

Implicit predictors in regularized data-driven predictive control

Manuel Klädtke and Moritz Schulze Darup

Abstract—We introduce the notion of implicit predictors, which characterize the input-(state)-output prediction behavior underlying a predictive control scheme, even if it is not explicitly enforced as an equality constraint (as in traditional model or subspace predictive control). To demonstrate this concept, we derive and analyze implicit predictors for some basic data-driven predictive control (DPC) schemes, which offers a new perspective on this popular approach that may form the basis for modified DPC schemes and further theoretical insights.

Index Terms—Data-driven control, predictive control, parametric optimization

I. INTRODUCTION

Data-driven predictive control (DPC) is an increasingly popular control approach that utilizes linear combinations of collected trajectory data to make predictions instead of relying on a system model (see, e.g., [1]–[3]). While exact predictions and equivalence to model predictive control (MPC) are typically only established for linear time-invariant (LTI) systems and exact data (with some nonlinear extensions available; see, e.g., [2], [4]), modified versions of DPC show promising results even in the absence of these requirements (e.g., [1], [3], [5], [6]). Most proposed DPC schemes use regularizations, which have demonstrated benefits and many different interpretations (see, e.g., [3], [5]), including relationships to other approaches such as subspace predictive control (SPC) [7]. However, many of these interpretations (e.g. [8, Thms. 3 and 4]) are based on limit behavior, i.e., regularization weights approaching zero or infinity. In this study, we aim to characterize the behavior of DPC for finite weights, as this is how it is typically applied in practice. To achieve a general characterization, we do not focus on a specific system class or noise properties but, instead, just make (reasonable) assumptions about the data. To make DPC predictions more accessible and relatable to traditional schemes such as MPC, where consistency with input-state-output predictors is explicitly enforced via equality constraints, we introduce the notion of implicit predictors (specified in Def. 1 below). Implicit predictors can be interpreted as the predictive behavior that is implicitly attributed to the data-generating system by the predictive scheme, and we view it as a central and elucidating object with much to learn from. To show the benefit of this concept, we derive and analyze implicit predictors for a basic regularized DPC scheme, focusing on two proposed

choices of (squared) 2-norm regularization (see [3]), and examine the effect of input and output constraints on these predictors. The paper is organized as follows. First, in Section II, we introduce the notion of implicit predictors, specify the objective of this paper, and summarize important preliminaries. In Section III, we derive implicit predictors for two types of regularized DPC and analyze the impact of constraints on these predictors. Finally, we conclude our work in Section IV and preview future opportunities we envision for implicit predictors.

II. PROBLEM STATEMENT AND PRELIMINARIES

One could argue that (multi-step) predictors, which establish a mapping from an initial state x_0 and an input (sequence) u_f to a predicted output (sequence) y_f , are at the heart of many predictive control schemes. Traditionally, such a predictor $\hat{y}_f(x_0, u_f)$ may be explicitly included in an optimal control problem (OCP) as an equality constraint $y_f = \hat{y}_f(x_0, u_f)$. However, even if not included explicitly (such as in DPC), one may still observe the predictions following a similar pattern, which we formalize as follows.

Definition 1: We call $\hat{y}(x_0, u_f)$ an *implicit predictor* for an OCP if including the constraint $y_f = \hat{y}(x_0, u_f)$ does not alter the (set of) minimizers (u_f^*, y_f^*) and the optimal value.

For DPC, we interpret this definition as the predictive behavior implicitly attributed to the data-generating system by the DPC scheme. Given this interpretation, we take a somewhat opposing viewpoint to the bias-variance hypothesis in [3, Sec. V.C]. In fact, we claim that the choice of regularizer (and other parameters of the OCP) fully specifies a model class given by the structure of its resulting implicit predictor. Therefore, this concept should be seen as a tool to analyze existing schemes and explain their behavior. Although this new viewpoint is mainly theoretical in nature, practitioners can use its results to evaluate whether the predictive behavior of a given scheme matches their prior knowledge of the true system properties, and thus select an appropriate scheme similar to traditional model selection. In this spirit, we analyze an existing DPC scheme without making assumptions on the data-generating system class or properties of measurement noise, allowing for a general characterization. To prepare this analysis, we recall preliminaries on DPC and its relation to MPC and SPC, while highlighting the role of predictors in these schemes.

Manuel Klädtke and M. Schulze Darup are with the Control and Cyber-physical Systems Group, Department of Mechanical Engineering, TU Dortmund University, Germany. E-mail correspondence to {manuel.klaedtke, moritz.schulzedarup}@tu-dortmund.de.

