

Data-driven design of explicit predictive controllers using model-based priors [★]

Valentina Breschi ^a, Andrea Sassella ^a, Simone Formentin ^a

^a*Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Piazza L. da Vinci 32, 20133 Milano, Italy.*

Abstract

In this paper, we propose a data-driven approach to derive explicit predictive control laws, without requiring any intermediate identification step. The keystone of the presented strategy is the exploitation of available priors on the control law, coming from model-based analysis. Specifically, by leveraging on the knowledge that the optimal predictive controller is expressed as a piecewise affine (PWA) law, we directly optimize the parameters of such an analytical controller from data, instead of running an on-line optimization problem. As the proposed method allows us to automatically retrieve also a model of the closed-loop system, we show that we can apply model-based techniques to perform a stability check prior to the controller deployment. The effectiveness of the proposed strategy is assessed on two benchmark simulation examples, through which we also discuss the use of regularization and its combination with averaging techniques to handle the presence of noise.

Key words: Data-driven control; learning-based control; predictive control; explicit predictive control

1 Introduction

For its capability of handling constraints, *Model predictive control* (MPC) is a widely employed technique for advanced control applications (see, *e.g.*, [14, 19, 21, 23, 24]). Due to the increasing complexity of the systems to be controlled, the model required by MPC is often no longer parameterized from the physics but learned from data in a *black-box* fashion. However, this identification step generally takes the lion's share of the time and effort required for control design. In the last years, research has thus benched into two main direction. On the one side, efforts has been focused on improving and easing learning procedures [9, 20]. On the other side, many approaches have been proposed to directly employ data for the design of predictive controllers, while bypassing any model identification step. Among existing works in this direction, we recall the foundational contributions of [5, 10], which have been extended to handle tracking problems [4], to deal with nonlinear sys-

tems [6], and to improve the performance in the presence of noise [11, 13], just to mention a few. For both traditional and data-driven predictive controls, the computational effort required to solve the related constrained optimization problem is known to be a potential limit for its application, especially for fast-sampling systems. Nonetheless, when the MPC problem is relatively small, *i.e.*, one has to control a low order system and/or the prediction horizon is relatively short, this limitation can be overcome by explicitly deriving its solution [3]. In fact, when the cost of the optimization problem is quadratic and the constraints are linear, the explicit solution of MPC is known to be a *Piece-Wise Affine* (PWA) state feedback law.

In this work, we propose an approach *to directly learn explicit predictive control laws from data*. More specifically, we initially build on the foundational results in [12, 26] and on model-based priors (namely, the aforementioned fact that the optimal solution of linear MPC is PWA [3]), to construct a data-driven, multi-step closed-loop predictor. The latter is exploited to construct a fully data-based optimization problem, yielding the estimation of the *parameters* of the optimal predictive control law. As a by-product, we obtain a data-driven description of the closed-loop behavior, which can be used in combination with existing model-based techniques (see, *e.g.*, [22]) to check the stability of the system controlled with the data-driven explicit controller, prior to its deployment.

[★] 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.

Email addresses: valentina.breschi@polimi.it (Valentina Breschi), andrea.sassella@polimi.it (Andrea Sassella), simone.formentin@polimi.it (Simone Formentin).

As far as we are aware, this is the first time a data-driven predictive control strategy is provided together with a preliminary assessment of its closed-loop performance.

We should mention here that an early attempt to obtain a data-driven counterpart of explicit MPC was already carried out in [25]. However, the strategy proposed therein relies on an implicit open-loop identification step. Another method was presented in [8], by relying on the behavioral predictor used in [5, 11]. Even though that method involves no open-loop identification phase and it does not require the state to be fully measurable, the latter does not leverage on priors coming from model-based analysis and design. Nonetheless, by leveraging on priors, we are here able to retrieve a data-based characterization of the closed-loop and, thus, practically assess its stability in a data-driven fashion prior to its deployment. This is instead not possible in [8], where stability cannot be directly assessed nor it is theoretically guaranteed in presence of noisy data.

The remainder of the paper is organized as follows. The targeted problem is formalized in Section 2, while all the steps required to obtain its data-driven formulation from priors are introduced in Section 3. The explicit control law is derived in Section 4, where we additionally discuss practical aspects for its implementation, certified deployment and noise handling. The effectiveness of the proposed strategy is assessed on two benchmark examples in Section 5. Section 6 concludes the work and indicates some directions for future research.

Notation We denote with \mathbb{N}_0 the set of natural numbers, that includes zero. Let \mathbb{R} , \mathbb{R}^n and $\mathbb{R}^{n \times m}$ be the set of real numbers, column vectors of dimension n and $n \times m$ dimensional real matrices, respectively. Given $B \in \mathbb{R}^{m \times n}$, its transpose is B^\top , its Moore-Penrose inverse is B^\dagger and, when $m = n$ its inverse is indicated as B^{-1} . Given a vector $v \in \mathbb{R}^n$, $[v]_{i:j}$ indicates its rows from i to j , with $i \leq j \leq n$. For a matrix $B \in \mathbb{R}^{n \times m}$, $[B]_{1:i, 1:j}$ denotes a sub-matrix comprising the first i rows and j columns of B , with $i \leq n$ and $j \leq m$. Identity matrices are denoted as I , while zero matrices and vectors will be denoted as 0 . If a matrix $Q \in \mathbb{R}^{n \times n}$ is positive definite (positive semi-definite), this is denoted as $Q \succ 0$ ($Q \succeq 0$). Given a vector $x \in \mathbb{R}^n$, the quadratic form $x^\top Q x$ is compactly indicated as $\|x\|_Q^2$. Given a signal $\{\nu_t \in \mathbb{R}^m\}_{t \in \mathbb{N}_0}$ and $0 < L < T$, we denote with $N_{0,L,T-1} \in \mathbb{R}^{mL \times T-L+1}$ the associated Hankel matrix

$$N_{0,L,T-1} = \begin{bmatrix} \nu(0) & \nu(1) & \cdots & \nu(T-L) \\ \nu(1) & \nu(2) & \cdots & \nu(T-L+1) \\ \vdots & \vdots & \ddots & \vdots \\ \nu(L) & \nu(L+1) & \cdots & \nu(T-1) \end{bmatrix}, \quad (1)$$

while, for $i, j \in \mathbb{N}_0$, we introduce

$$N_{i,T-j} = \left[\nu(i) \ \nu(i+1) \ \cdots \ \nu(T-j) \right], \quad i < T-j. \quad (2)$$

2 Problem formulation

Consider the class \mathcal{S} of discrete-time *linear, time invariant* (LTI), *controllable* systems with *fully measurable* state, *i.e.*,

$$\mathcal{S} : \begin{cases} x(t+1) = Ax(t) + Bu(t), \\ y^o(t) = x(t), \end{cases} \quad (3)$$

where $x(t) \in \mathbb{R}^n$ denotes the state of \mathcal{S} at time $t \in \mathbb{N}_0$, $u(t) \in \mathbb{R}^m$ is an exogenous input and $y^o(t) \in \mathbb{R}^n$ is the associated *noiseless* output. Let us consider the following *model predictive control* (MPC) problem:

$$\begin{aligned} & \underset{\{\tilde{u}(k)\}_{k=0}^{L-1}}{\text{minimize}} && \sum_{k=0}^{L-1} [\|\tilde{x}(k)\|_Q^2 + \|\tilde{u}(k)\|_R^2] + \|\tilde{x}(L)\|_P^2 && (4a) \\ & \text{s.t.} && \tilde{x}(k+1) = A\tilde{x}(k) + B\tilde{u}(k), \quad k=0, \dots, L-1, && (4b) \\ & && C_x \tilde{x}(k) + C_u \tilde{u}(k) \leq d, \quad k=0, \dots, L-1, && (4c) \\ & && \tilde{x}(0) = x(t). && (4d) \end{aligned}$$

