

Methods

Manuel Klädtke* and Moritz Schulze Darup

Towards explainable data-driven predictive control with regularizations

Ein Beitrag zur Erklärbarkeit datengetriebener prädiktiver Regelung mit Regularisierungen

<https://doi.org/10.1515/auto-2024-0161>

Received November 15, 2024; accepted January 17, 2025

Abstract: Data-driven predictive control (DPC), using linear combinations of recorded trajectory data, has recently emerged as a popular alternative to traditional model predictive control (MPC). Without an explicitly enforced prediction model, the effects of commonly used regularization terms – and the resulting predictions – can be opaque. This opacity may lead to practical challenges, such as reliance on empirical tuning of regularization parameters based on closed-loop performance, and potentially misleading heuristic interpretations of norm-based regularizations. However, by examining the structure of the underlying optimal control problem (OCP), more precise and insightful interpretations of regularization effects can be derived. In this paper, we demonstrate how to analyze the predictive behavior of DPC through implicit predictors and the trajectory-specific effects of quadratic regularization. We further extend these results to cover typical DPC modifications, including DPC for affine systems, offset regularizations, slack variables, and terminal constraints. Additionally, we provide a simple but general result on (recursive) feasibility in DPC. This work aims to enhance the explainability and reliability of DPC by providing a deeper understanding of these regularization mechanisms.

Keywords: data-driven control; predictive control; parametric optimization; explainability

Zusammenfassung: Datengetriebene prädiktive Regelung (DPC), welche im Gegensatz zur klassischen Modellprädiktiven Regelung (MPC) prädizierte Trajektorien durch Linearkombinationen aufgezeichneter Datentrajektorien generiert, hat sich als populäre Alternative etabliert. Durch das Fehlen eines expliziten Prädiktionsmodells sind die Effekte der verwendeten Regularisierungen und deren Auswirkungen auf das Prädiktionsverhalten häufig schwer durchschaubar. Dies kann zu praktischen Herausforderungen wie empirisches Auslegen von Regularisierungsparametern und möglicherweise irreführenden Interpretationen von Normbasierten Regularisierungen führen. In dieser Arbeit analysieren wir die Struktur des DPC unterliegenden Optimalsteuerungsproblems (OCP) und verschaffen so Einsichten in das implizite Prädiktionsverhalten und die Trajektorien-spezifischen Effekte von quadratischer Regularisierung. Diese Resultate erweitern wir auf affine DPC Ansätze, Offset-Regularisierungen, Schlupfvariablen und Terminalbeschränkungen. Dieser Beitrag zielt darauf ab, die Erklärbarkeit von DPC durch ein tieferes Verständnis der Regularisierungsmechanismen zu verbessern.

Schlagwörter: Datengetriebene Regelung; prädiktive Regelung; parametrische Optimierung; Erklärbarkeit

1 Introduction

Data-driven predictive control (DPC, [1]–[3]), utilizing linear combinations of recorded trajectory data to make predictions, rather than relying on an explicit system model, has recently emerged as a popular alternative to classical model predictive control (MPC, [4]). This control paradigm exemplifies a direct data-driven control scheme, contrasting with indirect (model-based) methods, as visualized in Figure 1. Originally, DPC is theoretically grounded in a result from behavioral system theory [5], known as the Willems' Fundamental Lemma. Furthermore, DPC yields an exact system

*Corresponding author: Manuel Klädtke, Department of Mechanical Engineering, Control and Cyberphysical Systems Group, TU Dortmund University, Leonhard-Euler-Straße 2, 44227 Dortmund, Germany, E-mail: manuel.klaedtke@tu-dortmund.de

Moritz Schulze Darup, Department of Mechanical Engineering, Control and Cyberphysical Systems Group, TU Dortmund University, Leonhard-Euler-Straße 2, 44227 Dortmund, Germany, E-mail: moritz.schulzedarup@tu-dortmund.de




Figure 1: Direct data-driven control schemes aim to design control directly from data. This is in contrast to the indirect (i.e., model-based) data-driven control design paradigm. Here, we aim for an indirect viewpoint (highlighted in blue) on the predictions made by direct schemes via implicit predictors.

representation and equivalence to MPC if the underlying system is deterministic LTI [2], [6]. While this exactness extends to certain classes of deterministic nonlinear systems [6], [7], it is typically lost in realistic cases involving noise, disturbances, and general nonlinearities. To address this, regularization terms are often added to the DPC objective function [2], initially motivated by their relation to distributional robustness [8]. Since then, regularized DPC schemes have shown promising performance, and further theoretical justification via robustness and stability results [1], [9]–[11]. However, the heuristic nature of regularizations often obscures their direct impact on the synthesis of predicted system trajectories from data, and since DPC operates without an explicitly enforced prediction model, the precise influence of regularizations on the resulting predictions can be challenging to discern. Therefore, our aim is to provide a deeper analysis that clarifies the interaction of control objective, constraints, and regularizations in DPC. Specifically, we propose two tools for this analysis, namely the trajectory-specific effect of regularizations (see Definition 1) and implicit predictors (see Definition 2), first introduced in [12]. The former reformulates regularization costs, translating their effect from auxiliary variables to the actual predicted system variables. The latter yields a model-based perspective on DPC by generating prediction mappings that align with DPC’s actual predictions (see Figure 1). We demonstrate the use of these analysis tools by summarizing previous results from [12]–[14] and extend them towards common modifications in DPC. The paper is organized as follows. First, in Section 2 we summarize fundamentals on direct data-driven predictions and regularized DPC, and discuss a numerical example, which is used for visualization throughout the paper. In Section 3, we further motivate our approach and explain its underlying assumptions. Section 4 presents the trajectory-specific effect of regularization and

summarizes related findings from [12]–[14]. In Section 5, we introduce implicit predictors and demonstrate their use by summarizing results from [12]–[14]. Section 6 expands these results to common DPC modifications, including affine systems, offset regularization, slack variables, and (terminal) equality constraints, and provides a general feasibility result. Finally, we conclude our work in Section 7 and preview future applications for our proposed analysis tools.