This is the conference paper accompanying a contribution to the IEEE Control Systems Letters. The DOI of the journal contribution is 10.1109/LC-SYS.2023.3285104

A. DPC and its relation to MPC

Instead of utilizing a discrete-time state-space model

$$x(k+1) = f(x(k), u(k), k) \quad (1a)$$

$$y(k) = g(x(k), u(k), k) \quad (1b)$$

with input $u \in \mathbb{R}^m$, state $x \in \mathbb{R}^n$, and output $y \in \mathbb{R}^p$ as in traditional MPC, predictions in DPC are realized based on previously collected trajectory data $(u^{(1)}, y^{(1)}), \dots, (u^{(\ell)}, y^{(\ell)})$ via linear combinations

$$\begin{pmatrix} u_{\text{pred}} \\ y_{\text{pred}} \end{pmatrix} = \begin{pmatrix} u^{(1)} \\ y^{(1)} \end{pmatrix} a_1 + \dots + \begin{pmatrix} u^{(\ell)} \\ y^{(\ell)} \end{pmatrix} a_\ell = \mathcal{D}a.$$

Here, the dimensions of the data matrix $\mathcal{D} \in \mathbb{R}^{L(m+p) \times \ell}$ and generator vector $a \in \mathbb{R}^\ell$ are specified by the length L of recorded (as well as predicted) trajectories and the number ℓ of data trajectories used for predictions. The rationale for this procedure is given by a result for linear time-invariant (LTI) systems, which can be described, e.g., by the specification

$$x(k+1) = Ax(k) + Bu(k) \quad (2a)$$

$$y(k) = Cx(k) + Du(k), \quad (2b)$$

of (1). For these LTI systems and assuming L is greater than the lag of the system, $\text{image}(\mathcal{D})$ is equivalent to the set of all possible system trajectories if and only if [9]

$$\text{rank}(\mathcal{D}) = Lm + n. \quad (3)$$

Note that the condition (3) not only signifies a minimum rank for data-driven predictions, but also the maximum rank that the data matrix \mathcal{D} can achieve for exact data. For the case that the individual trajectories $(u^{(i)}, y^{(i)})$ are time-shifted sections of a single long trajectory, a popular sufficient condition for (3) is given by Willems' fundamental lemma [10]. To include the current initial condition of the system as a starting point for predicted trajectories, the predicted I/O-sequence is typically partitioned into a past section (u_p, y_p) and a future section (u_f, y_f) with N_p respectively N_f time-steps yielding

$$\begin{pmatrix} u_{\text{pred}} \\ y_{\text{pred}} \end{pmatrix} = \begin{pmatrix} u_p \\ u_f \\ y_p \\ y_f \end{pmatrix} = \begin{pmatrix} U_p \\ U_f \\ Y_p \\ Y_f \end{pmatrix} a = \mathcal{D}a.$$

The past section of a predicted trajectory is then forced to match the I/O-data ξ recorded in the most recent N_p time-steps during closed-loop operation, i.e., the equality constraints

$$\xi = \begin{pmatrix} u_p \\ y_p \end{pmatrix} = \begin{pmatrix} U_p \\ Y_p \end{pmatrix} a = W_p a$$

force any predicted trajectory in (5) to start with the most recently witnessed behavior of the system. In this context, the past trajectory ξ can also be interpreted as the state of a (usually non-minimal) state-space realization of the system and properly specifies its initial condition if N_p is chosen larger or equal to its lag [11].

Remark 1: Although we have introduced the data-driven predictions in an I/O setting, they can be straightforwardly modified to a state-space setting [12], which we will use for visualization of a low dimensional example in Section III.

Now, consider a classical OCP

$$\min_{u(k), x(k), y(k)} \sum_{k=0}^{N_f-1} \|y(k)\|_Q^2 + \|u(k)\|_R^2 \quad (4)$$

$$\text{s.t. } x(0) = x_0,$$

$$x(k+1) = Ax(k) + Bu(k), \quad \forall k \in \{0, \dots, N_f - 2\},$$

$$y(k) = Cx(k) + Du(k), \quad \forall k \in \{0, \dots, N_f - 1\},$$

$$(u(k), y(k)) \in \mathcal{U}_k \times \mathcal{Y}_k, \quad \forall k \in \{0, \dots, N_f - 1\}$$

for MPC with prediction horizon N_f , positive definite weighing matrices $R \in \mathbb{R}^{m \times m}$, $Q \in \mathbb{R}^{p \times p}$, convex constraint sets $\mathcal{U}_k \subseteq \mathbb{R}^m$, $\mathcal{Y}_k \subseteq \mathbb{R}^p$, as well as equality constraints specifying the initial state condition x_0 and predicted behavior based on (2). It has been shown in [1] that DPC given by

$$\min_{u_f, y_f, a} \|y_f\|_Q^2 + \|u_f\|_R^2 \quad (5a)$$

$$\text{s.t. } \begin{pmatrix} \xi \\ u_f \\ y_f \end{pmatrix} = \begin{pmatrix} W_p \\ U_f \\ Y_f \end{pmatrix} a, \quad (5b)$$

$$(u_f, y_f) \in \mathcal{U} \times \mathcal{Y} \quad (5c)$$

with $\mathcal{Q} := \text{blkdiag}(Q, \dots, Q)$, $\mathcal{R} := \text{blkdiag}(R, \dots, R)$, and

$$\mathcal{Y} := \{y_f \in \mathbb{R}^{pN_f} \mid y(k) \in \mathcal{Y}_k, \forall k \in \{0, \dots, N_f - 1\}\}$$

$$\mathcal{U} := \{u_f \in \mathbb{R}^{mN_f} \mid u(k) \in \mathcal{U}_k, \forall k \in \{0, \dots, N_f - 1\}\}$$

based on exact data generated by an LTI system is equivalent to the MPC in (4). Remarkably, and although the original theory behind exact predictions via linear combination of data does not apply to arbitrary nonlinear systems or with noise and disturbances in the data, DPC has shown good closed-loop performance when applied to these cases. A common feature of these successful applications, is the addition of a regularization term $h(a)$ to the cost function, where different choices of $h(a)$ can have different interpretations for its intended effect on the predictions ([1], [3], [5], [6]).