The objective of (4) is to steer both the (predicted) state $\tilde{x}(k)$ and the input $\tilde{u}(k)$ to zero, over a prediction horizon of prefixed length $L > 0$. Optimality is indeed dictated by: (i) the distance of the predicted state from zero, penalized with $Q \succeq 0$ over the whole horizon except for the terminal state (weighted via $P \succeq 0$), and (ii) the control effort, penalized via $R \succ 0$. Meanwhile, a set of n_c *polyhedral constraints* dictated by (4c) has to be satisfied, with $C_x \in \mathbb{R}^{n_c \times n}$, $C_u \in \mathbb{R}^{n_c \times m}$ and $d \in \mathbb{R}^{n_c}$, while relying on the latest information on the system (see (4d)). Assume also that the matrices $A \in \mathbb{R}^{n \times n}$ and $B \in \mathbb{R}^{n \times m}$ characterizing the dynamics of \mathcal{S} are *unknown*, and that we can access a set of input/output data pairs $\mathcal{D}_T = \{\mathcal{U}_T, \mathcal{Y}_T\}$, where \mathcal{U}_T and \mathcal{Y}_T denote the available input and output sequences, respectively, satisfying the following assumptions.

Assumption 1 (Persistently exciting inputs) *The input sequence $\mathcal{U}_T = \{u(t)\}_{t=0}^T$ is persistently exciting of order $n+L$.*

Assumption 2 (Noisy outputs) *The output sequence $\mathcal{Y}_T = \{y(t)\}_{t=0}^T$ is corrupted by noise, namely*

$$y(t) = y^o(t) + w(t), \quad (5)$$

where $v(t)$ is the realization of a zero mean white noise with covariance $\Omega \in \mathbb{R}^{n \times n}$.

Assumption 3 (Sufficiently long dataset) *The length T of the dataset \mathcal{D}_T satisfies the following:*

$$T \geq (m+1)(n+L) - 1.$$

The goal of this work is to directly exploit the available data to find an *explicit* solution to (4), under Assumptions 1-3, while *bypassing* any identification step.

3 Exploiting priors for explicit DDPC

To attain our goal, it is fundamental to replace the model in (4b) with an expression that directly depends on the available data, by also not forgetting what we have learned from model-based predictive control. To start with, we thus recall the following lemma, explicitly stating the form of the optimal explicit predictive controller in our setting [3].

Lemma 1 (On the solution of (4)) *Let U denote the vector stacking the inputs over the prediction horizon, i.e.,*

$$U = U(x(t)) = \begin{bmatrix} \tilde{u}(0) \\ \tilde{u}(1) \\ \vdots \\ \tilde{u}(L-1) \end{bmatrix} \in \mathbb{R}^{mL}. \quad (6)$$

The optimal control sequence $U^ = U^*(x(t))$ solving (4) is a continuous piecewise affine (PWA) function of $x(t)$.*

Then, we generalize the one-step ahead predictor introduced in [12] to the multi-step case. To this end, it is important to recall that, thanks to Assumptions 1, 3 and the *Fundamental Lemma* [26], the following rank condition holds:

$$\text{rank} \left(\begin{bmatrix} X_{0,T-L-1} \\ U_{0,L,T-L-1} \end{bmatrix} \right) = n + mL. \quad (7)$$

We can now derive the multi-step data-based predictor as follows.

Theorem 1 (Data-based multi-step predictor)

Let Assumptions 1 and 3 hold. Given $x(t)$, the sequence of predicted states $\tilde{X} = \tilde{X}(x(t))$ admits the following data-based representation:

$$\tilde{X} = \begin{bmatrix} \tilde{x}(1) \\ \vdots \\ \tilde{x}(L) \end{bmatrix} = X_{1,L,T-L} \begin{bmatrix} X_{0,T-L-1} \\ U_{0,L,T-L-1} \end{bmatrix}^\dagger \begin{bmatrix} x(t) \\ U \end{bmatrix}. \quad (8)$$

Proof: The proof follows the steps of the one in [12, Appendix B]. Specifically, let $v \in \mathbb{R}^{n+mL}$ and $S \in \mathbb{R}^{n+mL \times T}$ be respectively defined as:

$$v := \begin{bmatrix} x(t) \\ U \end{bmatrix}, \quad S := \begin{bmatrix} X_{0,T-L-1} \\ U_{0,L,T-L-1} \end{bmatrix},$$

with S being full row rank, i.e., $\text{rank}(S) = n + mL$. By the Rouché-Capelli theorem, for any given v , the equality

$$v = S\alpha,$$

admits infinite solutions α of the form

$$\alpha = S^\dagger v + \Pi_S^\perp w, \forall w \in \mathbb{R}^T, \quad (9)$$

with $\Pi_S^\perp = (I - S^\dagger S)$ being the orthogonal projector onto the kernel of S . Meanwhile, based on the model in (4b), the predicted state sequence can be defined as a function of $x(t)$ and U as:

$$\tilde{X} = \underbrace{\begin{bmatrix} A \\ \vdots \\ A^{L-1} \end{bmatrix}}_{\xi} x(t) + \underbrace{\begin{bmatrix} B & 0 & \cdots & 0 \\ AB & B & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ A^{L-1}B & A^{L-2}B & \cdots & B \end{bmatrix}}_{\Gamma} U. \quad (10)$$

In turn, such a sequence can be recast as a function of α , i.e.,

$$\tilde{X} = \begin{bmatrix} \xi & \Gamma \end{bmatrix} S\alpha = X_{1,L,T-L}\alpha.$$

where the second equality straightforwardly follows from (10) and the definition of S . By replacing α with (9), we then obtain

$$\tilde{X} = X_{1,L,T-L} (S^\dagger v + \Pi_S^\perp w) = X_{1,L,T-L} S^\dagger v,$$

as $X_{1,L,T-L} \Pi_S^\perp = \begin{bmatrix} \xi & \Gamma \end{bmatrix} S \Pi_S^\perp = 0$, based on the definition of the projector. \square

This preliminary result allows us to exploit priors on the solution of (4) for the definition of the DDPC problem. In fact, according to Lemma 1, we can parameterize the control sequence U as:

$$U = \begin{cases} K_1 x(t) + f_1, & \text{if } H_1 x(t) \leq \ell_1, \\ \vdots \\ K_M x(t) + f_M, & \text{if } H_M x(t) \leq \ell_M, \end{cases} \quad (11)$$

where $K_i \in \mathbb{R}^{mL \times n}$ and $f_i \in \mathbb{R}^{mL}$ are the (unknown) feedback and affine gains characterizing the control law, $\{H_i, \ell_i\}_{i=1}^M$ dictates the associated polyhedral partition,

for $i = 1, \dots, M$, and the amounts of modes M is dictated by the number of possible combinations of active constraints. Therefore, for a given state $x(t)$, the input sequence is the affine function

$$U(x(t)) = Kx(t) + f, \quad (12)$$

with $K = K_{s(x(t))}$ and $f = f_{s(x(t))}$ denoting the gains associated to the active control law, and with

$$s(x(t)) = i \iff H_i x(t) \leq \ell_i, i \in \{1, \dots, M\}.$$

Based on this parameterization, we can compute a data-based closed-loop characterization of the predictor in (4b), as outlined in the following theorem.