2 Preliminaries and running example

2.1 Fundamentals of direct data-driven predictions

Instead of utilizing a discrete-time state-space model with input $\mathbf{u} \in \mathbb{R}^m$, state $\mathbf{x} \in \mathbb{R}^n$, and output $\mathbf{y} \in \mathbb{R}^p$ as in traditional MPC, predictions in DPC are realized based on previously collected trajectory data $(\mathbf{u}^{(1)}, \mathbf{y}^{(1)}), \dots, (\mathbf{u}^{(\ell)}, \mathbf{y}^{(\ell)}) \in \mathbb{R}^{mL} \times \mathbb{R}^{pL}$ via linear combinations

$$\begin{pmatrix} \mathbf{u}_{\text{gen}} \\ \mathbf{y}_{\text{gen}} \end{pmatrix} = \begin{pmatrix} \mathbf{u}^{(1)} \\ \mathbf{y}^{(1)} \end{pmatrix} a_1 + \dots + \begin{pmatrix} \mathbf{u}^{(\ell)} \\ \mathbf{y}^{(\ell)} \end{pmatrix} a_\ell = \mathbf{D} \mathbf{a}. \quad (1)$$

Here, the dimensions of the data matrix $\mathbf{D} \in \mathbb{R}^{L(m+p) \times \ell}$ and generator vector $\mathbf{a} \in \mathbb{R}^\ell$ are specified by the length L of recorded (and generated) trajectories and the number ℓ of data trajectories used for predictions. Furthermore, we refer to a pair $(\mathbf{u}, \mathbf{y})$ as an input/output-sequence (*I/O-sequence*), and use roman notation to distinguish individual variables like $\mathbf{u}, \mathbf{y}$ from their associated sequences $\mathbf{u}, \mathbf{y}$. Assuming exact data generated by an LTI system, which can be described, e.g., by the state-space representation

$$\mathbf{x}(k+1) = \mathbf{A}\mathbf{x}(k) + \mathbf{B}\mathbf{u}(k) \quad (2a)$$

$$\mathbf{y}(k) = \mathbf{C}\mathbf{x}(k) + \mathbf{D}\mathbf{u}(k), \quad (2b)$$

and assuming L is greater than the lag ι of the system, the image $\mathcal{R}(\mathbf{D})$ is equivalent to the set of all possible system trajectories (of length L) if and only if [15]

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

Note that the lag ι is an integer invariant [5] of the system, i.e., invariant with respect to the considered representation. In context of the state-space representation (2), it is also known as its observability index, which is defined as the smallest integer ι , for which $\text{rank}(\mathcal{O}_\iota) = \text{rank}(\mathcal{O}_{\iota+1})$ with

$$\mathcal{O}_i := \begin{pmatrix} \mathbf{C} \\ \mathbf{CA} \\ \vdots \\ \mathbf{CA}^{i-1} \end{pmatrix}$$

is satisfied. The generalized persistency of excitation condition (3) provides the theoretical foundation for the linear combinations (1), since

$$\begin{pmatrix} \mathbf{u}_{\text{gen}} \\ \mathbf{y}_{\text{gen}} \end{pmatrix} \in \mathcal{R}(\mathbf{D}) \iff \exists \mathbf{a} \text{ such that } \begin{pmatrix} \mathbf{u}_{\text{gen}} \\ \mathbf{y}_{\text{gen}} \end{pmatrix} = \mathbf{D}\mathbf{a}. \quad (4)$$

Representing system trajectories in this way is also known as an image representation of the system, as opposed to, e.g., a state-space representation (2). A popular sufficient condition for data to satisfy (3) is known as Willem's Fundamental Lemma [5], which has become synonymous with using image representations. We note that, contrary to the Fundamental Lemma, the generalized persistency of excitation condition (3) is both sufficient and necessary, and neither requires controllability nor a Hankel structure for the data matrix. To include the current initial condition of the system as a starting point for predicted trajectories, the generated I/O-sequence is typically partitioned into a past section $(\mathbf{u}_p, \mathbf{y}_p)$ and a future section $(\mathbf{u}_f, \mathbf{y}_f)$ with N_p respectively N_f time-steps yielding

$$\begin{pmatrix} \mathbf{u}_p \\ \mathbf{u}_f \end{pmatrix} = \mathbf{u}_{\text{gen}} = \begin{pmatrix} \mathbf{U}_p \\ \mathbf{U}_f \end{pmatrix} \mathbf{a} \quad \text{and} \quad \begin{pmatrix} \mathbf{y}_p \\ \mathbf{y}_f \end{pmatrix} = \mathbf{y}_{\text{gen}} = \begin{pmatrix} \mathbf{Y}_p \\ \mathbf{Y}_f \end{pmatrix} \mathbf{a}.$$

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 constraints

$$\xi = \begin{pmatrix} \mathbf{u}_p \\ \mathbf{y}_p \end{pmatrix} = \begin{pmatrix} \mathbf{U}_p \\ \mathbf{Y}_p \end{pmatrix} \mathbf{a} = \mathbf{W}\mathbf{a}$$

force any predicted trajectory to start with the most recently witnessed behavior of the system. Note that ξ is a (non-minimal) state of the LTI system (2), if N_p is chosen greater or equal to its lag ι . From now on, we omit the “future” subscript from $\mathbf{u}_f, \mathbf{y}_f, \mathbf{U}_f$, and $\mathbf{Y}_f$, since their “past” counterparts are already incorporated in $\mathbf{W}$ and ξ , eliminating any risk of confusion. For more concise notation, we also define

$$\mathbf{Z} := \begin{pmatrix} \mathbf{W} \\ \mathbf{U} \end{pmatrix}, \quad \mathbf{z} := \begin{pmatrix} \xi \\ \mathbf{u} \end{pmatrix}, \quad \text{and} \quad \mathbf{w} := \begin{pmatrix} \mathbf{z} \\ \mathbf{y} \end{pmatrix}.$$

Additionally, with a slight abuse of notation, we redefine $\mathbf{D} := (\mathbf{Z}^\top \quad \mathbf{Y}^\top)^\top$. The latter is simply a block-row

permuted version of the original data matrix introduced in (1), used to facilitate the corresponding partitioning of $\mathbf{u}_{\text{gen}}, \mathbf{y}_{\text{gen}}$ into $\xi, \mathbf{u}, \mathbf{y}$. To summarize these new notations, $\mathbf{W}$ contains data columns for the state ξ , $\mathbf{U}$ contains data columns for the future input sequence $\mathbf{u}$, and $\mathbf{Y}$ contains data columns for the future output sequence $\mathbf{y}$. Furthermore, $\mathbf{Z}$ contains data for the state-input sequence $\mathbf{z}$, and $\mathbf{D}$ contains data for the state-input-output sequence $\mathbf{w}$.

