2023 62nd IEEE Conference on Decision and Control (CDC)
December 13-15, 2023. Marina Bay Sands, Singapore

On the impact of regularization in data-driven predictive control

Valentina Breschi *Member, IEEE,* Alessandro Chiuso *Fellow, IEEE,*
Marco Fabris and Simone Formentin *Member, IEEE,*

***Abstract*— Model predictive control (MPC) is a control strategy widely used in industrial applications. However, its implementation typically requires a mathematical model of the system being controlled, which can be a time-consuming and expensive task. Data-driven predictive control (DDPC) methods offer an alternative approach that does not require an explicit mathematical model, but instead optimize the control policy directly from data. In this paper, we study the impact of two different regularization penalties on the closed-loop performance of a recently introduced data-driven method called γ-DDPC. Moreover, we discuss the tuning of the related coefficients in different data and noise scenarios, to provide some guidelines for the end user.**

***Index Terms*— Data driven control, Predictive control for linear systems, Uncertain systems**

I. Introduction

Model predictive control (MPC) is a popular control strategy that has been successfully applied in a wide range of applications [1]. However, a major limitation of MPC is that it requires a mathematical model of the system being controlled, which can be a costly and time-consuming task. This requirement has led to the development of data-driven predictive control (DDPC) methods, which aim to learn the control policy directly from data without the need for a mathematical model of the plant [2] [3]–[5].

Nonetheless, the data-based predictor used in DDPC is not exempt from shortcomings, due to the presence of noise on the measured data. Therefore, different techniques have been proposed to make the closed-loop performance less sensitive to such a noise, e.g., robust design in case hard power bounds are given [3], dynamic mode decomposition [6] and regularization [7]. The latter in particular can be used to prevent the data-based predictor to overfit the historical data, by tuning a few penalty coefficients. In the pioneering work [7], the design of such terms is discussed for different kinds of regularization, and the authors highlight the significant efforts required in terms of trial-and-error tuning, especially as far as some specific parameters are concerned. In [8], we showed that regularization may be avoided in case the data set is large enough and the DDPC problem is reformulated thanks to subspace identification tools, so as to shrink the number of decision variables, into the so-called γ-DDPC method. Finally, in [9], we have focused on finite size data sets and used asymptotic arguments to show that regularization might instead be useful to counteract the prediction error variance, due to the use of noisy data in the predictor. Two different regularization options have been introduced, and an on-line tuning of the associated penalizations has been proposed, based on the prior knowledge of the variance expression.

This paper's contribution is built upon [9], since our goal here is to analyze the joint tuning of the two regularization terms of γ-DDPC and analyze their impact on the closed-loop performance. In particular, we shall discuss the role of the driving input color (spectra) and some qualitative guidelines about regularization design will be drawn by means of extensive simulations on a benchmark linear system as well as on a challenging nonlinear problem, namely, wheel slip control in braking maneuvering. Finally, offline and on-line regularization tuning will be compared.

The remainder of the paper is as follows. In Section II, the predictive control problem setting is described, and the regularization tuning issue is mathematically formulated. Section III illustrates the considered regularization techniques for γ-DDPC and discusses the role of each term, also by means of two numerical case studies. The paper is ended by some concluding remarks.

Notation. Given a signal (say $u(t) \in \mathbb{R}^m$), the associated (block) Hankel matrix $U_{[t_0,t_1],N} \in \mathbb{R}^{m(t_1-t_0+1)\times N}$ is defined as:

$$U_{[t,s],N} := \frac{1}{\sqrt{N}} \begin{bmatrix} u(t) & u(t+1) & \cdots & u(t+N-1) \\ u(t+1) & u(t+2) & \cdots & u(0+N) \\ \vdots & \vdots & \ddots & \vdots \\ u(s) & u(s+1) & \dots & u(s+N-1) \end{bmatrix}, \tag{1}$$

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

$$U_t := \frac{1}{\sqrt{N}} \begin{bmatrix} u(t) & u(t+1) & \cdots & u(t+N-1) \end{bmatrix}. \tag{2}$$

II. Problem setting

Our goal is to design a controller for an *unknown* plant that can be modeled by *linear time-invariant* (LTI) discrete-time linear (stochastic) system $\mathcal{S}$. Without loss of generality, we consider its state space description in *innovation form*

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

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.

A. Chiuso and M. Fabris are with the Department of Information Engineering, University of Padova, Padua, Italy. V. Breschi is with Department of Electrical Engineering, Eindhoven University of Technology, Eindhoven, Netherlands. S. Formentin is with Dipartimento di Elettronica, Informatica e Bioingegneria (DEIB), Politecnico di Milano, Milan, Italy.