Theorem 2 (Closed-loop multi-step predictor)

Let Assumptions 1 and 3 hold. Given $x(t)$, the sequence of predicted states $\tilde{X} = \tilde{X}(x(t))$ can be equivalently expressed as:

$$\tilde{X} = X_{1,L,T-L}(G_K x(t) + G_f), \quad (13a)$$

with $G_K \in \mathbb{R}^{T-L \times n}$ and $G_f \in \mathbb{R}^{T-L}$ satisfying:

$$\begin{bmatrix} I \\ K \end{bmatrix} = \begin{bmatrix} X_{0,T-L-1} \\ U_{0,L,T-L-1} \end{bmatrix} G_K \quad (13b)$$

$$\begin{bmatrix} 0 \\ f \end{bmatrix} = \begin{bmatrix} X_{0,T-L-1} \\ U_{0,L,T-L-1} \end{bmatrix} G_f \quad (13c)$$

where K and f characterize the local control law (12). Accordingly, the input sequence is given by:

$$U = U_{0,L,T-L-1}(G_K x(t) + G_f). \quad (14)$$

Proof: For the Rouché-Capelli theorem, there exists an $T \times n$ matrix G_k and a T -dimensional vector G_f such that (13b)-(13c) hold. Meanwhile, replacing (12) into the open-loop predictor in (8), the predicted state sequence can be represented as:

$$\tilde{X} = X_{1,L,T-L} \begin{bmatrix} X_{0,T-L-1} \\ U_{0,L,T-L-1} \end{bmatrix}^\dagger \left(\begin{bmatrix} I \\ K \end{bmatrix} x(t) + \begin{bmatrix} 0 \\ f \end{bmatrix} \right).$$

By combining these two results, the closed-loop representation in (13) and the equivalent definition of the input sequence in (14) straightforwardly follow. \square

By leveraging on (13)-(14), we can now equivalently recast the predictive control task in (4) as an optimization

problem with the closed-loop matrices G_K, G_f being the decision variables, as follows¹:

$$\underset{G_K, G_f}{\text{minimize}} \quad \tilde{X}^\top \mathcal{Q} \tilde{X} + U^\top \mathcal{R} U \quad (15a)$$

$$\text{s.t.} \quad \tilde{X} = X_{1,L,T-L}(G_K x(t) + G_f) \quad (15b)$$

$$U = U_{0,L,T-L-1}(G_K x(t) + G_f), \quad (15c)$$

$$\begin{bmatrix} C_x & 0 \\ 0 & \mathcal{C}_x \end{bmatrix} \begin{bmatrix} x(t) \\ \tilde{X} \end{bmatrix} + \mathcal{C}_u U \leq \mathcal{D}, \quad (15d)$$

$$X_{0,T-L-1} G_K x(t) = x(t), \quad (15e)$$

$$X_{0,T-L-1} G_f = 0, \quad (15f)$$

In (15), $\mathcal{Q} = \text{diag}([Q, \dots, Q, P])$, $\mathcal{R} = \text{diag}([R, \dots, R])$, $\mathcal{C}_u = \text{diag}([C_u, \dots, C_u])$ and

$$\mathcal{C}_x = \begin{bmatrix} C_x & 0 & \dots & 0 & 0 \\ 0 & C_x & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & C_x & 0 \end{bmatrix}, \quad \mathcal{D} = \begin{bmatrix} d \\ \vdots \\ d \end{bmatrix}.$$

Note that, the last constraints (see (15e)-(15f)) are introduced for the problem to be consistent with the closed-loop representation in (13).

This shift from an open-loop predictor to its closed-loop counterpart allows us to directly learn the control law from data, and avoid any system identification step.

Remark 1 Both the equivalences in (8) and (13) exactly hold in a noiseless setting only. As such, problem (15) is equivalent to (4) only when the available batch of data is noise-free.

4 Learning Explicit DDPC

To derive its explicit solution, the problem in (15) is manipulated to obtain a *multi-parametric Quadratic Program* (mp-QP). As a preliminary step, we condense the unknowns of (15) into a single variable:

$$G(t) = \begin{bmatrix} G_K x(t) \\ G_f \end{bmatrix} \in \mathbb{R}^{2(T-L)}. \quad (16)$$

¹ Notice that the term in the cost that depends only on $x(t)$ has been neglected. This can be done without loss of generality, as the optimal solution does not change.

Accordingly, we can recast (15) as the following mp-QP

$$\underset{G(t)}{\text{minimize}} \quad (G(t))^\top \mathcal{H}_d G(t) \quad (17a)$$

$$\text{s.t.} \quad \Xi_d G(t) + \Psi x(t) \leq \mathcal{D}, \quad (17b)$$

$$\Theta_d G(t) = \begin{bmatrix} x(t) \\ 0 \end{bmatrix}, \quad (17c)$$

where

$$\mathcal{H}_d = \mathcal{X}^\top \mathcal{Q} \mathcal{X} + \mathcal{V}^\top \mathcal{R} \mathcal{V}, \quad (18a)$$

$$\Xi_d = \begin{bmatrix} C_x & 0 \\ 0 & C_x \end{bmatrix} \begin{bmatrix} 0 \\ \mathcal{X} \end{bmatrix} + C_u \mathcal{V}, \quad \Psi = \begin{bmatrix} C_x & 0 \\ 0 & C_x \end{bmatrix} \begin{bmatrix} I \\ 0 \end{bmatrix}, \quad (18b)$$

$$\Theta_d = \text{diag}([X_{0,T-L-1}, X_{0,T-L-1}]), \quad (18c)$$

and

$$\mathcal{X} = \begin{bmatrix} X_{1,L,T-L} & X_{1,L,T-L} \\ U_{0,L,T-L-1} & U_{0,L,T-L-1} \end{bmatrix},$$

By focusing on \mathcal{H}_d in (17a), it can be proven that this weighting matrix satisfies the following lemma.