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 [16]. To this end, consider

$$\begin{aligned} \xi &\triangleq \mathbf{x}_0 \in \mathbb{R}^n, & \mathbf{W} &\triangleq \mathbf{X}_0 \in \mathbb{R}^{n \times \ell}, \\ \mathbf{y} &\triangleq \mathbf{x} \in \mathbb{R}^{nN_f}, & \text{and} & \quad \mathbf{Y} &\triangleq \mathbf{X} \in \mathbb{R}^{nN_f \times \ell} \end{aligned}$$

leading to a data matrix $\mathbf{D}$ of dimension $n + (m + n)N_f \times \ell$. Our considerations and condition (3) with $\text{rank}(\mathbf{D}) = n + mN_f$ hold for both settings, and we employ the latter for visualization of a low dimensional numerical example used throughout the paper and detailed in Section 2.3.

Notably, in the ideal deterministic LTI setting with condition (3), the data matrix always has the rank deficiency

$$\text{rank}(\mathbf{D}) = \text{rank}(\mathbf{Z}), \quad (5)$$

such that the image representation (4) implies a unique (and exact) linear predictor mapping $\mathbf{y} = \hat{\mathbf{y}}(\xi, \mathbf{u})$. However, this rank deficiency, and with it the unique predictions, are typically lost in the presence of noise or nonlinearities, which is visualized in Figure 2. One remedy for this might be the use of (structured) low-rank approximations [17] for $\mathbf{D}$, which sometimes even implicitly occurs in DPC schemes, e.g., in γ -DDPC [18] with $\gamma_3 = \mathbf{0}$ (see [13]). Specific to the context of DPC, another remedy is given by the addition of a regularization term $h(\mathbf{a})$ to the objective function (6a) [1, 2], which is in the focus of this work.

2.2 Regularized DPC

The optimal control problem (OCP) that is solved for DPC in every time step can be stated as

$$\min_{\mathbf{u}, \mathbf{y}, \mathbf{a}} J(\xi, \mathbf{u}, \mathbf{y}) + h(\mathbf{a}) \quad (6a)$$

$$\text{s.t.} \quad \begin{pmatrix} \xi \\ \mathbf{u} \\ \mathbf{y} \end{pmatrix} = \begin{pmatrix} \mathbf{W} \\ \mathbf{U} \\ \mathbf{Y} \end{pmatrix} \mathbf{a}, \quad (6b)$$

$$(\mathbf{u}, \mathbf{y}) \in \mathcal{U} \times \mathcal{Y} \quad (6c)$$




Figure 2: Visualization of the data matrix $\mathbf{D}$ for the low-dimensional state-space example in Section 2.3. The green marks show data columns $(\mathbf{x}_0^{(i)}, \mathbf{u}^{(i)}, \mathbf{x}^{(i)})$ for (a) the ideal deterministic LTI setting and (b) the same data with added measurement noise. The gray subspaces depict $\mathcal{R}(\mathbf{D})$, showing $\text{rank}(\mathbf{D}) = \text{rank}(\mathbf{Z}) = 2$ in (a) and full rank $\mathbf{D}$ with $\text{rank}(\mathbf{D}) = 3 > \text{rank}(\mathbf{Z}) = 2$ in (b).

with control objective $J(\xi, \mathbf{u}, \mathbf{y})$, regularization $h(\mathbf{a})$, and input-output constraints $\mathcal{U} \times \mathcal{Y}$. Conditions for equivalence of DPC and MPC are well established for some special cases. In particular, as mentioned in the introduction, the equivalence with MPC for the LTI system (2) holds with exact data satisfying (3) and setting the regularization to $h(\mathbf{a}) = 0$ [2]. In the presence of noise or nonlinearities, the unregularized OCP may use the (unrealistic) additional degrees of freedom, available to it by (5) not being satisfied, to greedily minimize the objective function $J(\xi, \mathbf{u}, \mathbf{y})$. Given enough data (columns), it is established and realistic (see, e.g. [12], [18], [19], for a discussion) to make the following assumption.

Assumption 1. The data matrix $\mathbf{D}$ has full row-rank.

Note that full row-rank of $\mathbf{D}$ renders (6b) meaningless without regularization, since there is an $\mathbf{a}$ solving (6b) for any arbitrary left-hand side trajectory.

2.3 Running numerical example

We emphasize that this work does not propose a new DPC scheme; rather, it provides tools for analyzing the structure of the OCPs in existing schemes. Therefore, and in light of the following discussion in Section 3 regarding point (i), we do not consider extensive closed-loop simulations to be the best demonstration of our results. Such simulations can be found in the works we reference when analyzing the respective schemes. Instead, we use a low-dimensional system as a running example to visualize the structure of the data matrix $\mathbf{D}$ itself in Figure 2, and the structure of implicit predictors in Figures 3 and 4 (see page 373, page 378). Since using a past I/O sequence as a non-minimal state ξ yields a

minimum dimension of $(p + m)N_p \geq 2$, we instead treat the example system in a state-space setting (see Remark 1). The data-generating system is LTI with dimensions $n = m = 1$ and parameters $(\mathbf{A}, \mathbf{B}) = (2, -0.5)$. To generate Figure 2(a), we drew 500 i.i.d. samples $\mathbf{x}_0^{(i)}, \mathbf{u}^{(i)} \sim \mathcal{U}_{[-1,1]}$ and computed the subsequent state $\mathbf{x}^{(i)}$ via (2). For Figure 2(b), we added i.i.d. measurement noise $\mathcal{N}(0, 0.01)$ to each $\mathbf{x}_0^{(i)}, \mathbf{x}^{(i)}$. To visualize the implicit predictors and optimal parametric solutions in Figure 3, we chose $N_f = 1$ and set up the control objective (18) with reference $\mathbf{x}_{\text{ref}} = 0$, weight $\mathbf{Q} = 1$, and quadratic input costs $J_{\mathbf{u}}(\xi, \mathbf{u}) = \mathbf{u}^2$. The data matrix $\mathbf{D}$ is a block Hankel matrix (see, e.g., [5]), constructed from the input-state sequence

$$\begin{aligned} \mathbf{x}_{d,1} &= \begin{pmatrix} -0.1941 & 0.0048 & -0.2145 & -0.5427 & -1.3683 \end{pmatrix} \\ \mathbf{u}_{d,1} &= \begin{pmatrix} -0.7859 & 0.4483 & 0.2274 & 0.5659 \end{pmatrix}, \end{aligned}$$

which was generated by drawing an initial state $\mathbf{x}_0 \sim \mathcal{U}_{[-1,1]}$, applying the i.i.d. input sequence $\mathbf{u}_d^{(i)} \sim \mathcal{U}_{[-1,1]}$ via (2), and then adding i.i.d. measurement noise $\mathcal{N}(0, 0.01)$ to the resulting state sequence. To visualize the implicit predictors in Figure 4, we set $N_f = 2$, with reference $\mathbf{x}_{\text{ref}} = \begin{pmatrix} 0.5 & 0.5 \end{pmatrix}^\top$, weight $\mathbf{Q} = \mathbf{I}_2$, and arbitrary input costs, using a regularization weight of $\lambda = 1$. The data matrix $\mathbf{D}$ is again a block Hankel matrix, constructed from the input-state sequence

$$\begin{aligned} \mathbf{x}_{d,2} &= \begin{pmatrix} \mathbf{x}_{d,1} & -3.1384 & -5.9759 \end{pmatrix} \\ \mathbf{u}_{d,2} &= \begin{pmatrix} \mathbf{u}_{d,1} & 0.8037 & -0.6017 \end{pmatrix}, \end{aligned}$$

where the two additional samples are generated in the same way as before. These extra samples are needed to allow that the data matrix $\mathbf{D}$ satisfies Assumption 1 for the increased prediction horizon $N_f = 2$.