B. The role of predictors and SPC

In the case of linear MPC, its well-known multi-step predictor can be expressed as

$$\hat{y}_{\text{MPC}}(x_0, u_f) = \mathcal{O}x_0 + \mathcal{T}u_f, \quad (6)$$

where \mathcal{O} and \mathcal{T} are often referred to as the extended observability matrix and the impulse response matrix, respectively (see, e.g. [13]). Using (6), we can state the OCP (4) also as

$$\min_{u_f, y_f} \|y_f\|_Q^2 + \|u_f\|_R^2 \quad (7)$$

$$\text{s.t. } y_f = \mathcal{O}x_0 + \mathcal{T}u_f, \quad (u_f, y_f) \in \mathcal{U} \times \mathcal{Y},$$

which may be useful to eliminate optimization variables and to analyze the structure of its solution. Essentially, instead of using the state-space model (2) as a one-step predictor for each time-step in (4), consecutive applications of the one-step predictor yield the multi-step predictor (6) used in (7). Note that (6) trivially acts as an implicit predictor for the original MPC problem (4). By including the constraint $y_f = \hat{y}_{\text{MPC}}(x_0, u_f)$, other constraints and variables can be eliminated, resulting in (7). However, this removal should be seen separately to Definition 1 and is not our focus.

Instead of estimating a state-space model (2) for MPC, an alternative approach based on subspace identification [13] is

given by SPC [7]. Here, a linear multi-step predictor

$$\hat{y}_{\text{SPC}}(\xi, u_f) = K_{\text{SPC}} \begin{pmatrix} \xi \\ u_f \end{pmatrix} \quad (8)$$

can be estimated directly from data as the solution

$$K_{\text{SPC}} := \arg \min_K \left\| Y_f - K \begin{pmatrix} W_p \\ U_f \end{pmatrix} \right\|_F^2 = Y_f \begin{pmatrix} W_p \\ U_f \end{pmatrix}^+ \quad (9)$$

to a least squares problem with the Frobenius norm $\|\cdot\|_F$. Crucially, this form of SPC is related to DPC as follows.

Lemma 1: The SPC predictor (8) is an implicit predictor for the unregularized DPC problem (5) with exact trajectory data generated by an LTI system.

Proof: Follows from [14, Thm. 1]. ■

While the result in [14] is actually stated as (5) being equivalent to SPC, we rather rephrase it in the context of implicit predictors. This is because (5b) cannot be removed (without changing the set of feasible ξ) as it also implies $\xi \in \text{image}(W_p)$, which is not captured by the SPC predictor.

Introducing implicit predictors as explicit equality constraints can be useful to change the solution strategy for some OCPs. This, however, is not our focus in this work. Instead, we want to highlight this notion as a way to think about what kind of behavior is (implicitly) attributed to the system by DPC schemes, even though the predictor is not given as an equality constraint in any of the following cases.

III. IMPLICIT PREDICTORS IN REGULARIZED DPC

While in the case of deterministic LTI systems, adding additional trajectory data, i.e., columns to \mathcal{D} , cannot increase its rank past $\text{rank}(\mathcal{D}) \leq mL + n$, this is typically not the case for systems with noise or nonlinearities. Instead, adding trajectory data generated by an LTI system with some output noise and persistently exciting enough input data (see, e.g., [10]) will almost surely increase its rank until \mathcal{D} has full row rank for wide enough (at least square) data matrices [8, Lem. 3] and similar behavior can be observed for nonlinear systems with (or without) noise. Throughout this section, we will thus make the following assumption.

Assumption 1: The data matrix \mathcal{D} has full row rank.

We want to emphasize that Assumption 1 takes into account the presence of measurement noise, and we do not make any further assumptions about the class of system generating the data. Once the OCP tuning parameters (including regularization) are determined, the solution becomes deterministic with respect to the provided data \mathcal{D} (including potential output noise, which is included in W_p , Y_f and cannot be distinguished from the true output data), regardless of the type of system that generated it. Since the aim of this work is to characterize the predictive behavior and not to compare it to the (unknown) true system dynamic (as e.g., in [8]), this allows for a very general analysis. The assumption thus helps in shifting the focus away from what initially caused the lack of rank deficiency and towards what predictions are being made by DPC based on the given data (including noise). Furthermore, having more data than necessary in (3) is typically preferred, where some authors report good results with \mathcal{D} being square [15, Sect. 5.2.4].