Lemma 2 (Features of \mathcal{H}_d) *Under Assumptions 1 and 3, \mathcal{H}_d is positive semi-definite.*

Proof: This is a direct consequence of Assumptions 1 and 3, for which $\mathcal{V}^\top \in \mathbb{R}^{2(T-L) \times mL}$ is not full row rank. \square

As the cost should be strictly convex for a unique explicit solution to be retrieved, this feature of \mathcal{H}_d prevents us from deriving the explicit law. To overcome this limitation, we introduce a regularization term in the cost of (17), thus replacing the weight \mathcal{H}_d with:

$$\mathcal{H}_d^\gamma = \frac{1}{2} (\mathcal{H}_d + \gamma I), \quad (19)$$

where $\gamma > 0$ is a hyper-parameter to be tuned². The data-driven control problem then corresponds to the *regularized* mp-QP:

$$\underset{G(t)}{\text{minimize}} \quad (G(t))^\top \mathcal{H}_d^\gamma G(t) \quad (20a)$$

$$\text{s.t.} \quad \Xi_d G(t) + \Psi x(t) \leq \mathcal{D}, \quad (20b)$$

$$\Theta_d G(t) = \begin{bmatrix} x(t) \\ 0 \end{bmatrix}. \quad (20c)$$

² The cost has been normalized to ease the subsequent derivations.

4.1 Derivation of the explicit DDPC law

The introduction of the regularizer in (20) allows us to derive the explicit DDPC law through the manipulation of the *Karush-Kuhn-Tucker* (KKT) conditions associated with the new DDPC problem. To ease the computations, let us consider the following further assumption.

Assumption 4 (Non-degenerate constraints) *The active constraints of (20) are linearly independent.*

Based on this assumption, we now follow the same steps used to derive the explicit model-based predictive control law in [3].

The KKT conditions for the regularized DDPC problem in (20) are:

$$\mathcal{H}_d^\gamma G(t) + \Xi_d^\top \lambda + \Theta_d^\top \mu = 0, \quad (21a)$$

$$\lambda^\top (\Xi_d G(t) + \Psi x(t) - \mathcal{D}) = 0, \quad (21b)$$

$$\lambda \geq 0, \quad (21c)$$

$$\Xi_d G(t) + \Psi x(t) - \mathcal{D} \leq 0, \quad (21d)$$

$$\Theta_d G(t) - \begin{bmatrix} x(t) \\ 0 \end{bmatrix} = 0, \quad (21e)$$

where λ and μ are the Lagrange multipliers associated with inequality and equality constraints in (20b) and (20c), respectively. Let us focus on the i -th set of active constraints only, distinguishing between the Lagrange multipliers associated with a given active and inactive inequality constraints. We respectively denote them as $\tilde{\lambda}_i$ and $\bar{\lambda}_i$. It is straightforward to notice that the combination of (21b) and (21c) leads to the following condition on $\bar{\lambda}_i$:

$$\bar{\lambda}_i = 0.$$

By merging (21b) and (21e) for the i -th set of active constraints, it is also straightforward to show that the optimal solution $G_i(t)$ satisfies

$$\Phi_{d,i} G_i(t) - \tilde{S}_i x(t) - \tilde{W}_i = 0, \quad (22)$$

where

$$\Phi_{d,i} = \begin{bmatrix} \tilde{\Xi}_{d,i} \\ \Theta_d \end{bmatrix}, \quad \tilde{S}_i = \begin{bmatrix} -\tilde{\Psi}_i \\ I \\ 0 \end{bmatrix}, \quad \tilde{W}_i = \begin{bmatrix} \tilde{\mathcal{D}}_i \\ 0 \end{bmatrix}$$

and $\tilde{\Xi}_{d,i}$, $\tilde{\Psi}_i$ and $\tilde{\mathcal{D}}_i$ are the rows of Ξ_d , Ψ and \mathcal{D} coupled with the considered set active constraints. By leveraging on (21a), we can now express our optimization variable $G_i(t)$ as a function of the Lagrange multipliers, *i.e.*,

$$G_i(t) = -(\mathcal{H}_d^\gamma)^{-1} \underbrace{\begin{bmatrix} \tilde{\Xi}_{d,i}^\top & \Theta_d^\top \end{bmatrix}}_{\Phi_{d,i}^\top} \tilde{\Lambda}_i, \quad (23)$$

where

$$\tilde{\Lambda}_i = \begin{bmatrix} \tilde{\lambda}_i \\ \mu \end{bmatrix}.$$

We can now replace the latter into (22) to obtain an explicit expression of the Lagrange multipliers as functions of the matrices characterizing (20):

$$\tilde{\Lambda}_i = - \underbrace{\left[\Phi_{d,i} (\mathcal{H}_d^\gamma)^{-1} \Phi_{d,i}^\top \right]^{-1}}_{\Upsilon_{d,i}} \left(\tilde{S}_i x(t) + \tilde{W}_i \right). \quad (24)$$

In turn, this allows us to explicitly retrieve $G_i(t)$ as:

$$G_i(t) = (\mathcal{H}_d^\gamma)^{-1} \Phi_{d,i}^\top \Upsilon_{d,i} \left(\tilde{S}_i x(t) + \tilde{W}_i \right), \quad (25)$$

and the associated optimal input sequence as

$$U_i(x(t)) = \mathcal{V} (\mathcal{H}_d^\gamma)^{-1} \Phi_{d,i}^\top \Upsilon_{d,i} \left(\tilde{S}_i x(t) + \tilde{W}_i \right). \quad (26)$$

Thus, the input to be fed to the system when the i -th set of constraints is active is defined as:

$$u_i(x(t)) = [U_i(x(t))]_{1:m}. \quad (27)$$

Through (21c) and (21d), we can finally define the polyhedral region associated with the considered combination of active constraints, which is dictated by the following inequalities:

$$\Upsilon_{d,i} \left(\tilde{S}_i x(t) + \tilde{W}_i \right) \leq 0, \quad (28a)$$

$$\Xi_d (\mathcal{H}_d^\gamma)^{-1} \Phi_{d,i}^\top \Upsilon_{d,i} \left(\tilde{S}_i x(t) + \tilde{W}_i \right) + \Psi x(t) - \mathcal{D} \leq 0. \quad (28b)$$

The complete data-driven expression for (11) is then straightforwardly obtained by following the above steps for all possible combinations of the active constraints. This operation ultimately yields an optimal input sequence $U(x(t))$ of the form:

$$U(x(t)) = \begin{cases} U_1(x(t)), & \text{if } \mathcal{F}_{d,1} x(t) \leq \mathcal{E}_{d,1}, \\ \vdots \\ U_M(x(t)), & \text{if } \mathcal{F}_{d,M} x(t) \leq \mathcal{E}_{d,M}, \end{cases} \quad (29)$$

where M is given by the number of possible combinations of active constraints, U_i corresponds to (26), for all $i \in \{1, \dots, M\}$, while $\{\mathcal{F}_{d,i}, \mathcal{E}_{d,i}\}_{i=1}^M$ can be easily obtained from (28). Consequently, the input to be fed to \mathcal{S} starting from $x(t)$ can be retrieved by evaluating the PWA law

$$u(x(t)) = \begin{cases} u_1(x(t)), & \text{if } \mathcal{F}_{d,1} x(t) \leq \mathcal{E}_{d,1}, \\ \vdots \\ u_M(x(t)), & \text{if } \mathcal{F}_{d,M} x(t) \leq \mathcal{E}_{d,M}, \end{cases} \quad (30)$$

with $u_i(x(t))$ given by (27), for $i = 1, \dots, M$.