3 Motivation and setting

Looking at recent DPC literature, we have observed the following two trends:

- (i) Regularization weights are typically tuned empirically, by evaluating closed-loop performance [18], [20], [21].
- (ii) Heuristic interpretations of norm-based regularizations, e.g., $h(\mathbf{a}) = \lambda \|\mathbf{a}\|_2^2$ or $h(\mathbf{a}) = \lambda \|\mathbf{a}\|_1$, are (mostly unchanged) transferred to DPC [1], [20].

We acknowledge that both observations are grounded in intuitive and practically valid heuristics. However, (i) only offers a quantitative indicator of control performance being

improved/diminished by regularizations. It may obscure the underlying qualitative interactions between constraints, control objective, and regularization that led to this effect. Some of the already established theoretical investigations in these qualitative effects include the interpretation of DPC with regularization as a convex relaxation of other (indirect) schemes [20] or in the context of distributional robustness [8], and our work aims to add on to these.

Regarding (ii), such interpretations may be unspecific in the context of predictive control and might even be misleading. For example, it is claimed in [22] for regularization via $h(\mathbf{a}) = \lambda \|\mathbf{a}\|_2^2$ that too large λ deteriorate tracking performance, since “too small choices of the vector $\mathbf{a}$ shift the input and output to which the closed-loop converges towards zero” [22, Section 4]. While we agree that tracking performance is deteriorated, the inputs and outputs are generally not shifted towards zero by small $\mathbf{a}$ for $\xi \neq \mathbf{0}$, since the chosen optimal $\mathbf{a}^*$ must also satisfy $\mathbf{W}\mathbf{a}^* = \xi$. Instead, they are shifted towards a (multistep) predictor $\hat{\mathbf{y}}(\xi, \mathbf{u})$ and (multistep) controller $\hat{\mathbf{u}}(\xi)$ “favored” by the regularization $h(\mathbf{a})$. Both can be precisely characterized by our proposed tools, as demonstrated in Section 4.1 for linear DPC and Section 6.1 for the affine DPC setting considered in [22].

The motivation of our work is to provide a structured analysis that explains the interactions between constraints, control objectives, and regularization in DPC. Although the current results focus on analyzing existing schemes rather than proposing new ones, they offer deeper insights into these interactions and reveal potential pitfalls. This creates a foundation for future improvements in DPC schemes. Furthermore, our analysis is intentionally agnostic to the specific class of the data-generating system. Rather than focusing on a particular system type, our setting is tailored to the structure of data matrices generated by them. This often reduces to Assumption 1, which is typically (almost surely) satisfied for both non-deterministic LTI, and non-deterministic nonlinear systems. On the contrary, in practical cases, Assumption 1 typically only fails if too little data are used (i.e., $\mathbf{D}$ is tall) or if the input data lack sufficient excitation. This generality comes at a price. Without such assumptions on the true system, one cannot expect to evaluate how closely the analyzed predictions match its true behavior or how the optimal controls generated by DPC interact in closed-loop. Nonetheless, such assumptions can be readily added to tailor these results to a specific system class, incorporating prior system knowledge. Moreover, our findings can already help explain certain phenomena observed when applying DPC in closed-loop, such as those seen in [20, Figure 2], which we discuss in Section 4.2.

Our analysis relies on two key conceptual tools. One is given by the novel concept of *implicit predictors* (depicted in Figure 1 and specified in Definition 2 further below), which aims to describe the predictive behavior of OCPs that may not have an explicitly enforced prediction model. The introduction of implicit predictors in [12] has also led to a parametric characterization of regularization costs in terms of their *trajectory-specific effects* (formalized in Definition 1). Since then, similar analyses have appeared in [13], [14], [19], [23], offering useful generalizations and intuitive interpretations of the underlying structures. Due to its independent value beyond implicit predictors, we treat this as a separate, second tool.

We first introduce the trajectory-specific effect of regularizations in Section 4, since it facilitates the following introduction of implicit predictors in Section 5. As an exemplary demonstration of these tools, we summarize results from [12]–[14] for their use in DPC with quadratic regularization in Sections 4.1, 4.2, and 5.1. We then extend these results in Section 6 by analyzing the effects of common modifications to the DPC problem. These analyses cover the extension to affine DPC in Section 6.1, the inclusion of an offset in the regularization in Section 6.2, the inclusion of slack variables in Section 6.3, and the inclusion of additional (terminal) equality constraints in Section 6.4. While not technically a modification to DPC, we also give two short but very general results on (recursive) feasibility in DPC in Section 6.5.

4 Trajectory-specific effect of regularization

In the ideal deterministic LTI setting without regularization, the variable $\mathbf{a}$ is only used in (6b) as an expression for the image representation introduced in (4), i.e.,

$$\begin{pmatrix} \xi \\ \mathbf{u} \\ \mathbf{y} \end{pmatrix} \in \mathcal{R}(\mathbf{D}) \iff \exists \mathbf{a} \text{ such that } \begin{pmatrix} \xi \\ \mathbf{u} \\ \mathbf{y} \end{pmatrix} = \mathbf{D}\mathbf{a}. \quad (7)$$

However, adding a regularization $h(\mathbf{a})$ introduces another meaning to $\mathbf{a}$, which is not based on the image representation and the underlying behavioral system theory but on heuristics. Intuitively, $h(\mathbf{a})$ adds a price tag to every $\mathbf{a}$, which is also transferred to the trajectory tuple $(\xi, \mathbf{u}, \mathbf{y})$ generated by $\mathbf{D}\mathbf{a}$. Note that, if $\mathbf{D}$ is wide, there might be multiple $\mathbf{a}$ generating the same trajectory and thus multiple price tags associated with it. However, since we are optimizing over $\mathbf{a}$ in (6), it is natural to only consider the

lowest cost associated with any trajectory. This is formally captured by the following definition.