As the assumption aligns with this case, we believe it is a reasonable starting point but future work will consider data matrices with ranks between the practical minimum given by (3) and the maximum given by Assumption 1. Crucially, due to Assumption 1 there is a vector a satisfying (5b) for any triple (ξ, u_f, y_f) , resulting in arbitrary and meaningless predictions if the scheme is not suitably modified. However, adding a regularization $h(a)$ leading to the regularized DPC scheme

$$\min_{u_f, y_f, a} \|y_f\|_{\mathcal{Q}}^2 + \|u_f\|_{\mathcal{R}}^2 + h(a) \quad \text{s.t.} \quad (5b)-(5c) \quad (10)$$

has shown good results with respect to its predictive capabilities. Note that, while (5b) does not restrict the choice of any (ξ, u_f, y_f) in terms of feasibility due to Assumption 1, it still defines a relation to a , which can be used to express the effect of $h(a)$ in the following way.

Lemma 2: Under Assumption 1, the regularized DPC problem (10) is equivalent to

$$\min_{u_f, y_f} \|y_f\|_{\mathcal{Q}}^2 + \|u_f\|_{\mathcal{R}}^2 + h^*(\xi, u_f, y_f) \quad \text{s.t.} \quad (5c) \quad (11)$$

with unique

$$h^*(\xi, u_f, y_f) := \min_a h(a) \quad \text{s.t.} \quad (5b). \quad (12)$$

Proof: The OCP (10) can be trivially decoupled with the outer problem (11) and inner problem (12) by optimizing over one of the optimization variables while treating the others as parameters. The constraint (5c) is irrelevant to the inner problem because (u_f, y_f) act as parameters there. On the other hand, the constraint (5b) is irrelevant to the outer problem because a is eliminated by solving the inner problem and (5b) does not restrict the choice of (ξ, u_f, y_f) due to Assumption 1. While the minimizer $a^*(\xi, u_f, y_f)$ to (12) may be non-unique depending on the choice of $h(a)$ and the number ℓ of data trajectories (i.e., the width of \mathcal{D}), the resulting optimal value of the inner problem $h^*(\xi, u_f, y_f) = h(a^*(\xi, u_f, y_f))$ is trivially unique by optimality. ■

While any trajectory (ξ, u_f, y_f) can be predicted via linear combination of collected data trajectories, the regularizer $h(a)$ assigns to every such trajectory a (possibly non-unique) optimal generator vector $a^*(\xi, u_f, y_f)$ and associated unique cost $h^*(\xi, u_f, y_f)$. Therefore, the regularizer $h(a)$ is, ideally, chosen in a way that incentivizes predictions that seem likely based on the given data in \mathcal{D} by assigning them low (or even zero) cost while avoiding unlikely predictions via high costs.

We will now analyze this behavior for different choices of $h(a)$ and characterize the associated implicit predictors. Note that we treat the influence of equality constraints (5c) on the implied predictor separately in Section III-C.

A. (Squared) 2-norm regularization

Commonly seen in regularized DPC is (squared) 2-norm regularization, where we choose $h(a) = \lambda_a \|a\|_2^2$ with weighing parameter λ_a . Here, the inner problem (12) is a quadratic minimization with linear equality constraints and positive definite Hessian $2\lambda_a I$, which yields the unique minimizer

$$a^*(\xi, u_f, y_f) = \begin{pmatrix} W_p \\ U_f \\ Y_f \end{pmatrix}^+ \begin{pmatrix} \xi \\ u_f \\ y_f \end{pmatrix} \quad (13)$$

and associated cost

$$h^*(\xi, u_f, y_f) = \lambda_a \begin{pmatrix} \xi \\ u_f \\ y_f \end{pmatrix}^\top \left(\begin{pmatrix} W_p \\ U_f \\ Y_f \end{pmatrix} \begin{pmatrix} W_p \\ U_f \\ Y_f \end{pmatrix}^\top \right)^{-1} \begin{pmatrix} \xi \\ u_f \\ y_f \end{pmatrix}.$$

Utilizing a block LDU decomposition of the weighing matrix, one can further show that

$$h^*(\xi, u_f, y_f) = \lambda_a (y_f - \hat{y}_{\text{SPC}})^\top \mathcal{Q}_{\text{reg}} (y_f - \hat{y}_{\text{SPC}}) + \lambda_a \begin{pmatrix} \xi \\ u_f \end{pmatrix}^\top \left(\begin{pmatrix} W_p \\ U_f \end{pmatrix} \begin{pmatrix} W_p \\ U_f \end{pmatrix}^\top \right)^{-1} \begin{pmatrix} \xi \\ u_f \end{pmatrix}, \quad (14)$$

where we introduced

$$\mathcal{Q}_{\text{reg}} := (Y_f (I - \Pi) Y_f^\top)^{-1},$$

$$\Pi := \begin{pmatrix} W_p \\ U_f \end{pmatrix}^\top \left(\begin{pmatrix} W_p \\ U_f \end{pmatrix} \begin{pmatrix} W_p \\ U_f \end{pmatrix}^\top \right)^{-1} \begin{pmatrix} W_p \\ U_f \end{pmatrix}, \quad (15)$$

and occasionally omit the arguments of $\hat{y}_{\text{SPC}}(\xi, u_f)$ for brevity. The expression (14) highlights the connection between DPC and SPC for this regularization. However, this regularizer not only incentivizes predictions that align with the SPC predictor but also introduces a bias associated with the additional quadratic cost term in (ξ, u_f) . This (typically unwelcome) bias justifies the introduction of a projection in the regularizer [3], which we will analyze in Section III-B.