Corresponding author: S. Formentin, simone.formentin@polimi.it

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.

Under the *unrealistic* assumption that the system matrices (A, B, C, D, K) are known, the predictive constrained tracking control problem of interest for this paper (for a given reference $y_r(t)$ and a prediction horizon T) can be formulated as follows

$$\underset{\{u(k)\}_t^{t+T-1}}{\text{minimize}} \ \frac{1}{2} \sum_{k=t}^{t+T-1} \ell(u(k), \hat{y}(k), y_r(k)) \tag{4a}$$

$$\text{s.t. } \hat{x}(k+1) = A\hat{x}(k) + Bu(k), \ k \in [t, t+T), \tag{4b}$$

$$\hat{y}(k) = C\hat{x}(k) + Du(k), \ k \in [t, t+T), \tag{4c}$$

$$\hat{x}(t) = \hat{x}_{init}, \tag{4d}$$

$$u(k) \in \mathcal{U}, \ \hat{y}(k) \in \mathcal{Y}, \ k \in [t, t+T), \tag{4e}$$

where $k \in \mathbb{Z}$, $\hat{x}_{init}$ is the state-estimate at time t, which can be obtained by running a conventional Kalman filter given the input-output measurements available up to time t, and y_r is the reference signal. Also, $\ell(\cdot)$ is a convex loss function, penalizing both the tracking performance and the control effort, *e.g.,*

$$\ell(u(k), \hat{y}(k), y_r(k)) = \|\hat{y}(k) - y_r(k)\|_Q^2 + \|u(k) - u_r(k)\|_R^2, \tag{4f}$$

where 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.

A standard assumption in *data-driven* control is that the system matrices (A, B, C, D, K) are *not known*, and only a finite sequence of input/output data $\mathcal{D}_{N_{data}} = \{u(j), y(j)\}_{j=1}^{N_{data}}$. We would like to stress that in our framework measured data are by assumption noisy, in the sense that there is no LTI system that, with the given input $u(t)$, produces exactly the measured output in a deterministic way.

In this paper, we follow that data-driven predictive control problem formulation provided in [9], [10], and we refer to those papers for a connection with the recent related literature such as [3], [7].

To this purpose, we need to introduce the Hankel matrices, including past and future values of inputs and outputs, with respect to time t. In particular, with obvious use of the subscripts P and F, we define:

$$U_F := U_{[\rho, \rho+T-1], N}, \ Y_F := Y_{[\rho, \rho+T-1], N}, \tag{5}$$

where $N := N_{data} - T - \rho$ and ρ is the "past horizon".

Based on (3) the Hankel Y_F can be written as

$$Y_F = \Gamma X_\rho + \mathcal{H}_d U_F + \mathcal{H}_s E_F, \tag{6a}$$

where E_F is the Hankel of future innovations,

$$\Gamma = \begin{bmatrix} C \\ CA \\ CA^2 \\ \vdots \\ CA^{T-1} \end{bmatrix}, \tag{6b}$$

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

$$\mathcal{H}_d = \begin{bmatrix} D & 0 & 0 & \dots & 0 \\ CB & D & 0 & \dots & 0 \\ CAB & CB & D & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ CA^{T-2}B & CA^{T-3}B & CA^{T-4}B & \dots & D \end{bmatrix}, \tag{6c}$$

$$\mathcal{H}_s = \begin{bmatrix} I & 0 & 0 & \dots & 0 \\ CK & I & 0 & \dots & 0 \\ CAK & CK & I & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ CA^{T-2}K & CA^{T-3}K & CA^{T-4}K & \dots & I \end{bmatrix}. \tag{6d}$$

Let us now define $z(t)$ as the joint input/output process

$$z(t) := \begin{bmatrix} u(t) \\ y(t) \end{bmatrix}, \tag{7}$$

with the associated Hankel matrix being $Z_P := Z_{[0, \rho-1], N}$. The orthogonal projection of Y_F onto the row space of Z_P and U_F turns out to be given by