Algorithm 1 Offline construction of the explicit law

Input: Dataset \mathcal{D}_T ; penalties $Q, P \succeq 0$; $R \succ 0$; horizon $N > 0$; constraint matrices C_x, C_u, d ; regularization parameter $\gamma > 0$.

1. **Construct** the data-based matrices $X_{1,L,T-L}, U_{0,L,T-L-1}, X_{0,T-L-1}$.
 2. **Build** $\mathcal{H}_d^\gamma, \Xi_d, \Psi, \Theta_d$ in (18a) based on the cost and constraints of the DDPC problem.
 3. **Find** all possible combinations of active constraints.
 4. **For each** combination, **isolate** the matrices $\tilde{\Xi}_d, \tilde{W}$ and \tilde{S} characterizing (22)
 5. **If not** all rows of $\tilde{\Xi}_d$ are **linearly independent**, **handle** the degeneracy, *e.g.*, as in [3].
 6. **Find** the PWA explicit law by retrieving (26)-(28) **for all** possible combinations of active constraints.
 7. **Merge** polyhedral regions as in [2].
 8. **Extract** the first component of the optimal input sequence $U(x(t))$.
-

Output: Optimal input $u(x(t))$.

Remark 2 (On Assumption 4) *Although introduced to ease computations, we remark that Assumption ?? is not restrictive. Indeed, degenerate cases can be straightforwardly handled via existing approaches, e.g., see [3].*

Remark 3 (Data-driven and model-based) *Within a noiseless setting, the results in Theorem 2 and the one-to-one correspondence between the chosen parameterization of the control law in (11) and its model-based counterpart guarantee the equivalence between (4) and (15). Therefore, when there is no noise, the data-driven explicit controller coincides with the E-MPC law as $\gamma \rightarrow 0$.*

4.2 Implementing Explicit DDPC

Based on the available batch of data and the features of the considered predictive control problem, the explicit DDPC law can be completely retrieved *offline* from the available measurements, as summarized in Algorithm 1.

Given the data, one has to initially construct the Hankel matrices needed to build the DDPC problem (see steps 1-2). Once all the possible combinations of active constraints have been detected at step 3 and degenerate scenarios have been handled (see step 5), at step 6 the local controllers and the associated polyhedral regions are retrieved according to (26)-(28). Lastly, at step 7, the optimal control sequence is simplified, by merging polyhedral regions whenever possible. After this step, the explicit optimal input can simply be retrieved by extracting the first element of the input sequence $U(x(t))$ (see step 8).

Once the explicit DDPC law has been retrieved, the computation of the optimal input at each time instant simply

consists of a function evaluation. Specifically, one has to (i) search for the polyhedral region the current state $x(t)$ belongs to, and (ii) apply the corresponding parametric law. We stress that this computational advantage is retained for simple control problems only (*i.e.*, for short prediction horizon and small systems). Indeed, the complexity of the PWA law is known to rapidly increase [7] with the one of the DDPC problem to be solved, analogously to the model-based case.

4.3 Explicit data-driven predictive control and closed-loop stability

When designing a controller in a data-driven setting, it is crucial to check the stability of the resulting closed-loop system before the controller deployment. Towards this objective, we now show how the peculiar features of the explicit data-driven predictive control can be leveraged in combination with existing techniques to devise an off-line, data-driven stability test. To this end, let us assume that the i -th set of constraints is active and consider the following multi-step ahead closed-loop model:

$$\hat{X}(t) = \mathcal{X}(\mathcal{H}_d^\gamma)^{-1} \Phi_{d,i}^\top \Upsilon_{d,i} \tilde{S}_i x(t) + \mathcal{X}(\mathcal{H}_d^\gamma)^{-1} \Phi_{d,i}^\top \Upsilon_{d,i} \tilde{W}_i, \quad (31)$$

obtained by combining (13) with the result of our explicit derivation in (25), where \hat{X} stacks the state predicted by the learned closed-loop model, *i.e.*,

$$\hat{X}(t) = \begin{bmatrix} \hat{x}(t+1) \\ \vdots \\ \hat{x}(t+L) \end{bmatrix}.$$

From (31), we can then isolate the learned one-step ahead closed-loop model, namely

$$\hat{x}(t+1) = A_{d,i}^{cl} x(t) + f_{d,i}^{cl}, \quad (32a)$$

where

$$A_{d,i}^{cl} = [\mathcal{X}(\mathcal{H}_d^\gamma)^{-1} \Phi_{d,i}^\top \Upsilon_{d,i} \tilde{S}_i]_{1:n,1:n}, \quad (32b)$$

$$f_{d,i}^{cl} = [\mathcal{X}(\mathcal{H}_d^\gamma)^{-1} \Phi_{d,i}^\top \Upsilon_{d,i} \tilde{W}_i]_{1:n}. \quad (32c)$$

When performed for each mode $i \in \{1, \dots, M\}$, these manipulations allow us to retrieve the data-driven closed-loop transition matrix $A_{d,i}^{cl}$ for each $i \in \{1, \dots, M\}$. Retrieving these matrices ultimately enables us to apply model-based techniques, *e.g.*, the ones presented in [22], to shed a light on the features of the final closed-loop. As an example, one can search for a matrix $\mathcal{P} \in \mathbb{R}^{n \times n}$ satisfying the following sufficient conditions for asymptotic stability:

$$\mathcal{P} > 0 \quad (33a)$$

$$(A_{d,i}^{cl})^\top \mathcal{P} A_{d,i}^{cl} - \mathcal{P} < 0, \quad i = 1, \dots, M, \quad (33b)$$

or, alternatively, look for the set of matrices $\{\mathcal{P}_i \in \mathbb{R}^{n \times n}\}_{i=1}^M$ verifying the following LMIs:

$$\mathcal{P}_i > 0, \quad i = 1, \dots, M, \quad (34a)$$

$$(A_{d,i}^{cl})^\top \mathcal{P}_j A_{d,i}^{cl} - \mathcal{P}_j < 0, \quad i, j = 1, \dots, M, \quad (34b)$$

which are also sufficient conditions for asymptotic stability.

Remark 4 (On the tuning of P in (4a)) When \mathcal{S} is known to be open-loop stable, the terminal weight $P \succeq 0$ in (4a) is generally selected as the solution of the Lyapunov equation

$$P = A^\top P A + Q.$$

This equation can be directly translated into its data-driven counterpart by exploiting [12] as follows:

$$P = A_d P A_d + Q, \quad (35)$$

with

$$A_d = X_{1,T} \begin{bmatrix} X_{0,T-1} \\ U_{0,1,T-1} \end{bmatrix}^\dagger \begin{bmatrix} I \\ 0 \end{bmatrix},$$

thus providing a data-based approach for the selection of this parameter.

Remark 5 (Hyper-parameter tuning) The possibility of performing an off-line data-based stability check on the data-driven explicit law can be useful to preliminarily assess the effects of different choices of the tuning parameters Q , R and P in (4a) and γ in (19), allowing one to discard the ones resulting in a failure of the data-driven stability tests.

4.4 Regularization and noise handling

As highlighted in Remark 1, all the equivalences we rely on to derive the explicit predictive control law are verified when \mathcal{D}_T is noiseless. However, in practice, Ω in Assumption 2 is generally a non-zero matrix.