Even though predictions that align with the SPC predictor are incentivized with this regularization, the SPC predictor does not act as an implicit predictor in this case. In contrast to the unregularized deterministic case, where the rank deficiency of \mathcal{D} implies a subspace for predictions, the implicit predictor is instead implied by optimality.

Theorem 3: Consider the regularized DPC problem (10) with regularizer $h(a) = \lambda_a \|a\|_2^2$ and without additional input and output constraints. Under Assumption 1,

$$\hat{y}_{\text{DPC}}(\xi, u_f) = (\lambda_a \mathcal{Q}_{\text{reg}} + \mathcal{Q})^{-1} \lambda_a \mathcal{Q}_{\text{reg}} \hat{y}_{\text{SPC}}(\xi, u_f) \quad (16)$$

is an implicit predictor for this problem.

Proof: The implicit predictor (16) is the minimizer

$$\hat{y}_{\text{DPC}}(\xi, u_f) = \arg \min_{y_f} \|y_f\|_{\mathcal{Q}}^2 + \|u_f\|_{\mathcal{R}}^2 + h^*(\xi, u_f, y_f)$$

to an inner optimization problem, where u_f and ξ act as parameters. Since dropping any terms independent of y_f does not change the minimizer, we can simplify the problem as

$$\arg \min_{y_f} \lambda_a (y_f - \hat{y}_{\text{SPC}})^\top \mathcal{Q}_{\text{reg}} (y_f - \hat{y}_{\text{SPC}}) + y_f^\top \mathcal{Q} y_f$$

$$= \arg \min_{y_f} y_f^\top (\lambda_a \mathcal{Q}_{\text{reg}} + \mathcal{Q}) y_f - 2 \hat{y}_{\text{SPC}}^\top \lambda_a \mathcal{Q}_{\text{reg}} y_f,$$

which yields an unconstrained quadratic minimization problem with the minimizer given by (16). Now, since $\hat{y}_{\text{DPC}}(\xi, u_f)$ is the parametric minimizer for any (ξ, u_f) , the minimizers (u_f^*, y_f^*) to the regularized DPC problem must naturally satisfy the relation $y_f^* = \hat{y}_{\text{DPC}}(\xi, u_f^*)$ for any ξ . Hence, including the equality constraint $y_f = \hat{y}_{\text{DPC}}(\xi, u_f)$ with the regularized DPC problem does not change its optimal value or minimizers, making $\hat{y}_{\text{DPC}}(\xi, u_f)$ an implicit predictor of this OCP. ■

While this implicit predictor does not match the SPC predictor, it still represents a subspace just like the latter and is visualized in Fig. 1 for different regularization weight-

ings λ_a . For very high λ_a , the predictor starts aligning with the SPC predictor (and even matches it for $\lambda_a \rightarrow \infty$) as one would expect from the costs given by (14). For lower λ_a , the implicit predictor tilts towards $\lim_{\lambda_a \rightarrow 0} \hat{y}_{\text{DPC}}(\xi, u_f) = 0$ as seen in Fig. 1.a, which is also reasonable, since it represents the optimal solution to $\arg \min_{y_f} \|y_f\|_{\mathcal{Q}}^2$. These behaviors perfectly match the viewpoint that ‘‘control and identification regularize each other’’ [3] in DPC. However, we emphasize that this means the (implicitly) predicted system behavior of regularized DPC is linear and either consistent with the SPC predictor (for $\lambda_a \rightarrow \infty$) or more optimistic than SPC in the sense that it biases these most likely predictions (in the least squares sense, see (9)) towards more favorable predictions given by the control objective. While this optimism in the predictions may indeed be one reason why DPC performs better than SPC in some case studies (e.g. [3, Fig. 2]), it should be recognized that this also leads to potentially undesirable predictive behavior, which we will highlight when discussing additional constraints in Section III-C.

B. Projection-based (squared) 2-norm regularization

While we only introduced Π in (14) for notation purposes, we next exploit its role as a projection matrix. However, instead of Π itself, we will focus on $I - \Pi$, which is an orthogonal projector on $\ker \left(\begin{pmatrix} W_p \\ U_f \end{pmatrix} \begin{pmatrix} W_p \\ U_f \end{pmatrix}^\top \right)$. Intuitively, these projection matrices can be thought of as splitting any $a = \Pi a + (I - \Pi)a$ into a part Πa that aligns with the SPC solution and a part $(I - \Pi)a$ that does not. In [3], the regularizer $h(a) = \lambda_a \|(I - \Pi)a\|_2^2$ was introduced because, although it penalizes deviations from the SPC predictor just as $h(a) = \lambda_a \|a\|_2^2$, it does not introduce any further bias like the latter in (14). In the following, we will phrase this important result in our framework and subsequently analyze the associated implicit predictor. We first note that the inner problem given by (12) is, again, a quadratic minimization problem with linear constraints and the same (unique) minimizer as before in (13). This can easily be verified by plugging $a^*(\xi, u_f, y_f)$ into the KKT system of (12). Uniqueness of $a^*(\xi, u_f, y_f)$ may not be immediately obvious, since the Hessian $2\lambda_a (I - \Pi)^\top (I - \Pi) = 2\lambda_a (I - \Pi)$ is only positive semi-definite due to $(I - \Pi)$ being a projector on a subspace of \mathbb{R}^ℓ . However, since