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

where the last term vanishes[1] (in probability) as $1/\sqrt{N}$. This means that, when the matrices (A, B, C, D, K) are *unknown*, future outputs can still be predicted directly from data. In fact, given any (past) joint input and output trajectory and future control inputs

$$z_{init} := \begin{bmatrix} z(t-\rho) \\ \vdots \\ z(t-2) \\ z(t-1) \end{bmatrix}, \quad u_f := \begin{bmatrix} u(t) \\ u(t+1) \\ \vdots \\ u(t+T-1) \end{bmatrix}, \tag{9}$$

the prediction $\hat{y}_f$ of future outputs y_f

$$y_f := \begin{bmatrix} y(t) \\ y(t+1) \\ \vdots \\ y(t+T-1) \end{bmatrix}, \tag{10}$$

based on past inputs z_{init} and future inputs u_f can be obtained from[2]

$$\begin{bmatrix} z_{init} \\ u_f \\ \hat{y}_f \end{bmatrix} = \begin{bmatrix} Z_P \\ U_F \\ \hat{Y}_F \end{bmatrix} \alpha + O_P(1/\sqrt{N}), \tag{11}$$

with $\alpha \in \mathbb{R}^N$ to be optimized as in, e.g., [3], [10], [11].

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

[2] Conditions on ρ for this to hold are provided in [10].

Following subspace identification [12] ideas, the orthogonal projection (8) can be written exploiting the LQ decomposition of the data matrices. In particular, let us define

$$\begin{bmatrix} Z_P \\ U_F \\ Y_F \end{bmatrix} = \begin{bmatrix} L_{11} & 0 & 0 \\ L_{21} & L_{22} & 0 \\ L_{31} & L_{32} & L_{33} \end{bmatrix} \begin{bmatrix} Q_1 \\ Q_2 \\ Q_3 \end{bmatrix}. \tag{12}$$

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 $i = 1, \dots, 3$, $Q_i Q_j^\top = 0$, $i \neq j$. The orthogonal projection (8) can be written in the form:

$$\hat{Y}_F = L_{31} Q_1 + L_{32} Q_2$$

With this notation, following the same rationale of [9], [10], we can further reformulate (11) as:

$$\begin{bmatrix} z_{init} \\ u_f \\ \hat{y}_f \end{bmatrix} = \begin{bmatrix} Z_P \\ U_F \\ \hat{Y}_F \end{bmatrix} \alpha = \begin{bmatrix} L_{11} & 0 \\ L_{21} & L_{22} \\ L_{31} & L_{32} \end{bmatrix} \underbrace{\begin{bmatrix} Q_1 \\ Q_2 \end{bmatrix} \alpha}_{\gamma} + O_P(1/\sqrt{N}). \tag{13}$$

and the parameters

$$\gamma = \begin{bmatrix} \gamma_1 \\ \gamma_2 \end{bmatrix}, \tag{14}$$

become the new decision variables. In addition, in [9] it was suggested to add a (slack) optimization variable γ_3 to model the projection error in (8) and avoid overfitting. In particular, the prediction (with slack) can be written as:

$$\bar{y}_f = \underbrace{\begin{bmatrix} L_{31} & L_{32} \end{bmatrix} \begin{bmatrix} \gamma_1 \\ \gamma_2 \end{bmatrix}}_{=\hat{y}_f} + L_{33}\gamma_3$$

We refer the reader to [9] for a sound statistical motivation of this particular expression of the slack $L_{33}\gamma_3$. In particular, since L_{33} is generically of full rank, constraints/regularization should be imposed on the slack optimization variable γ_3.

A data-driven predictive controller with the same objectives and constraints of (4) can be formulated as follows [10]

$$\min_{\gamma_2, \gamma_3} \frac{1}{2} \sum_{k=t}^{t+T-1} \ell(u(k), \bar{y}(k), y_r(k)) + \Psi(\gamma_1, \gamma_2, \gamma_3) \tag{15a}$$

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

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

with

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

and

$$\gamma_1^\star = L_{11}^{-1} z_{init}, \tag{17}$$

where z_{init} is defined as in (9) and the choice of γ_1 straightforwardly follows from the initial conditions (showing the advantages of using γ instead of α as the decision vector).

The purpose of this paper is *to study the design and impact of the regularization term $\Psi(\gamma_1, \gamma_2, \gamma_3)$ within a noisy stochastic environment, and provide the end user with useful hints on how to tune such a penalty term.*

III. The role of regularization

In [9], it has been argued that the average variance of the error on the future output predictions $\hat{y}_f$ due to the finite data projection errors in (8), is proportional to $\|\gamma_1\|^2 + \|\gamma_2\|^2$. Since, in the optimization problem (15), γ_1 is determined by the initial conditions, it only remains to regularize γ_2 so as to avoid an (unnecessarily) high variance on the predictor and, therefore, poor control performance. In this paper, we consider also an alternative regularization term that penalizes directly the control input effort (in addition to the control penalty already embedded in the control cost), and discuss its relation with regularization on γ_2. Differently from [9], we consider this jointly with presence of a slack variable γ_3 and thus a related regularization. These considerations lead to the following two forms of the regularization term $\Psi(\gamma_1, \gamma_2, \gamma_3)$ in (15):