To cope with noisy data, we follow the footsteps of [5, 13] and propose to leverage on the regularization term introduced in (19). Indeed, as in standard ridge regression [18], this additional element of the cost steers all the components of $G(t)$ towards small values. In turn, this potentially limits the impact of noise on the constraints in (20b)-(20c) and, thus, on the final explicit law. This shrinkage effect is modulated by the regularization parameter γ , with the reduction in the magnitude of $G(t)$ being stronger whenever large values of γ are considered. At the same time, γ implicitly changes the balance between the penalties in the original data-driven control problem in (15), with excessively high values of γ potentially driving the explicit data-driven controller far away from its model-based counterpart. The

choice of this hyper-parameters thus becomes an important *tuning-knob* of the approach, requiring one to trade-off between handling noise and keeping explicit DDPC as close as possible to the implicit DDPC problem.

Although the stability checks in (33)-(34) can be used to have a preliminary assessment on the effect of different choices of γ , at the moment this balance can only be attained through closed-loop trials for several values of γ . Such a procedure allows one to ultimately select the hyper-parameter that best fits one's needs, at the price of requiring closed-loop experiments that can be rather safety-critical in practice, especially when a simulator of the system is not available.

Whenever multiple experiments can be performed by feeding the plant with the same input sequence \mathcal{U}_T , the burden associated to the choice of γ can be alleviated by exploiting the features of Assumption 2 itself. In this scenario, one can indeed replace \mathcal{D}_T with the averaged dataset $\bar{\mathcal{D}}_T = \{\mathcal{U}_T, \bar{\mathcal{Y}}_T\}$, where $\bar{\mathcal{Y}}_T = \{\bar{y}_t\}_{t=0}^T$ and

$$\bar{y}_t = \frac{1}{N} \sum_{i=1}^N y_t^{(i)}, \quad (36)$$

with $y_t^{(i)}$ denoting the output of the i -th experiment. Since the noise is assumed to be zero mean, the law of large numbers asymptotically yields

$$\bar{y}_t \xrightarrow{N \rightarrow \infty} y_t^0. \quad (37)$$

As such, when the number N of experiments increases, the role of γ in handling noise is progressively less dominant. In this case, γ should then be used *only* to make the DDPC problem well-defined. Any small $\gamma > 0$ is acceptable for this purpose.

5 Numerical examples

The performance of the explicit predictive controller are now assessed on two benchmark examples: (i) the regulation to zero of the stable open-loop system of [3], for the case when the state is fully measurable; and (ii) the altitude control of a quadcopter. Since the last example features an open-loop unstable linearized plant, data are collected in closed-loop, by assuming that the drone is stabilized by an (unknown) pre-existing controller. In both the examples, the level of noise acting on the measured states is assessed through the averaged Signal-to-Noise-Ratio (SNR):

$$\overline{\text{SNR}} = \frac{1}{n} \sum_{i=1}^n 10 \log_{10} \left(\frac{\sum_{t=0}^T (x_i(t) - w_i(t))^2}{\sum_{t=0}^T (w_i(t))^2} \right), \text{ [dB]} \quad (38)$$

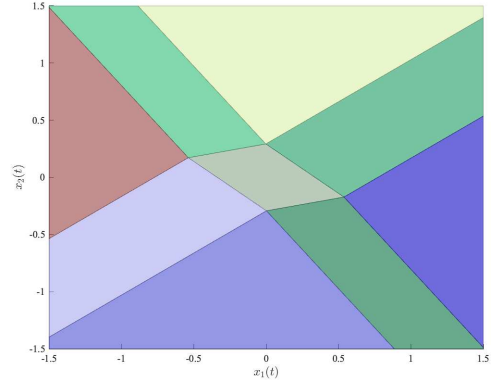
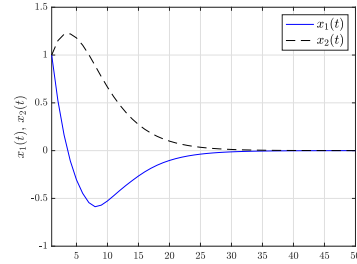
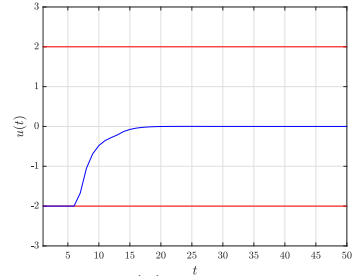


Figure 1. Open-loop stable benchmark: polyhedral partition of the explicit data-driven law.



(a) State trajectories.



(b) Input

Figure 2. Open-loop stable benchmark: state and input trajectories obtained with the explicit data-driven law.

where $x_i(t)$ and $w_i(t)$ denote the i -th components of the state and the measurement noise, respectively. All computations are carried out on an Intel Core i7-7700HQ processor, running MATLAB 2019b.

5.1 Open-loop stable benchmark system

Let us consider the benchmark system described by:

$$\mathcal{S} : x(t+1) = \begin{bmatrix} 0.7326 & -0.0861 \\ 0.1722 & 0.9909 \end{bmatrix} x(t) + \begin{bmatrix} 0.0609 \\ 0.0064 \end{bmatrix} u(t). \quad (39)$$

Our goal is to regulate both components of the state to zero, while enforcing the following box-constraint on the

input:

$$-2 \leq u(k) \leq 2. \quad (40)$$

Towards this goal, we collect a set \mathcal{D}_T of $T = 100$ input/state pairs, by feeding \mathcal{S} with an input sequence uniformly distributed within the interval $[-5, 5]$. According to Assumption 2, the measured states are corrupted by an additive zero-mean white noise sequence, with variance yielding $\overline{\text{SNR}} = 20$ dB. The parameters characterizing the DDPC problem to be solved are selected as in [3], namely $L = 2$, $Q = I$, $R = 0.01$ and P is chosen as the solution of the data-driven Lyapunov in (35). By setting $\gamma = 10$, the partition associated with the explicit data-driven predictive controller³ is the one shown in Figure 1, which approximately correspond to that reported in [3]⁴. Prior to the controller deployment, we have performed the data-based closed-loop stability check in (33), resulting in⁵

$$\mathcal{P} = \begin{bmatrix} 24.8695 & 10.5595 \\ 10.5595 & 43.2657 \end{bmatrix} \succ 0.$$

This indicates that the explicit law preserves the stability of the open-loop system. Figure 2 report the trajectories of the state and the optimal input obtained over a noise-free closed-loop tests with the explicit data-driven law, which confirm its effectiveness and validate the result of the data-driven stability check.

We now assess the sensitivity of the explicit controller to the choice of γ over 20 Monte-Carlo realizations of the batch datasets \mathcal{D}_T , for different noise levels. This evaluation is performed by looking at the cost of the controller over the same noiseless closed-loop test of length $T_v = 50$ considered previously, *i.e.*,

$$\mathcal{J} = \sum_{t=0}^{T_v} \|x(t)\|_Q^2 + \|u(t)\|_R^2. \quad (41)$$

As shown in Figure 3, the value of γ that leads to the minimum closed-loop cost tends to decrease when the noise level increases, supporting our considerations in Section 4.4. At the same time, by properly selecting γ we attain a cost $\mathcal{J} = 16.37 \pm 0.58$, which is generally close to the oracle $\mathcal{J}^{\mathcal{O}} = 15.77$, which is the one achieved by the *oracle* law, *i.e.*, the model-based predictive controller obtained by exploiting the true model of \mathcal{S} . These results additionally show that, for increasing levels of noise, the choice of γ becomes more challenging, since the range of values leading to the minimum \mathcal{J} progressively shrinks. Note that, when γ is excessively small the optimal input is always zero and \mathcal{S} evolves freely⁶.