$$\ker \begin{pmatrix} W_p \\ U_f \\ Y_f \end{pmatrix} \subset \ker \begin{pmatrix} W_p \\ U_f \end{pmatrix}$$

and $(I - \Pi)$ projects onto the latter kernel, we have $(I - \Pi)a \neq 0$ for any nonzero a in the former kernel, i.e., the Hessian and the constraint matrix only share the trivial nullspace by construction, making the KKT matrix nonsingular and thus $a^*(\xi, u_f, y_f)$ unique [16, Sect. 10.1.1]. The additional cost introduced by the regularizer is therefore given as

$$h^*(\xi, u_f, y_f) = \lambda_a \begin{pmatrix} \xi \\ u_f \\ y_f \end{pmatrix}^\top \begin{pmatrix} W_p \\ U_f \\ Y_f \end{pmatrix}^{+, \top} (I - \Pi) \begin{pmatrix} W_p \\ U_f \\ Y_f \end{pmatrix}^+ \begin{pmatrix} \xi \\ u_f \\ y_f \end{pmatrix},$$

which can be shown to be equivalent to

$$h^*(\xi, u_f, y_f) = \lambda_a (y_f - \hat{y}_{\text{SPC}})^\top \mathcal{Q}_{\text{reg}} (y_f - \hat{y}_{\text{SPC}}) \quad (17)$$

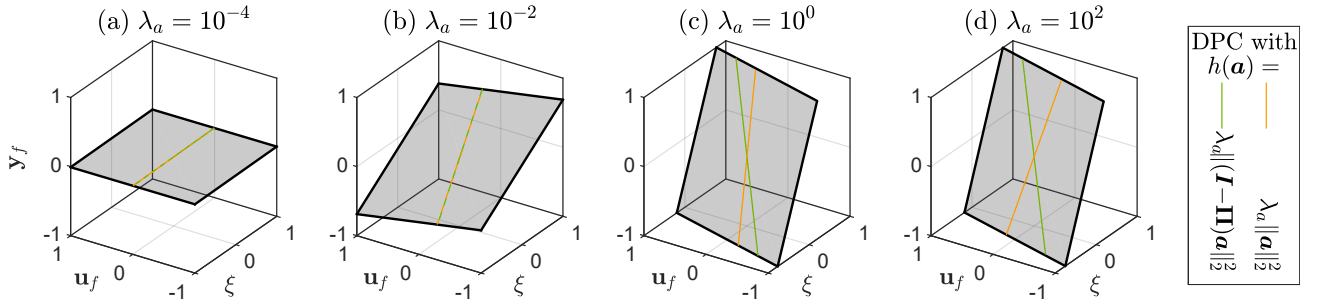


Fig. 1. The implicit predictor (grey) is equal for both analyzed unconstrained DPC schemes and its structure is given by a subspace that “tilts” between $\hat{y}_{\text{DPC}}(\xi, u_f) = 0$ and the SPC predictor (8) depending on λ_a . The parametric DPC solutions $(\xi, u_f^*(\xi), y_f^*(\xi))$ for the different regularizations (green/orange) evolve on this subspace, confirming its validity.

by, again, utilizing block LDU decompositions of the inverses involved in the weighing matrix. The resulting cost is similar to (14), which should not be too surprising, since the projection-based regularization was introduced to remove the bias related to additional costs in (ξ, u_f) and only penalize deviations from $\hat{y}_{\text{SPC}}(\xi, u_f)$, which is exactly captured by (17). Furthermore, this expression allows us to state the following result regarding the implicit predictor, which is, again, implied by optimality.

Theorem 4: Consider the regularized DPC problem (10) with regularizer $h(a) = \lambda_a \|(I - \Pi)a\|_2^2$ and without additional input and output constraints. Under Assumption 1, (16) is an implicit predictor for this problem.

Proof: The proof can be carried out analogously to Theorem 3 by noting that the only difference between the regularizer costs (17) and (14) is irrelevant to the implied predictor since it does not depend on y_f . ■

This result may be surprising to readers familiar with both DPC regularizations, since the minimizers (u_f^*, y_f^*) can differ drastically depending on the chosen regularizer. However, it shows that the predictive aspect of both DPC schemes is the same in the sense that, while (u_f^*, y_f^*) may differ (especially for increasing λ_a , as seen in Fig. 1.d), the relation $y_f^* = \hat{y}_{\text{DPC}}(\xi, u_f^*)$ holds for both as visualized in Fig. 1.