(a) **Regularization on γ_2 and slack γ_3**

$$\Psi_{\gamma_2}(\gamma_1, \gamma_2, \gamma_3) := \beta_2 \|\gamma_2\|^2 + \beta_3 \|\gamma_3\|^2; \tag{18}$$

(b) **Regularization on input u_f and slack γ_3**

$$\begin{aligned} \Psi_u(\gamma_1, \gamma_2, \gamma_3) :=& \beta_2 \|u_f\|^2 + \beta_3 \|\gamma_3\|^2 \\ =& \beta_2 \|L_{21}\gamma_1 + L_{22}\gamma_2\|^2 + \beta_3 \|\gamma_3\|^2; \end{aligned} \tag{19}$$

where (β_2, β_3) are hyper-parameters to be determined.

A. Theoretical analysis

We first state a Theorem the establishes the connection between (18) and (19).

Theorem 1: If the training input sequence $u(t)$ in the Hankel matrices U_F and U_P is (zero mean) white with variance $\sigma^2 I$, the regularization terms Ψ_{γ_2} in (18) and Ψ_u in (19) are asymptotically (in N) equivalent up to a rescaling of the weight β_2.

Proof: Under the assumption that $u(t)$ is white noise, then the future inputs are uncorrelated with past input and output data, so that the projection $\hat{U}_F := \Pi_{Z_P}(U_F)$ of U_F on the joint past Z_P tends to zero as $O_P(1/\sqrt{N})$, more precisely

$$\hat{U}_F := L_{21} Q_1. \tag{20}$$

Since $Q_1 Q_1^\top = I$, it follows that $L_{21} = O_P(1/\sqrt{N})$. In addition, since u is white, its sample covariance matrix $U_F U_F^\top$ converges to $\sigma^2 I$, i.e.

$$U_F U_F^\top = \underbrace{L_{21} L_{21}^\top}_{O_P(1/N)} + L_{22} L_{22}^\top \overset{N \to \infty}{\longrightarrow} \sigma^2 I \tag{21}$$

Equations (20) and (21) imply that, asymptotically in N, $L_{21} \simeq 0$ and $L_{22} \simeq \sigma I$. Therefore we have:

$$\begin{aligned} \Psi_u(\gamma_1, \gamma_2, \gamma_3) \quad &:= \quad \beta_2 \|L_{21}\gamma_1 + L_{22}\gamma_2\|^2 + \beta_3 \|\gamma_3\|^2 \\ &\simeq \quad \beta_2 \sigma^2 \|\gamma_2\|^2 + \beta_3 \|\gamma_3\|^2 \end{aligned} \tag{22}$$

showing that, up to the rescaling of the weight β_2, this is equivalent to $\Psi_{\gamma_2}(\gamma_1, \gamma_2, \gamma_3)$ ■

This result has two important implications:

- When the (training) input is white, regularization on γ_2 is equivalent to a penalty on the future input energy, which is typically present in the control cost. As such, we can argue that, in this case, the control cost has an *indirect but important* effect in counteracting the effect of the noise variance in the predictor.
- When the training input is not white, the control energy cost *is not* equivalent to penalizing the norm of γ_2, which on the other hand should be penalized to limit the effect of noise variance. The simulation results in the next section indeed confirm that, when noise input is not white, regularization on γ_2 (i.e. Ψ_{γ_2}) has to be included.

B. Experimental analysis

In this section we shall illustrate, exploiting two numerical examples (one linear and one nonlinear[3]), the role of different regularization terms in the optimal control problem (15). In particular, following the rationale proposed in [7], we evaluate the closed-loop performance over T_v feedback steps as measured by the performance index:

$$J(u,y) = \frac{1}{T_v}\sum_{t=0}^{T_v-1}\left(\|u(t)-u_r(t)\|_R^2 + \|y(t)-y_r(t)\|_Q^2\right). \quad (23)$$

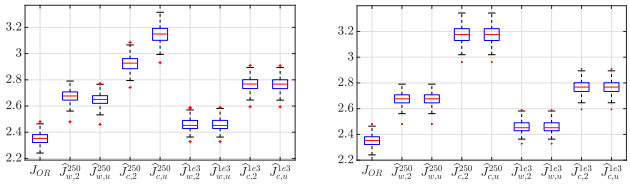

(a) Performance indexes (b) Optimal performance under constraint $\beta_2 = 0$

Fig. 1: (a): comparison between the Kalman-filter-based oracle performance J_{OR} and the minimum cost realizations $\hat{J}^{N_{data}}_{n_s,rg}$ for Σ_L over 100 Monte Carlo runs; (b): Optimal performance under the constraint $\beta_2 = 0$.