Definition 1. We call the solution $h^*(\xi, \mathbf{u}, \mathbf{y})$ to the optimization problem

$$h^*(\xi, \mathbf{u}, \mathbf{y}) := h(\mathbf{a}^*(\xi, \mathbf{u}, \mathbf{y})) = \min_{\mathbf{a}} h(\mathbf{a}) \quad (8a)$$

$$\text{s.t.} \quad \begin{pmatrix} \xi \\ \mathbf{u} \\ \mathbf{y} \end{pmatrix} = \begin{pmatrix} \mathbf{W} \\ \mathbf{U} \\ \mathbf{Y} \end{pmatrix} \mathbf{a} \quad (8b)$$

the *trajectory-specific effect* of the regularization $h(\mathbf{a})$ given the data $\mathcal{D}$.

Importantly, note that $(\xi, \mathbf{u}, \mathbf{y})$ appear as parameters in (1) and not as optimization variables. Therefore, additional constraints such as (6c) are irrelevant. That is, $h^*(\xi, \mathbf{u}, \mathbf{y})$ is valid for all $(\xi, \mathbf{u}, \mathbf{y})$ satisfying (7), and therefore also for those, which additionally need to satisfy (6c). The relevance of $h^*(\xi, \mathbf{u}, \mathbf{y})$ to the DPC problem comes from the fact that (1) naturally appears as an inner optimization problem in (6). That is, (6) is equivalent to

$$\min_{\mathbf{u}, \mathbf{y}} J(\xi, \mathbf{u}, \mathbf{y}) + h^*(\xi, \mathbf{u}, \mathbf{y}) \quad (9a)$$

$$\text{s.t.} \quad \begin{pmatrix} \xi \\ \mathbf{u} \\ \mathbf{y} \end{pmatrix} \in \mathcal{R}(\mathcal{D}), \quad (9b)$$

$$(\mathbf{u}, \mathbf{y}) \in \mathcal{U} \times \mathcal{Y}. \quad (9c)$$

Note that we have deliberately replaced (6b) via (7) to highlight the fact that $\mathbf{a}$ can be fully eliminated, since it is just an auxiliary variable, after all. That is, the image representation (7) acts just as before, but the additional heuristic costs introduced by $h(\mathbf{a})$ (with which we started this section) are now fully explained by their trajectory-specific effect $h^*(\xi, \mathbf{u}, \mathbf{y})$. This allows for much more intuitive interpretations, which we demonstrate by summarizing results from [12]–[14] on the trajectory-specific effect of quadratic regularization $h(\mathbf{a}) = \lambda \|\mathbf{a}\|_2^2$ in Section 4.1, and its projection-based variant $h(\mathbf{a}) = \lambda \|\Pi_{\perp} \mathbf{a}\|_2^2$ in Section 4.2, where we also discuss its link to γ -DDPC [18]. One may notice that, due to its popularity and its conformity for analytic solutions, all results in this paper focus on (projection-based) quadratic regularization. However, we want to emphasize the generality of Definition 1, which is suited for *any* choice of regularization $h(\mathbf{a})$. In particular, similar analyses of trajectory-specific effects for 1-norm $h(\mathbf{a}) = \lambda \|\mathbf{a}\|_1$ [2] or general p-norm $h(\mathbf{a}) = \lambda \|\mathbf{a}\|_p$ [20] regularizations should be equally insightful, but are still an open problem.

4.1 Trajectory-specific effect of quadratic regularization

It was first shown in [12] with additional details in [13, Prop. 1] that, under Assumption 1, the trajectory-specific effect of quadratic regularization $h(\mathbf{a}) = \lambda \|\mathbf{a}\|_2^2$ is given by

$$\lambda \|\mathbf{a}^*(\xi, \mathbf{u}, \mathbf{y})\|_2^2 = \lambda \|\mathbf{y} - \hat{\mathbf{y}}_{\text{LS}}(\xi, \mathbf{u})\|_{\mathcal{Q}_{\text{reg}}}^2 \quad (10a)$$

$$+ \lambda \|\mathbf{u} - \hat{\mathbf{u}}_{\text{LS}}(\xi)\|_{\mathcal{R}_{\text{reg}}}^2 \quad (10b)$$

$$+ \lambda \|\xi\|_{(\mathbf{W}\mathbf{W}^T)^{-1}}^2. \quad (10c)$$

Here, $\hat{\mathbf{y}}_{\text{LS}}(\xi, \mathbf{u})$ and $\hat{\mathbf{u}}_{\text{LS}}(\xi)$ are the (multistep) predictor/controller “favored” by the regularization, as initially mentioned in Section 3. The regularization pushes predicted $\mathbf{y}$ and $\mathbf{u}$ towards them, which is visualized in Figure 3 for the low-dimensional state-space example discussed in Section 2.3. By parametrically solving (1) and decomposing block-matrix expressions (see [12] for details), one can show that both are given by linear mappings, which can be represented as

$$\hat{\mathbf{y}}_{\text{LS}}(\xi, \mathbf{u}) := \mathbf{G}_{\text{LS}} \mathbf{z} \quad (11)$$

$$\hat{\mathbf{u}}_{\text{LS}}(\xi) := \mathbf{K}_{\text{LS}} \xi \quad (12)$$

via the least-squares solutions

$$\mathbf{G}_{\text{LS}} := \arg \min_{\mathbf{G}} \|\mathbf{Y} - \mathbf{G}\mathbf{Z}\|_{\text{F}}^2 = \mathbf{Y}\mathbf{Z}^+ \quad (13a)$$

$$\mathbf{K}_{\text{LS}} := \arg \min_{\mathbf{K}} \|\mathbf{U} - \mathbf{K}\mathbf{W}\|_{\text{F}}^2 = \mathbf{U}\mathbf{W}^+, \quad (13b)$$

where $\|\cdot\|_{\text{F}}$ denotes the Frobenius norm. Note that $\hat{\mathbf{y}}_{\text{LS}}(\xi, \mathbf{u})$ is equivalent to the subspace predictor used as an equality constraint in subspace predictive control (SPC, [24]). The involved weighing matrices are given by

$$\mathcal{Q}_{\text{reg}} := (\mathbf{E}_{\mathbf{y}} \mathbf{E}_{\mathbf{y}}^T)^{-1} \quad \text{and} \quad \mathcal{R}_{\text{reg}} := (\mathbf{E}_{\mathbf{u}} \mathbf{E}_{\mathbf{u}}^T)^{-1}, \quad (14)$$