Remark 2: We can easily extend Theorems 3 and 4 towards reference tracking cost $\|y_f - y_{\text{ref}}\|_{\mathcal{Q}}^2$, which adds a linear term in the cost function and thus yields an affine implicit predictor, which is biased towards the unconstrained minimum y_{ref} .

C. Predictions affected by (output) constraints

As previously noted, the implied predictors associated with regularized DPC tend to predict more optimistically (with respect to the control objective) than the SPC predictor. This phenomenon will be further highlighted in this section, where we investigate the effect of additional constraints. In the following, we will restrict our attention to a polyhedral set \mathcal{Y} , which is a common assumption in (linear) predictive control. Similar results can be expected for differently shaped sets, although they are not as nice to characterize accurately. Moreover, we introduce $\mathcal{X} \subseteq \text{dom}(\xi)$, which should not be seen as a constraint on an optimization variable (because ξ is a parameter of the OCP), but instead can be viewed as the




Fig. 2. The implicit predictor (grey) for the analyzed DPC schemes with (output) constraints is a PWA function. For increasing λ_a , the parametric DPC solutions (orange/green, see Fig. 1 for a legend) tend to stay on the segment that matches the unconstrained predictor (if possible).

set of realistically occurring initial conditions in practical operation.

Theorem 5: Consider the regularized DPC problem (10) with regularizer $h(a) = \lambda_a \|a\|_2^2$ or $h(a) = \lambda_a \|(I - \Pi)a\|_2^2$ and with constraint sets \mathcal{U}, \mathcal{Y} , where the latter is a polyhedron. Under Assumption 1, an implicit predictor $\hat{y}_{\text{DPC}}(\xi, u_f)$ can be characterized as the minimizer to the multiparametric quadratic program (mpQP)

$$\arg \min_{y_f} y_f^\top (\lambda_a \mathcal{Q}_{\text{reg}} + \mathcal{Q}) y_f - 2\hat{y}_{\text{SPC}}^\top(\xi, u_f) \lambda_a \mathcal{Q}_{\text{reg}} y_f \quad (18)$$

s.t. $y_f \in \mathcal{Y}$

with parameters (ξ, u_f) , which is a continuous piecewise affine (PWA) function on partitions of the domain $\mathcal{X} \times \mathcal{U}$.

Proof: Derivation of the mpQP can be carried out analogously to Theorems 3 and 4, by noting that only the output constraints are explicitly relevant to the implied predictor and other constraints only specify the domain of its parametric solution. The structure of the minimizer follows from standard mpQP results (see, e.g., [17, Cor. 5.2]). ■

While PWA functions are not uncommon in (linear) MPC (mainly in relation to its explicit solution [17]), we want to stress that the PWA function characterized in Theorem 5 is in its interpretation very different to an optimal control law $u_f^*(\xi)$, since it characterizes the predictive behavior we associate with regularized DPC. While this implicit predictor coincides with (16) on the set

$\mathcal{Z} := \{(\xi, u_f) \in \mathcal{X} \times \mathcal{U} \mid (16) \text{ satisfies } \hat{y}_{\text{DPC}}(\xi, u_f) \in \mathcal{Y}\}$, i.e., where the unconstrained solution naturally satisfies the output constraints, it changes drastically wherever this coincidence does not occur, as seen in Fig. 2. However, at this point, a discussion on the uniqueness of implicit predictors is due. While the implicit predictors derived so far are kept as general as possible (i.e., they are independent of input cost weights \mathcal{R} and constraints \mathcal{U}), by Definition 1 agreement with the predictor needs to occur only at optimal points (u_f^*, y_f^*) for it to be valid. In other words, if the actual solution to the constrained regularized DPC problem (10) naturally stays on \mathcal{Z} , i.e., we have $(\xi, u_f) \in \mathcal{Z}$ for all $\xi \in \mathcal{X}$, then its predictive behavior is equally well explained by both the linear predictor (16) and the PWA solution to (18). Moreover, the DPC solution favors this phenomenon for increasing λ_a , as highlighted in Fig. 2, since deviations of the unconstrained predictor, in this case, are increasingly penalized by the regularizer. However, it should be obvious that (even with very large λ_a) a match of both predictors cannot occur for such ξ , where (16) cannot be satisfied for any $(u_f, y_f) \in \mathcal{U} \times \mathcal{Y}$. In that case, the DPC solution will definitely change its predictive behavior in favor of satisfying the output constraints in accordance with (18). In other words, regularized DPC as in (10) will simply provide a “very unlikely” solution if all “more likely” solutions (in the sense of proximity to its unconstrained predictions or even the SPC predictor) are infeasible. Consequentially, from the viewpoint of an explicit solution to the OCP [17], where \mathcal{X} acts as a constraint of the parameter ξ , its set of feasible states is equal to \mathcal{X} regardless of other constraints \mathcal{U}, \mathcal{Y} .