1) Benchmark LTI system: Consider the SISO, 4-th order system in [13] (Σ_L in the sequel) with a prediction horizon $T = 20$. To assess the impact of the *training* data on closed-loop performance, we consider four data sets of two different lengths N_{data} (either 250 or 1000), obtained either with white noise input (denoted with $n_s = w$) or with a low-pass (obtained filtering white noise with a discrete-time low-pass filter with cut-off angular frequency 1.8 rad/s) input sequence (denoted with $n_s = c$). White noise is added to the output to guarantee a signal-to-noise ratio of 15 dB.

The data-driven optimal control problem (15) is solved fixing $Q = 10^3$ and $R = 10^{-2}$, setting the output reference $y_r(t) = \sin(5\pi t/(T + T_v - 1))$ and the input reference $u_r(t) = 0$, with $T_v = 50$. The two different regularization strategies discussed in Section III are denoted with the shorthands $rg = 2$ and $rg = u$ for (18) and (19), respectively.

The "oracle" value of (β_2, β_3) leading to **the minimum cost** (23), here denoted as $J^{N_{data}}_{n_s,rg} = J(u^{N_{data}}_{n_s,rg}, y^{N_{data}}_{n_s,rg})$ to account for the different data set lengths, input and regularization strategies, are searched over a rectangular logarithmic-spaced grid G_{23} with 7 points per decade, so that $G_{23} \subseteq \{0\} \cup [10^{-4}, 10^0] \cup \{+\infty\} \times \{0\} \cup [10^{-3}, 10^0] \cup \{+\infty\}$.

We perform 100 Monte Carlo experiments[4] (i.e. 100 different training data sets with the output of Σ_L corrupted by white noise) to tune β_2 and β_3 for all the four considered training scenarios and possible regularization. For each Monte Carlo run, and for each set of possible parameters in the grid G_{23}, the closed loop performance index (23) is computed by averaging over 100 closed loop experiments (all with the same control law but different closed measurements errors) the corresponding performance index $J^{N_{data}}_{n_s,rg}(i)$, i.e.

$$\bar{J}^{N_{data}}_{n_s,rg} = \frac{1}{100}\sum_{i=1}^{100} J^{N_{data}}_{n_s,rg}(i), \quad (24)$$

The optimal values of β_2 and β_3 over the grid G_{23} is obtained by finding the minimum $\hat{J}^{N_{data}}_{n_s,rg} = \min_{(\beta_2,\beta_3)\in G_{23}}\{\bar{J}^{N_{data}}_{n_s,rg}\}$.

The results are reported in Fig. 1, based on which we can make the following general considerations:

- As expected based on Theorem 1, when the input is white noise, the two types of regularization provide the same performance (minor differences for $N_{data} = 250$ are due to sample variability).
- The "optimal" (oracle) closed loop performance obtained with the two different regularization strategies differ when the input is not white. In particular, the penalty (18) that acts directly on γ_2 and thus controls the predictor variance provides the best performance, particularly so for small data sets where the effect of noise has more impact.
- Based on the comparison between Fig.1(a) and Fig. 1(b) (in which β_2 is constrained to zero), we can observe that the impact of β_2 is significant in the low-data regime (equivalent to large noise in the predictor), whereas for larger data sets its impact can be neglected and the optimal performances exploiting only β_3 match those obtained optimizing jointly β_2 and β_3.

In light of the above observations, the general validity of (15) constrained to either $\beta_2 = 0$ or $\beta_3 = \infty$ devised in [9] is strengthened, as different types of data set are given. The effectiveness of these γ-DDPC schemes is also reinforced since it is evident that, in most cases, the operative tuning of either the sole parameter β_2 or the sole parameter β_3 is worth to be carried out in practice.

2) Wheel slip control problem: We now consider the problem of designing a wheel slip controller, steering the vehicle slip $\lambda(t) \in [0, 1]$ to a constant target value λ_r. The design is carried out by focusing on quasi-stationary

[3]By working in a specific operating regime, the control of an unknown nonlinear system can be tackled as that of an uncertain linear system.

[4]The past horizon ρ is determined using Akaike's information criterion.