where

$$\mathbf{E}_{\mathbf{y}} := \mathbf{Y} - \mathbf{G}_{\text{LS}} \mathbf{Z} \quad \text{and} \quad \mathbf{E}_{\mathbf{u}} := \mathbf{U} - \mathbf{K}_{\text{LS}} \mathbf{W}$$

are the residual matrices associated with the least-square problems (13), and the inverses exist under Assumption 1. Therefore, $\mathcal{Q}_{\text{reg}}$ and $\mathcal{R}_{\text{reg}}$ can be interpreted as inverses of scaled (because they are not normalized w.r.t. the amount of data columns ℓ) empirical second moment matrices for the output prediction error $\mathbf{E}_{\mathbf{y}}$ and “input prediction error” $\mathbf{E}_{\mathbf{u}}$ of the least-squares solutions $\mathbf{G}_{\text{LS}}, \mathbf{K}_{\text{LS}}$ based on the data $\mathcal{D}$. Hence, quadratic regularization $h(\mathbf{a}) = \lambda \|\mathbf{a}\|_2^2$ pushes $\mathbf{y}$ and $\mathbf{u}$ towards the least-squares estimates $\hat{\mathbf{y}}_{\text{LS}}(\xi, \mathbf{u})$ and $\hat{\mathbf{u}}_{\text{LS}}(\xi)$. It does so more (less) harshly in directions, where the

latter are believed to be more (less) accurate, based on the available data. The same second moment (or covariance)-based interpretation appears in [19], with the distinction that the state-input dependent part in [19, Eq. (21)] is not yet decomposed into the two terms (10b) and (10c), similar to [12]. However [19], also shows that the results naturally generalize to rank-deficient data matrices, i.e., $\mathbf{D}$ not satisfying Assumption 1, by simply using the Moore–Penrose inverse in (14). In that case, deviations $\mathbf{e}_y = \mathbf{y} - \hat{\mathbf{y}}_{LS}(\xi, \mathbf{u})$ are not penalized in singular directions of $\mathbf{E}_y \mathbf{E}_y^\top$, i.e., directions in which no deviation from $\hat{\mathbf{y}}_{LS}(\xi, \mathbf{u})$ has been observed in the data matrix $\mathbf{D}$. However, due to (7), one can show that $\mathbf{y}$ may not deviate at all in these directions. That is, in unexplored directions of the state-input-output space, the penalization of deviations in (10) is replaced with hard constraints, since deviations from it are not possible in $\mathcal{R}(\mathbf{D})$. The same considerations apply for deviations $\mathbf{e}_u = \mathbf{u} - \hat{\mathbf{u}}_{LS}(\xi)$ if $\mathbf{Z}$ does not have full row-rank.

When discussing the role of these cost terms, first note that the last cost term (10c) is irrelevant to the OCP, since ξ is a parameter determined in closed-loop and not an optimization variable. Regarding the usefulness of (10a), we believe that $\mathbf{y}$ being pushed towards the least-squares (multistep) predictor $\hat{\mathbf{y}}_{LS}(\xi, \mathbf{u})$ in (10a) is quite intuitive. However, while $\hat{\mathbf{y}}_{LS}(\xi, \mathbf{u})$ is conceptually not a bad choice for a predictor, it should be noted that (without further modifications such as in [25]) it is typically neither causal nor time-invariant. Furthermore, although it is conceptually nice to include additional statistical information via $\mathbf{Q}_{reg}$ and $\mathbf{R}_{reg}$, this information is utilized very greedily in DPC. That is, since the optimal solution is a trade-off between control objective and regularization cost, the DPC cost objective (6a) promotes deviations from $\hat{\mathbf{y}}_{LS}(\xi, \mathbf{u})$ primarily in uncertain directions in order to achieve lower costs for the control objective $J(\xi, \mathbf{u}, \mathbf{y})$. We believe that it remains an interesting avenue for future research to explore under what conditions penalizing predictions $\mathbf{y}$ towards a (potentially well- or ill-suited) predictor $\hat{\mathbf{y}}(\xi, \mathbf{u})$, as in (10a), outperforms enforcing it as an equality constraint $\mathbf{y} = \hat{\mathbf{y}}(\xi, \mathbf{u})$. While the intuition behind penalizing the input $\mathbf{u}$ towards $\hat{\mathbf{u}}_{LS}(\xi)$ in the same fashion may not be immediately clear, it was shown in the γ -DDPC setting that a cost term equivalent to (10b) (proven in [12]) can indeed be utilized to increase DPC performance [26]. Furthermore, a functionally similar input penalty appears in another framework focused on minimizing the final control error (FCE) [27]. Intuitively, the performance improvement can be explained by the term (10b) pushing the chosen input $\mathbf{u}$ towards the best-explored region of the state-input-space, i.e., where the most confident predictions

can be made based on the available data. For a more precise quantification of the relation between prediction error variance and $\|\gamma_2\|_2^2$, we refer to [26].

Finally, we want to highlight the discrepancy between (10a) and (10b) in terms of proper tuning for λ . The output-related cost term (10a) needs a large weight λ because, without the previously discussed rank deficiency (5), it is the sole factor keeping output predictions $\mathbf{y}$ from being greedily and unrealistically (i.e., without considering the data in $\mathbf{D}$) pushed toward the minimum of the control objective $J(\xi, \mathbf{u}, \mathbf{y})$. This effect is visualized in Figure 3(a). On the other hand, the input-related cost term (10b) may not have a large weight λ , since this would favor always sticking to the “best-explored” choice $\hat{\mathbf{u}}_{LS}(\xi)$ over choosing an input sequence $\mathbf{u}$ that benefits the control objective $J(\xi, \mathbf{u}, \mathbf{y})$, which is visualized in Figure 3(d) and (e). These considerations highlight the need for isolation of the individual cost terms. In the following section, we present two modified DPC schemes in which, as revealed through an analysis of the trajectory-specific effect of regularization, this isolation of regularization cost terms naturally occurs.