Finally, we want to highlight that the implicit predictor in Theorem 5 is independent of the input constraints, since \mathcal{U} only affects its domain but, in contrast to \mathcal{Y} , not its structure. Therefore, the linear implicit predictor (16) remains valid in the absence of output constraints (i.e., $\mathcal{Y} = \mathbb{R}^{pN_f}$), even if additional inputs constraints $\mathcal{U} \subset \mathbb{R}^{mN_f}$ are considered.

D. A remark on the numerical example

To visualize the implicit predictors, we deliberately chose a very low dimensional example with $n = p = m = N_f = Q = \mathcal{R} = 1$, $\ell = 3$ and treated it in a state-space setting (see Rem. 1). Furthermore, the constraint sets involved in Fig. 2 are given by $\mathcal{U} = \mathcal{Y} = [-1, 1]$. The data generating system is LTI with parameters $(A, B, C, D) = (2, -0.5, 1, 0)$ and zero-mean Gaussian measurement noise with variance $\sigma^2 = 0.01$. However, we would like to emphasize that the structure of the implicit predictor does not depend on the data generating system class but only the data itself and the OCP parameters.

IV. CONCLUSIONS AND OUTLOOK

By introducing the notion of implicit predictors, we related the input-(state)-output prediction behavior of regularized DPC schemes to more traditional predictive control schemes such as MPC and SPC, where the predictor is explicitly enforced as an equality constraint. The structure of these implicit predictors seems very relevant to us since they can

be interpreted as the behavior that the data-driven predictions attribute to the system based on given data and independently of its actual (unknown) behavior. To demonstrate this concept, we derived implicit predictors for a basic regularized DPC scheme and analyzed their structure for two popular choices of (squared) 2-norm regularization and its dependence on constraints.

In future work, we will continue this structural analysis with general ranks beyond Assumption 1 and further common DPC modifications such as slack variables [1], [6], (terminal) equality constraints [6], and different regularization choices such as the 1-norm [1], [3] or general p -norms. Although currently limited to providing a novel perspective on current DPC practices, we believe that these analyses have the potential to establish a foundation for new schemes that promote advantageous properties of the implicit predictor, and for new theoretical insights, including DPC stability proofs such as in [6], but based on the implicit predictor.

REFERENCES

- [1] J. Coulson, J. Lygeros, and F. Dörfler. Data-enabled predictive control: In the shallows of the DeePC. *18th European Control Conference*, pp. 307–312, 2019.
- [2] J. Berberich and F. Allgöwer. A trajectory-based framework for data-driven system analysis and control. *2020 European Control Conference*, pp. 1365–1370, 2020.
- [3] F. Dörfler, J. Coulson, and I. Markovsky. Bridging direct and indirect data-driven control formulations via regularizations and relaxations. *IEEE Transactions on Automatic Control*, 68(2):883–897, 2023.
- [4] M. Alsalti, J. Berberich, V. G. Lopez, F. Allgöwer, and M. A. Müller. Data-based system analysis and control of flat nonlinear systems. *2021 60th IEEE Conference on Decision and Control*, pp. 1484–1489, 2021.
- [5] J. Coulson, J. Lygeros, and F. Dörfler. Regularized and distributionally robust data-enabled predictive control. *2019 IEEE 58th Conference on Decision and Control*, pp. 2696–2701, 2019.
- [6] J. Berberich, J. Köhler, M. Müller, and F. Allgöwer. Data-driven model predictive control with stability and robustness guarantees. *IEEE Transactions on Automatic Control*, 66(4):1702–1717, 2021.
- [7] W. Favoreel, B. De Moor, and M. Gevers. Spc: Subspace predictive control. *IFAC Proceedings Volumes*, 32(2):4004–4009, 1999. 14th IFAC World Congress 1999, Beijing, China.
- [8] V. Breschi, A. Chiuso, and S. Formentin. Data-driven predictive control in a stochastic setting: a unified framework. *Automatica*, 152, 2023.
- [9] I. Markovsky and F. Dörfler. Identifiability in the behavioral setting. *IEEE Transactions on Automatic Control*, 68(3):1667–1677, 2023.
- [10] J. C. Willems, P. Rapisarda, I. Markovsky, and B. De Moor. A note on persistency of excitation. *Systems & Control Letters*, 54(4):325–329, 2005.
- [11] I. Markovsky and P. Rapisarda. Data-driven simulation and control. *International Journal of Control*, 81:1946 – 1959, 2008.
- [12] C. De Persis and P. Tesi. Formulas for data-driven control: Stabilization, optimality, and robustness. *IEEE Transactions on Automatic Control*, 65:909–924, 2020.
- [13] P. van Overschee and B. de Moor. *Subspace identification for linear systems: Theory, implementation, applications*. Kluwer Academic Publishers, 1996.
- [14] F. Fiedler and S. Lucia. On the relationship between data-enabled predictive control and subspace predictive control. *2021 European Control Conference*, pp. 222–229, 2021.
- [15] I. Markovsky and F. Dörfler. Behavioral systems theory in data-driven analysis, signal processing, and control. *Annual Reviews in Control*, 52:42–64, 2021.
- [16] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [17] F. Borrelli, A. Bemporad, and M. Morari. *Predictive Control for Linear and Hybrid Systems*. Cambridge University Press, 2017.