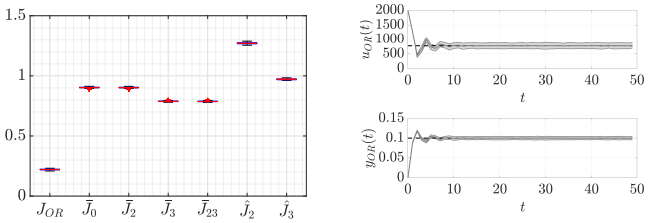

(a) Performance indexes (b) MPC-based oracle

Fig. 2: (a): comparison of the performance indexes obtained with different γ-DDPC strategies (bar and hat notation indicating offline and online approaches respectively) and a model-based oracle. The subscript $a \in \{0, 2, 3, 23\}$ on J refers to the regularization scheme (respectively: no regularization, β_2, β_3 or both); (b): input/output tracking obtained from an MPC-based oracle. Mean (line) and 1.95 times the standard deviation (shaded area) of the closed-loop input/output trajectories; the reference input and output are indicated with black dashed lines.

operating condition (the parameters of the vehicle, its velocity and the road profile are assumed to be constant). In both data collection and closed-loop testing, the behavior of the braking system (from now on indicated as Σ_{NL}) is simulated based on the *nonlinear* model in [14]. We indicate with $T_b(t)$ [Nm] the *controllable* braking torque and set the system parameters to the same values used in [15]. Although this dynamics is clearly nonlinear, it is possible to identify two main operating regions of the system[5], where the behavior of the slip can be approximated as linear. To comply with our framework (see (3)), we thus consider both data collection and simulation tests where the vehicle generally operates in a low-slip regime. Accordingly, data are gathered by performing closed-loop experiments with the benchmark controller introduced in [16], selecting a slip reference uniformly chosen at random in the interval $[0, 0.15]$ and collecting data at a sampling rate of 100 [Hz]. In particular, the output $y_{trn}(t)$ of the employed training data set is generated by exploiting a closed-loop experiment wherein the output is corrupted by a zero-mean white noise process with variance $\sigma_n^2 = 10^{-6}$ and, also, zero-mean white noise with variance $10^8\sigma_n^2$ is added to the input $u_{trn}(t)$ provided by the controller.

Meanwhile, the reference slip for the closed-loop tests is $\lambda_r = 0.1$, corresponding to a reference braking torque $T_{b,r} = 768.9$ [Nm]. To improve the tracking performance in closed-loop, apart from the terms weighting the tracking error and the difference between the predicted and reference torque, respectively weighted by $Q = 10^3$ and $R = 10^{-7}$, the cost of the γ-DDPC problem (15a) is augmented with a term penalizing abrupt variations of the input (weighted as 10^{-4}), a term penalizing the integral of the tracking error (weighted as 10^5), and two terms further penalizing the difference between the slip and torque references and their actual value over the last step of the prediction horizon (weighted as 10^3 and $5 \cdot 10^{-6}$, respectively). The following constraint is also added at each feedback step $t \geqslant 0$ for $s = 0, \dots, T-2$:

$$\begin{cases} q(t+s) = q(t+s-1) + y_r(t+s-1) - \hat{y}(t+s-1); \\ q(t-1) = y_r(t-1) - y(t-1); \end{cases} \tag{25}$$

to account for the known dynamics of the integrator. Nonetheless, performances are still assessed via the index in (23) over a closed-loop test of $T_v = 50$ steps. A Monte Carlo campaign with 100 iterations is run on the above setup, corrupting the output of Σ_{NL} with a white noise having signal-to-noise ratio 40 dB. For each of the 100 tests, the regularization parameters β_2 and β_3 are both selected from a grid over $[10^{-4}, 10^4]$ comprising of 15 logarithmic-spaced points. For the joint optimization, the squared grid $\{\bar{\beta}_2, \bar{\beta}_2/10, \bar{\beta}_2/100, 10^8\} \times \{0, \bar{\beta}_3, 10\bar{\beta}_3, 100\bar{\beta}_3\}$ composed by the optimal values $(\bar{\beta}_2, \bar{\beta}_3)$ obtained via offline γ-DDPC is instead taken into account. Fig. 2(a) depicts the distributions of the performance index in (23) as the selected regularization strategy varies considering (β_2, β_3) tuned either offline or online and comparing γ-DDPC with a MPC-based oracle (see also Fig. 2(b)). In particular, the input-output trajectories of all γ-DDPC strategies can be summarized in Fig. 3. Although the MPC-based oracle displays evident preeminence, it is worthwhile to appreciate that all these trajectories are characterized by solid performances (rise time of at most 5 steps, settling time of about 25 steps, maximum overshoot of 40% or less with no cross into the unstable region), with such traits indicating that γ-DDPC schemes remain competitive even in nonlinear scenarios. Noticeably, the online strategies (implementable on real applications) shown in Fig. 3(e)-3(f) share similar performances with the corresponding offline strategies in Fig. 3(b) - 3(c), especially that relying on the online tuning of parameter β_3. Moreover, within the setup of this numerical example, one observes that the performance of the offline strategy based on β_2 strictly matches with that of the scheme lacking of regularization (Fig. 3(a)); whereas, the performance of the offline strategy based on β_3 strictly matches that of the scheme in which a joint optimization of both β_2 and β_3 (Fig. 3(d)) is carried out. Hence, under this setup and with the data collected in this numerical example, it emerges once again that the optimal tuning based on the sole penalty parameter β_3 (i.e., setting $\beta_2 = 0$) can be considered *in practice* for high-data regimes. This, in turn, may lead to significantly diminish the computational burden associated to the tuning of the penalty parameters whenever a real implementation based on the proposed regularized scheme (15) is considered and a big training data set is available.

