$$

with $\ell(u(k), \hat{y}(k), y_r(k))$ given by (16). Also in this case β_3 can be tuned, via a linear search, to be small enough⁴ so as to guarantee that

$$\|\gamma_3^*(\beta_3)\|^2 \simeq T \frac{\|\gamma_1^*\|^2 + \|\gamma_2^*(\beta_3)\|^2}{N}, \quad (26)$$

where $\gamma_2^*(\beta_3)$ and $\gamma_3^*(\beta_3)$ denote the optimal parameter γ_2 and γ_3 that solve (25) as a function of β_3 .

⁴ Note that here $\beta_3 \rightarrow \infty$ implies that $\gamma_3 = 0$ and thus no slack would be introduced.

We now report the results of some numerical simulations to illustrate our theoretical findings, by considering the benchmark single-input, single-output, 4-th order, linear time-invariant system in Landau et al. (1995). Similarly to Dorfler et al. (2022), we collect one noise-free input/output time series of length $N_{data} = 250$, by applying a random Gaussian input of unitary variance. From this noise-free data set, $n_{MC} = 1000$ independent noisy data sets are constructed by adding Gaussian noise with signal-to-noise ratio of 13 dB, to be used as training data for the closed loop experiments. The latter are carried out in a *noisy* scenario, i.e., wit noise added to the output and thus fed back in the loop.

In the experiments the output reference⁵ is $y_r(t) = \sin(5\pi t/(T + T_v - 1))$ the prediction horizon $T = 20$ and the number of feedback steps is $T_v = 50$. The performance index we use is chosen as

$$J_a = T_v^{-1} \sum_{t=0}^{T_v-1} \|y(t) - y_r(t)\|_Q^2 + \|u(t)\|_R^2, \quad (27)$$

where $(Q, R) = (qI_p, rI_m)$, with $(q, r) = (2000, 0.01)$ and $m = p = 1$. The subscript $a = 2$ refers to schemes regularizing γ_2 as discussed in Section 5.1, while $a = 3$ is associated with approaches tuning γ_3 (see Section 5.2). To benchmark closed-loop performance, we consider a Kalman-filter-based model predictive scheme that exploits the true model parameters, and denote its performance index as J_a , $a = KF$. We also report the corresponding control effort

$$J_{u,a} = T_v^{-1} \sum_{t=0}^{T_v-1} \|u(t)\|^2,$$

and relative tracking error

$$J_{y,a} = \left(\sum_{t=0}^{T_v-1} \|y(t) - y_r(t)\|^2 \right) / \left(\sum_{t=0}^{T_v-1} \|y_r(t)\|^2 \right),$$

where a is either 2, 3, or KF .

We benchmark our tuning strategies for β_2 and β_3 against an oracle that exploits an estimate of the close-loop cost obtained as follows by running 1000 closed-loop Monte Carlo trials and then averaging: $J_{AV,a} = n_{MC}^{-1} \sum_{j=1}^{n_{MC}} J_a^{(j)}$. $J_a^{(j)}$ denotes the j -th realization of the cost J_a and it is computed for β_a ($a = 2, 3$) ranging on a logarithmically spaced grid G_a of $|G_a| = 200$ fixed points for $\beta_2 \in G_2 \subset [10^0, 10^4]$ and $\beta_3 \in G_3 \subset [10^{-4}, 10^0]$. The corresponding minimizers are denoted as β_a , $a \in \{2, 3\}$ (see Fig. 1(a))⁶.

Closed loop performance are evaluated on n_{MC} Monte Carlo experiments with the following four configurations.

- (a) The regularized problem in Section 5.1 with $\beta_2(t) = \bar{\beta}_2$. The corresponding costs are denoted as \bar{J}_2 .
- (b) The regularized problem in Section 5.2 with $\beta_3(t) = \bar{\beta}_3$. The corresponding costs are denoted as \bar{J}_3 .
- (c) The regularized problem in Section 5.1 is solved with $\beta_2(t)$ tuned at each closed loop step enforcing condition (24). The corresponding costs are denoted as \hat{J}_2 .

⁵ The data-driven predictive controller is designed with preview of the reference to be tracked.

⁶ Note that the values resulting from this procedure cannot be retrieved in practice without performing a set of (possibly unsafe) closed loop experiments.