³ The partition is plotted thanks to the Hybrid Toolbox [1].

⁴ The negligible differences with respect to the model-based partition are due to the noise on the batch data.

⁵ The LMIs in (33) are solved with CVX [16, 17].

⁶ This behavior is also observed for $\gamma \leq 10^{-6}$, $\gamma \leq 10^{-4}$ and

We additionally evaluate the effect of averaging, by looking at the performance index \mathcal{J} in (41) over 30 Monte-Carlo data-collections for an increasing number N of repeated experiments of length $T = 100$. The measurements are affected by noise, yielding $\overline{\text{SNR}} = 20$ dB. Figure 4 shows that the use of the averaged dataset $\overline{\mathcal{D}}_T$ has a similar effect to a reduction of the noise level. Indeed, for increasing N the optimal γ slowly shifts towards smaller values, thus further implying the gradual reduction in the impact of γ on noise handling.

5.2 Altitude control of a quadcopter

As a final case study, we consider a nonlinear system, namely the problem of controlling the altitude of a quadcopter, to perform landing or take-off maneuvers. To this end, we exploit the same simulator used in [15] to collect the data and to carry out the closed-loop experiments with the learned explicit law. Let $z(t)$ [m] be the altitude of the quadcopter, $v_z(t)$ [m/s] be its vertical velocity and $(\theta(t), \phi(t), \psi(t))$ [deg] its roll, pitch and yaw angles at time t . Both the the altitude $z(t)$ [m] and the vertical velocity $v_z(t)$ are assumed to be measured, with the measurement being corrupted by a zero-mean white noise, resulting in $\overline{\text{SNR}} = 30$ dB over these two outputs. As this system is open-loop unstable, the data collection phase is carried out in closed-loop for 20 [s] at a sampling rate of 40 [Hz], by using the four *proportional derivative* (PD) controllers introduced in [15]. The altitude set point used at this stage is generated uniformly at random in the interval $[0, 4]$ [m]. The set points for all the attitude angles are instead selected as slowly variable random signals within the interval $[-0.2, 0.2]$ [rad]. These choices yield a dataset \mathcal{D}_T of length $T = 800$, that satisfies Assumption 1 and allows us to retain information on possible non-zero angular configurations.

The three attitude controllers introduced in [15] are further retained in testing to keep the attitude angles at zero and to decouple the altitude dynamics from that of the other state variables. Within this setting, the explicit data-driven law is designed by imposing $L = 5$, $Q = P = \text{diag}([1, 0.1])$, $R = 10^{-5}$ and $\gamma = 1$. To mitigate the effect of the gravitational force, the design and closed-loop deployment of the designed explicit controller are carried out by pre-compensating it. As a result, the input to be optimized is

$$u(t) = \frac{u_z(t)}{m} - g, \quad (42)$$

where $m = 0.5$ [kg] is the mass of the quadcopter, $g = 9.81$ [m/s] is the gravitational acceleration and $u_z(t)$ is the input prior to the compensation. A similar approach is adopted for the control problem to fit our framework

$\gamma \leq 10^{-3}$ when $\overline{\text{SNR}} = 40$ dB, $\overline{\text{SNR}} = 20$ dB and $\overline{\text{SNR}} = 10$ dB, respectively.

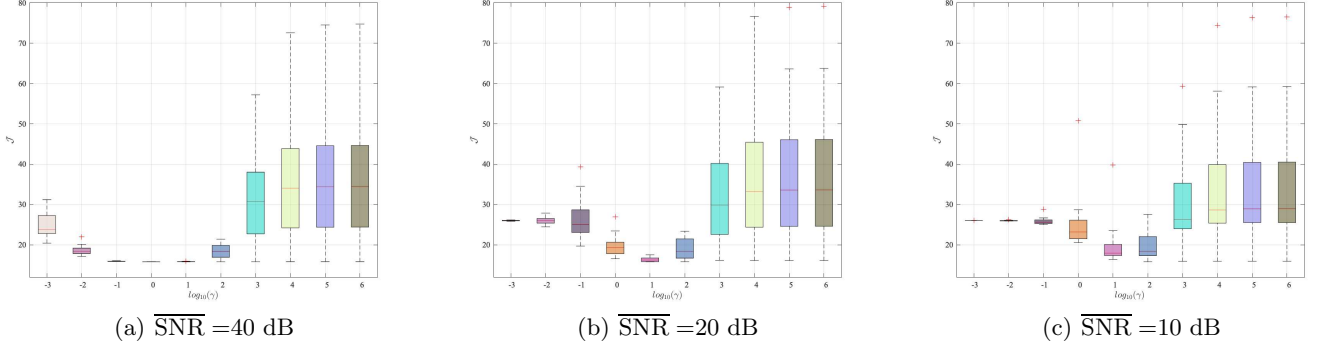


Figure 3. Open-loop stable benchmark: \mathcal{J} in (41) vs $\log_{10}(\gamma)$ for and increasing levels of noise over 30 realizations of \mathcal{D}_T .

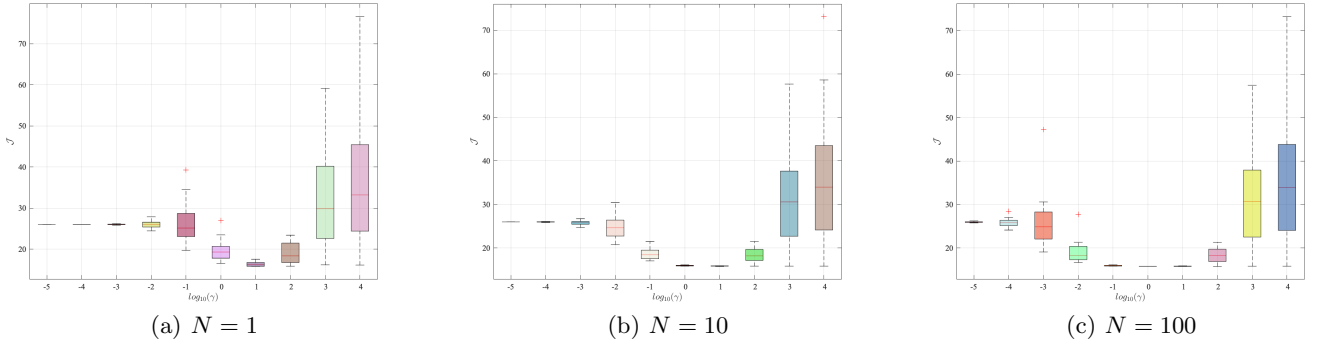


Figure 4. Open-loop stable benchmark: \mathcal{J} in (41) vs $\log_{10}(\gamma)$ for an increasing number of repeated experiments over 30 realizations of \mathcal{D}_T .