The above comparison further highlights that regularized DDPC approaches can be competitive w.r.t. traditional model-based controllers and that γ-DDPC solution with the online tuning proposed in [9] can be robustly effective also when dealing with nonlinear systems.

[5]In this case, these two regions are limited by the slip $\bar{\lambda} = 0.17$, for slip values lower than 0.17; while it becomes unstable for higher slips.

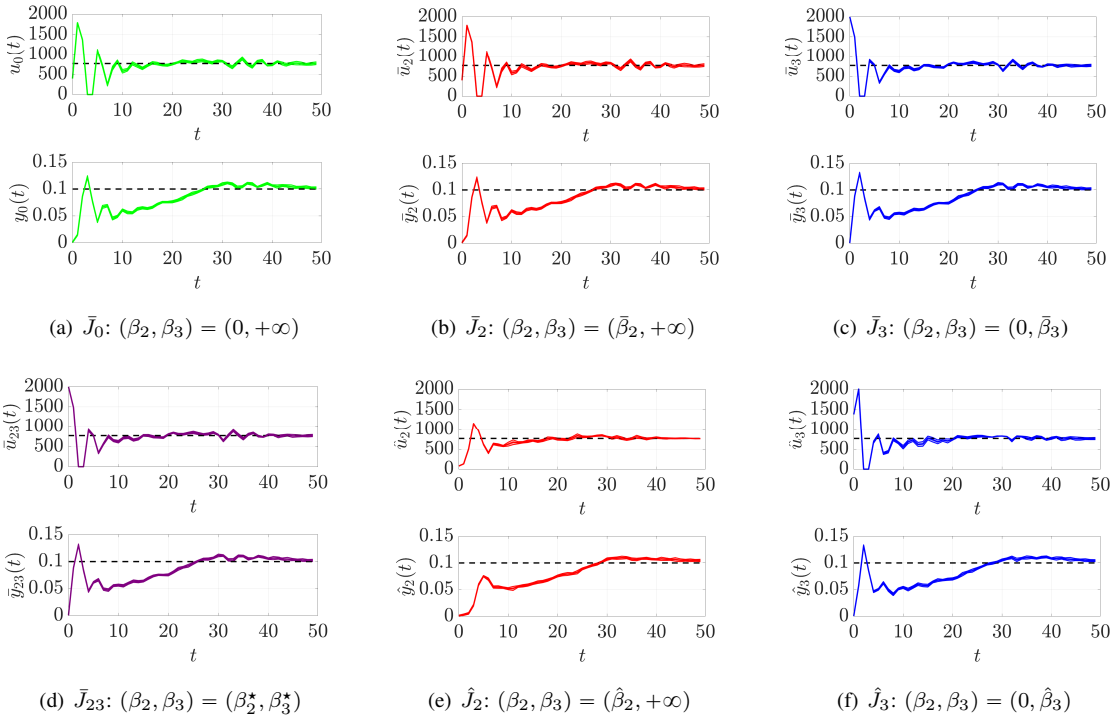

(a) $\bar{J}_0$: $(\beta_2, \beta_3) = (0, +\infty)$ (b) $\bar{J}_2$: $(\beta_2, \beta_3) = (\bar{\beta}_2, +\infty)$ (c) $\bar{J}_3$: $(\beta_2, \beta_3) = (0, \bar{\beta}_3)$