4.2 Isolating trajectory-specific effects via projections or γ -DDPC

In [20], the orthogonal projection matrices

$$\mathbf{\Pi} := \mathbf{Z}^+ \mathbf{Z} \quad \text{and} \quad \mathbf{\Pi}_\perp := \mathbf{I} - \mathbf{\Pi}$$

were introduced to yield a regularization $h(\mathbf{a}) = \lambda \|\mathbf{\Pi}_\perp \mathbf{a}\|_2^2$ for which the closed-loop behavior is consistent with the subspace predictor as $\lambda \rightarrow \infty$. While it is noted that norm-based regularizers such as $h(\mathbf{a}) = \lambda \|\mathbf{a}\|_2^2$ “are not consistent and bias the optimal solution” [20, Section IV.C], the nature of this bias is not further explored. In [12], this nature is uncovered by showing that (under Assumption 1) the trajectory-specific effect of projection-based quadratic regularization $h(\mathbf{a}) = \lambda \|\mathbf{\Pi}_\perp \mathbf{a}\|_2^2$ is given by

$$\lambda \|\mathbf{\Pi}_\perp \mathbf{a}^*(\xi, \mathbf{u}, \mathbf{y})\|_2^2 = \lambda \|\mathbf{y} - \hat{\mathbf{y}}_{LS}(\xi, \mathbf{u})\|_{\mathbf{Q}_{reg}}^2. \quad (15)$$

Since $\|\mathbf{a}\|_2^2 = \|\mathbf{\Pi} \mathbf{a}\|_2^2 + \|\mathbf{\Pi}_\perp \mathbf{a}\|_2^2$, we also immediately see that the remaining cost terms (10b) and (10c) are associated with $\mathbf{\Pi}$. That is, (under Assumption 1) the trajectory-specific effect of projection-based quadratic regularization $h(\mathbf{a}) = \lambda \|\mathbf{\Pi} \mathbf{a}\|_2^2$ is given by

$$\lambda \|\mathbf{\Pi} \mathbf{a}^*(\xi, \mathbf{u}, \mathbf{y})\|_2^2 = \lambda \|\mathbf{u} - \hat{\mathbf{u}}_{LS}(\xi)\|_{\mathbf{R}_{reg}}^2 + \lambda \|\xi\|_{(\mathbf{w} \mathbf{w}^\top)^{-1}}^2.$$

This separation explains the observations in [20, Figure 2]. There, the performance of $h(\mathbf{a}) = \lambda \|\mathbf{a}\|_2^2$ deteriorates for $\lambda \rightarrow \infty$ since the chosen input $\mathbf{u}$ is aggressively pushed towards $\hat{\mathbf{u}}_{\text{LS}}(\xi)$, while neglecting the control objective $J(\xi, \mathbf{u}, \mathbf{y})$. In contrast, the projection-based regularization $h(\mathbf{a}) = \lambda \|\Pi_{\perp} \mathbf{a}\|_2^2$ avoids this problem by fully dropping the input-related term (10b). However, the “closed-loop view” (discussed in Section 3) in [20, Figure 2] hides the fact that two functionally very different cost terms (10a) and (10b) are responsible for the observed performance. Furthermore [20], only considers the two options of either using both (10a) and (10b) with the same weight λ via $h(\mathbf{a}) = \lambda \|\mathbf{a}\|_2^2$, or only using (10a) via $h(\mathbf{a}) = \lambda \|\Pi_{\perp} \mathbf{a}\|_2^2$. However, an individually weighted mix of the two terms may provide the best of both worlds, which is coincidentally explored in the regularization schemes proposed for γ -DDPC [18].

In γ -DDPC, the constraint (6b) and variable $\mathbf{a}$ are replaced via LQ decomposition as

$$\begin{pmatrix} \mathbf{W} \\ \mathbf{U} \\ \mathbf{Y} \end{pmatrix} = \begin{pmatrix} L_{11} & 0 & 0 & 0 \\ L_{21} & L_{22} & 0 & 0 \\ L_{31} & L_{32} & L_{33} & 0 \end{pmatrix} \begin{pmatrix} Q_1 \\ Q_2 \\ Q_3 \\ Q_4 \end{pmatrix},$$

$$\boldsymbol{\gamma} := \begin{pmatrix} \gamma_1 \\ \gamma_2 \\ \gamma_3 \\ \gamma_4 \end{pmatrix} = \begin{pmatrix} Q_1 \\ Q_2 \\ Q_3 \\ Q_4 \end{pmatrix} \mathbf{a},$$

where the diagonal blocks L_{ii} for $i \in \{1, 2, 3\}$ are non-singular (under Assumption 1) and the matrices Q_i have orthonormal rows. Furthermore, Q_4 and γ_4 are typically omitted, since they do not affect the generated trajectory. The idea behind γ -DDPC can be summarized as reparameterizing the OCP with a lower dimensional variable $\boldsymbol{\gamma}$, and decoupling the matching of the initial condition ξ , since $\gamma_1 = L_{11}^{-1} \xi$ is uniquely determined [18]. The proposed regularization strategies for γ -DDPC are based on a mix of quadratic regularization $\tilde{h}(\boldsymbol{\gamma}) = \lambda_2 \|\gamma_2\|_2^2 + \lambda_3 \|\gamma_3\|_2^2$, or constraining $\gamma_3 = \mathbf{0}$ and using only $\tilde{h}(\boldsymbol{\gamma}) = \lambda_2 \|\gamma_2\|_2^2$ [23], [26]. The connection between this regularization of $\boldsymbol{\gamma}$ -variables and projection-based regularization of $\mathbf{a}$ was revealed in [13] by analyzing their trajectory-specific effect.

Corollary 1. ([13]) *Under Assumption 1, regularization of the $\boldsymbol{\gamma}$ -variables can be equivalently expressed by the trajectory-specific effect*

$$\begin{aligned} \|\gamma_1\|_2^2 &= \|\xi\|_{(ww^T)^{-1}}^2, \\ \|\gamma_2\|_2^2 &= \|\mathbf{u} - \hat{\mathbf{u}}_{\text{LS}}(\xi)\|_{\mathcal{R}_{\text{reg}}}^2, \\ \|\gamma_3\|_2^2 &= \|\mathbf{y} - \hat{\mathbf{y}}_{\text{LS}}(\xi, \mathbf{u})\|_{\mathcal{Q}_{\text{reg}}}^2. \end{aligned}$$

We note that the difference in notation w.r.t. to Definition 1 only comes from the fact that $(\gamma_1, \gamma_2, \gamma_3)$ are uniquely determined by $(\xi, \mathbf{u}, \mathbf{y})$, and therefore $\gamma_i^*(\xi, \mathbf{u}, \mathbf{y}) = \gamma_i$. From Corollary 1, we can see that both the projection-based regularization and the regularization of γ_i isolate specific terms of the trajectory-specific effect of standard quadratic regularization. Their relation can be summarized as follows

$$\begin{aligned} & \underbrace{\|\Pi_{\perp} \mathbf{a}^*(\xi, \mathbf{u}, \mathbf{y})\|_2^2}_{\|\gamma_3\|_2^2} + \underbrace{\|\Pi \mathbf{a}^*(\xi, \mathbf{u}, \mathbf{y})\|_2^2}_{\|\gamma_2\|_2^2} + \underbrace{\|\xi\|_{(ww^T)^{-1}}^2}_{\|\gamma_1\|_2^2} \\ &= \|\mathbf{a}^*(\xi, \mathbf{u}, \mathbf{y})\|_2^2. \end{aligned}$$