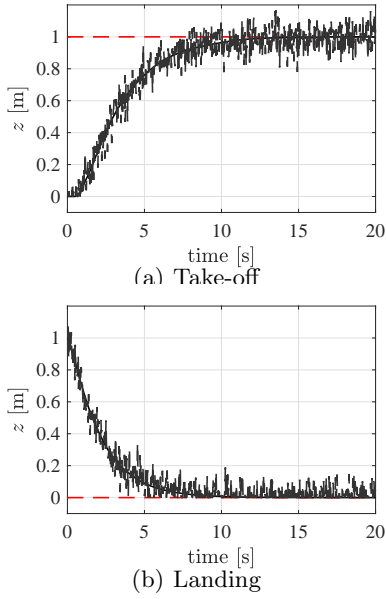


Figure 5. Altitude control: measured (dotted-dashed gray line) and actual (black line) altitude vs reference (dashed red line).

in both landing and take-off scenarios. We thus consider

the reduced state

$$x(t) = \begin{bmatrix} z(t) - \bar{z} \\ v_z(t) \end{bmatrix}, \quad (43)$$

where \bar{z} [m] is the altitude set point. To avoid potential crashes of the quadcopter, in designing the explicit law we impose the following constraint on the state of the system:

$$x_1(t) \geq -\bar{z}, \quad (44)$$

which, in turn, guarantees the altitude to be always non-negative. Meanwhile, the pre-compensated input is constrained to the interval:

$$-9.81 \leq u(t) \leq 9.564, \quad (45)$$

where the lower bound corresponds to a null input and the upper limit is dictated by the maximum power of the motors⁷.

The performance of the learned explicit law attained in take-off and landing are reported in Figure 5. Here we consider closed-loop tests in which the altitude and the

⁷ The reader is referred to [15] for additional details on the system.

vertical velocity are noisy, with the noise acting on the closed-loop measurements sharing the features of that corrupting the batch ones. Despite the noise acting on the initial condition at each step, both maneuvers are successfully performed, thus showing the effectiveness of the retrieved explicit data-driven laws.

6 Conclusions

By leveraging on the known PWA nature of the explicit MPC law within linear quadratic predictive control, in this paper we propose an approach to derive such an explicit controller from data only, without undertaking a full modeling/identification step. Thanks to the formalization of the problem, well-known model-based techniques can be straightforwardly adapted to check the stability of the closed-loop system before deploying the controller.

Future research will be devoted to extend these preliminary results to cases in which the state is not fully measurable, to exploit priors to guarantee practical closed-loop stability by design. Future work will also be devoted to formalize the connections between the explicit solution proposed in this paper and the one introduced in [8], consequently providing a comparative analysis of the two approaches.

References

- [1] A. Bemporad. Hybrid Toolbox - User's Guide, 2004. <http://cse.lab.imtlucca.it/~bemporad/hybrid/toolbox>.
- [2] A. Bemporad, K. Fukuda, and F.D. Torrisi. Convexity recognition of the union of polyhedra. *Computational Geometry*, 18(3):141–154, 2001.
- [3] A. Bemporad, M. Morari, V. Dua, and E.N. Pistikopoulos. The explicit linear quadratic regulator for constrained systems. *Automatica*, 38(1):3–20, 2002.
- [4] J. Berberich, J. Köhler, M.A. Müller, and F. Allgöwer. Data-driven tracking MPC for changing setpoints. *IFAC-PapersOnLine*, 53(2):6923–6930, 2020. 21th IFAC World Congress.
- [5] 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*, 66(4):1702–1717, 2021.
- [6] J. Berberich, J. Köhler, M.A. Müller, and F. Allgöwer. Linear tracking MPC for nonlinear systems part II: The data-driven case. In *arXiv/2105.08567*, 2021.
- [7] F. Borrelli, M. Baotić, J. Pekar, and G. Stewart. On the complexity of explicit mpc laws. In *2009 European Control Conference (ECC)*, pages 2408–2413, 2009.
- [8] V. Breschi, A. Sassella, and S. Formentin. On the design of regularized explicit predictive controllers from input-output data. In *arXiv/2110.11808*, 2021.
- [9] A. Chiuso and G. Pillonetto. System identification: A machine learning perspective. *Annual Review of Control, Robotics, and Autonomous Systems*, 2(1):281–304, 2019.
- [10] J. Coulson, J. Lygeros, and F. Dörfler. Data-enabled predictive control: In the shallows of the DeePC. In *2019 18th European Control Conference (ECC)*, pages 307–312, 2019.
- [11] J. Coulson, J. Lygeros, and F. Dörfler. Regularized and distributionally robust data-enabled predictive control. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 2696–2701, 2019.
- [12] C. De Persis and P. Tesi. Formulas for data-driven control: Stabilization, optimality, and robustness. *IEEE Transactions on Automatic Control*, 65(3):909–924, 2019.
- [13] F. Dörfler, J. Coulson, and I. Markovskiy. Bridging direct & indirect data-driven control formulations via regularizations and relaxations. In *arXiv/2101.01273*, 2021.
- [14] T. Faulwasser, M.A. Müller, and K. Worthmann. Recent advances in model predictive control: Theory, algorithms, and applications. 2021.
- [15] S. Formentin and M. Lovera. Flatness-based control of a quadrotor helicopter via feedforward linearization. In *2011 50th IEEE Conference on Decision and Control and European Control Conference*, pages 6171–6176, 2011.
- [16] M. Grant and S. Boyd. Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, and H. Kimura, editors, *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, pages 95–110. Springer-Verlag Limited, 2008. http://stanford.edu/~boyd/graph_dcp.html.
- [17] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>, March 2014.
- [18] T. Hastie, R. Tibshirani, and J.H. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer series in statistics. Springer, 2009.
- [19] D. Hrovat, S. Di Cairano, H.E. Tseng, and I.V. Kolmanovskiy. The development of model predictive control in automotive industry: A survey. In *2012 IEEE International Conference on Control Applications*, pages 295–302, 2012.
- [20] L. Ljung. Perspectives on system identification. *Annual Reviews in Control*, 34(1):1–12, 2010.
- [21] D.Q. Mayne. Model predictive control: Recent developments and future promise. *Automatica*, 50(12):2967–2986, 2014.
- [22] D. Mignone, G. Ferrari-Trecate, and M. Morari. Stability and stabilization of piecewise affine and hybrid systems: an lmi approach. In *Proceedings of the 39th IEEE Conference on Decision and Control (Cat. No.00CH37187)*, volume 1, pages 504–509 vol.1, 2000.
- [23] S.J. Qin and T.A. Badgwell. A survey of industrial model predictive control technology. *Control Engineering Practice*, 11(7):733–764, 2003.
- [24] J.B. Rawlings. Tutorial overview of model predictive control. *IEEE control systems magazine*, 20(3):38–52, 2000.
- [25] A. Sassella, V. Breschi, and S. Formentin. Learning explicit predictive controllers: theory and applications. In *arXiv/2108.08412*, 2021.
- [26] J.C. Willems, P. Rapisarda, I. Markovskiy, and B.L.M. De Moor. A note on persistency of excitation. *Systems & Control Letters*, 54(4):325–329, 2005.