(d) $\bar{J}_{23}$: $(\beta_2, \beta_3) = (\beta_2^\star, \beta_3^\star)$ (e) $\hat{J}_2$: $(\beta_2, \beta_3) = (\hat{\beta}_2, +\infty)$ (f) $\hat{J}_3$: $(\beta_2, \beta_3) = (0, \hat{\beta}_3)$

Fig. 3: For all diagrams: mean (line) and 1.95 times the standard deviation (shaded area) of the closed-loop input/output trajectories; the reference input and output are indicated with black dashed lines. (a): γ-DDPC with no regularization; (b)-(c): offline regularization strategies employing $\bar{\beta}_2$ and $\bar{\beta}_3$ separately; (d): offline regularization strategies employing $\bar{\beta}_2$ and $\bar{\beta}_3$ jointly; (e)-(f): online regularization strategies employing $\hat{\beta}_2$ and $\hat{\beta}_3$ separately.

IV. Concluding Remarks and Future Directions

Several regularization strategies for Data Driven Predictive Control (γ-DDPC) have been discussed and evaluated in terms of closed-loop performance. It has been proved that when the input is white, regularizing γ_2 and penalizing control energy are equivalent. Numerical examples further illustrate that the tuning of the penalty parameters in the γ-DDPC can be decoupled without dramatically impacting the performance corresponding to a (more costly) joint regularization wherein both β_2 and β_3 are accounted for.

Future work will be devoted to the extensive experimental assessment of the considered regularization strategies, as well as to a theoretical analysis of the optimization of the sole β_3.

References

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

[2] S. Baros, C.-Y. Chang, G. E. Colón-Reyes, and A. Bernstein, "Online data-enabled predictive control," *Automatica*, vol. 138, p. 109926, 2022.

[3] J. Berberich, J. Köhler, M. A. Müller, and F. Allgöwer, "Data-driven model predictive control with stability and robustness guarantees," *IEEE Transactions on Automatic Control*, vol. 66, no. 4, pp. 1702–1717, 2020.

[4] V. Breschi, A. Sassella, and S. Formentin, "On the design of regularized explicit predictive controllers from input-output data," *IEEE Transactions on Automatic Control*, 2022.

[5] A. Sassella, V. Breschi, and S. Formentin, "Data-driven design of explicit predictive controllers," in *2022 IEEE 61st Conference on Decision and Control (CDC)*. IEEE, 2022, pp. 2821–2826.

[6] ——, "Noise handling in data-driven predictive control: a strategy based on dynamic mode decomposition," in *Learning for Dynamics and Control Conference*. PMLR, 2022, pp. 74–85.

[7] F. Dorfler, J. Coulson, and I. Markovsky, "Bridging direct & indirect data-driven control formulations via regularizations and relaxations," *IEEE Transactions on Automatic Control*, 2022.

[8] V. Breschi, A. Chiuso, and S. Formentin, "The role of regularization in data-driven predictive control," *arXiv preprint arXiv:2203.10846*, 2022.

[9] V. Breschi, M. Fabris, S. Formentin, and A. Chiuso, "Uncertainty-aware data-driven predictive control in a stochastic setting," *arXiv preprint arXiv:2211.10321*, 2022.

[10] V. Breschi, A. Chiuso, and S. Formentin, "Data-driven predictive control in a stochastic setting: a unified framework," *Automatica*, vol. 152, no. 110961, 2023.

[11] J. Coulson, J. Lygeros, and F. Dörfler, "Data-enabled predictive control: In the shallows of the deepc," in *2019 18th European Control Conference (ECC)*. IEEE, 2019, pp. 307–312.

[12] P. Van Overschee and B. De Moor, *Subspace Identification for Linear Systems*. Kluwer Academic Publications, 1996.

[13] I. Landau, D. Rey, A. Karimi, A. Voda, and A. Franco, "A flexible transmission system as a benchmark for robust digital control," *European Journal of Control*, vol. 1, no. 2, pp. 77–96, 1995.

[14] S. Formentin, C. Novara, S. Savaresi, and M. Milanese, "Active braking control system design: The D^2 -IBC approach," *IEEE/ASME Transactions on Mechatronics*, vol. 20, no. 4, pp. 1573–1584, 2015.

[15] A. Sassella, V. Breschi, S. Formentin, and S. Savaresi, "A data-driven switching control approach for braking systems with constraints," *Nonlinear Analysis: Hybrid Systems*, vol. 46, p. 101220, 2022.

[16] S. M. Savaresi and M. Tanelli, *Active braking control systems design for vehicles*. Springer Science & Business Media, 2010.