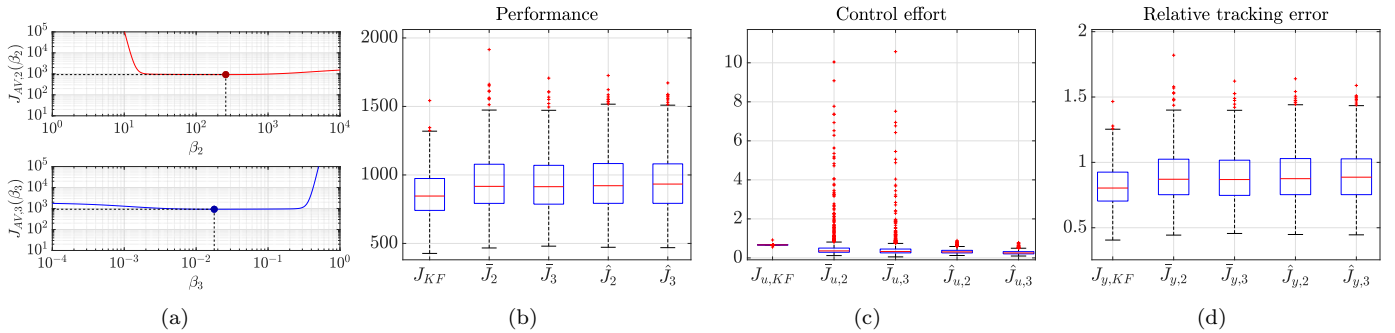


Fig. 1. (a): Oracle tuning: average closed-loop costs $J_{AV,a}$ and their minimizers $\bar{\beta}_a$, $a = 2, 3$. (b),(c),(d): Distributions of the closed loop performance, control effort and relative tracking error over 1000 Monte Carlo runs. Costs with bars refer to “oracle” tuning whereas costs with hats refer to online (feasible) strategy.

(d) The regularized problem in Section 5.2 with $\beta_3(t)$ tuned at each close loop step enforcing condition (26). The corresponding costs are denoted as \hat{J}_3 .

In the sequel, we refer to the selection $\beta_a(t) = \bar{\beta}_a$ in (a) and (b) as the “offline” tuning strategies; whereas, (c) and (d) are the “online” tuning approaches (as they do not require additional closed-loop experiments) proposed in this paper.

Fig. 1(b)-1(d) show the boxplots n_{MC} Monte Carlo runs of the realized closed loop costs. Remarkably, the offline selection of $\beta_a(t) = \bar{\beta}_a$ and the online strategies using (24) and (26) perform comparably.

Furthermore, Fig. 1(a) confirms that, on average, the regularizing role of the parameters β_2 and β_3 is crucial for the stability (and optimal performance) of the closed loop. It is worth observing that β_2 and β_3 play dual roles. Indeed, $\beta_2 \rightarrow 0$ and $\beta_3 \rightarrow +\infty$ correspond to no regularization on the control problem, whereas $\beta_2 \rightarrow +\infty$ and $\beta_3 \rightarrow 0$ correspond to “maximal” regularization. Lastly, from Fig. 1(a) it is clear that the closed-loop control cost diverges (a clear sign of instability), when no regularization is performed, i.e., $\beta_2 \rightarrow 0$ or $\beta_3 \rightarrow +\infty$.

7. CONCLUSIONS

Leveraging the statistical analysis of the (non-parametric) predictor used in data-driven, model-free, predictive control problems, in this paper we have proposed two regularization approaches to account for finite sample effects in the design of data-driven predictive controllers within a stochastic setting. We have also discussed corresponding online tuning strategies for the selection of the regularization penalties. The proposed tuning rationale allows for the design of the controller without the need for additional closed-loop experiments.

Simulation results confirm the effectiveness of the online strategies in face of uncertainties, showing that their performance is practically indistinguishable from an oracle-type tuning based on offline closed-loop experiments. Future work will include a thorough evaluation of the proposed on-line tuning strategy, as well as a formal analysis of the closed loop stability.

REFERENCES

Berberich, J., Köhler, J., Müller, M.A., and Allgöwer, F. (2020). Data-driven model predictive control with

stability and robustness guarantees. *IEEE Transactions on Automatic Control*, 66(4), 1702–1717.

Borrelli, F., Bemporad, A., and Morari, M. (2017). *Predictive control for linear and hybrid systems*. Cambridge University Press.

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.

Ding, B. (2017). Robust and adaptive model predictive control of nonlinear systems. *IEEE Control Systems Magazine*, 37(1), 125–127.

Dorfler, F., Coulson, J., and Markovskiy, I. (2022). Bridging direct & indirect data-driven control formulations via regularizations and relaxations. *IEEE Transactions on Automatic Control*.

Favoreel, W., Moor, B.D., and Gevers, M. (1999). Spc: Subspace predictive control. *IFAC Proceedings Volumes*, 32(2), 4004–4009. 14th IFAC World Congress 1999, Beijing, Chia, 5-9 July.

Fiedler, F. and Lucia, S. (2021). On the relationship between data-enabled predictive control and subspace predictive control. In *2021 European Control Conference (ECC)*, 222–229.

Krishnan, V. and Pasqualetti, F. (2021). On direct vs indirect data-driven predictive control. In *2021 60th IEEE Conference on Decision and Control (CDC)*, 736–741. IEEE.

Landau, I., Rey, D., Karimi, A., Voda, A., and Franco, A. (1995). A flexible transmission system as a benchmark for robust digital control. *European Journal of Control*, 1(2), 77–96.

Moonen, M., De Moor, B., Vandeberghe, L., and Vandewalle, J. (1989). On- and off-line identification of linear state-space models. *Int. J. of Control*, 49(1), 219–232.

Willems, J.C., Rapisarda, P., Markovskiy, I., and De Moor, B.L. (2005). A note on persistency of excitation. *Systems & Control Letters*, 54(4), 325–329.

Willems, J. and Polderman, J. (2013). *Introduction to Mathematical Systems Theory: A Behavioral Approach*. Texts in Applied Mathematics. Springer New York. URL <https://books.google.it/books?id=qoLSBwAAQBAJ>.