As a guideline for practitioners, we advocate to use a mixed regularization $h(\mathbf{a}) = \lambda_2 \|\Pi \mathbf{a}\|_2^2 + \lambda_3 \|\Pi_{\perp} \mathbf{a}\|_2^2$ (or, equivalently in the γ -DDPC framework $\tilde{h}(\boldsymbol{\gamma}) = \lambda_2 \|\gamma_2\|_2^2 + \lambda_3 \|\gamma_3\|_2^2$). This approach generalizes both the quadratic regularization $h(\mathbf{a}) = \lambda \|\mathbf{a}\|_2^2$, and the purely projection based $h(\mathbf{a}) = \lambda \|\Pi_{\perp} \mathbf{a}\|_2^2$ by choosing $\lambda = \lambda_2 = \lambda_3$ or $\lambda = \lambda_3, \lambda_2 = 0$, respectively. To avoid output predictions blindly following the control objective, λ_3 should be chosen very large, and (except for potential numerical solver instabilities) there are no adverse effects to be expected even for $\lambda_3 \rightarrow \infty$. On the contrary, λ_2 should be chosen very carefully to avoid the inputs blindly following the least-squares controller $\hat{\mathbf{u}}_{\text{LS}}(\xi)$. While $\lambda_2 = 0$ (recovering the projection-based regularization $h(\mathbf{a}) = \lambda_3 \|\Pi_{\perp} \mathbf{a}\|_2^2$) can be considered a safe choice in this regard, it also foregoes the potential performance improvements associated with proper tuning of this term. Here, we refer to [26] for more specific tuning guidelines. Finally, as also observed in [19], adding or removing trajectory columns from $\mathbf{D}$ after the tuning requires a re-scaling of all regularization weights

$$\lambda_{i,\text{new}} = \frac{\ell_{\text{new}}}{\ell_{\text{old}}} \lambda_{i,\text{old}}$$

to retain a quantitatively similar effect, since $\mathcal{Q}_{\text{reg}}, \mathcal{R}_{\text{reg}}$ are not normalized w.r.t. the number of trajectory columns.




Figure 3: Implicit predictor, optimal parametric solutions, and least-squares mappings for the DPC problem discussed in Section 2.3 (a–d). The optimal parametric DPC solutions $(\mathbf{x}_0, \mathbf{u}^*(\mathbf{x}_0), \mathbf{x}^*(\mathbf{x}_0))$ for the different regularizations $h(\mathbf{a}) = \lambda \|\mathbf{a}\|_2^2$ (orange) and $h(\mathbf{a}) = \lambda \|\mathbf{I} \mathbf{I}_\perp \mathbf{a}\|_2^2$ (green) naturally evolve on the implicit predictor $\hat{\mathbf{x}}_{\text{DPC}}(\mathbf{x}_0, \mathbf{u})$ (gray). (e) Visualization of the least-square solutions $\hat{\mathbf{x}}_{\text{LS}}(\mathbf{x}_0, \mathbf{u})$ (gray) and $\hat{\mathbf{u}}_{\text{LS}}(\mathbf{x}_0)$. The latter is shown via the tuple $(\mathbf{x}_0, \hat{\mathbf{u}}_{\text{LS}}(\mathbf{x}_0), \hat{\mathbf{x}}_{\text{LS}}(\mathbf{x}_0, \hat{\mathbf{u}}_{\text{LS}}(\mathbf{x}_0)))$ (orange).

Remark 2. Technically, the data matrix $\mathbf{D}$ is already normalized w.r.t. $\sqrt{\ell}$ in the γ -DDPC framework, such that the aforementioned re-scaling is not necessary. For clarity of presentation and direct comparison to the projection-based regularization, we have omitted this normalization. Note that all considerations remain valid with(out) normalization, since it does not change the image representation (7) and the same effect in terms of quadratic regularization can be achieved by rescaling weights λ_i with ℓ .

5 Implicit predictors in regularized DPC

Although the absence of the rank deficiency (5) allows for non-unique output predictions, and Assumption 1 even allows for *any* triple $(\xi, \mathbf{u}, \mathbf{y})$ to be generated from $\mathbf{D}$ via linear combinations, one can still observe certain patterns in the optimal solutions $(\mathbf{y}^*(\xi), \mathbf{u}^*(\xi), \xi)$ to (6). In particular, it seems that the output predictions implicitly follow a (multistep) predictor mapping $\mathbf{y}^* = \hat{\mathbf{y}}(\xi, \mathbf{u}^*(\xi))$, even though such a mapping is not explicitly enforced as an equality constraint in (6). These considerations led to the following notion of implicit predictors introduced in [12].

Definition 2. ([12]) We call $\hat{\mathbf{y}}(\xi, \mathbf{u})$ an *implicit predictor* for an OCP if including the constraint $\mathbf{y} = \hat{\mathbf{y}}(\xi, \mathbf{u})$ does not alter the (set of) minimizers $(\mathbf{u}^*, \mathbf{y}^*)$ and the optimal value.

Hence, an implicit predictor $\hat{\mathbf{y}}_{\text{DPC}}(\xi, \mathbf{u})$ for the DPC problem (6) can be interpreted as a model that generates the output predictions $\mathbf{y}$ based on $(\xi, \mathbf{u})$ consistent with the solutions to (6), even though such a model is never explicitly enforced. It is meant as a descriptive object (not prescriptive, since it is not enforced in (6)), which can be studied to

explain the predictive behavior of DPC schemes and make improvements, accordingly. The interaction of this interpretation with the direct and indirect data-driven control paradigms is visualized in Figure 1. While Definition 2 is more conceptual than constructive, a valid implicit predictor for the DPC problem (6) can be constructed by solving the optimization problem

$$\hat{\mathbf{y}}_{\text{DPC}}(\xi, \mathbf{u}) = \arg \min_{\mathbf{y}} \min_{\mathbf{a}} J(\xi, \mathbf{u}, \mathbf{y}) + h(\mathbf{a}) \quad (16a)$$

$$\text{s.t.} \quad \begin{pmatrix} \xi \\ \mathbf{u} \\ \mathbf{y} \end{pmatrix} = \begin{pmatrix} \mathbf{W} \\ \mathbf{U} \\ \mathbf{Y} \end{pmatrix} \mathbf{a}, \quad (16b)$$

$$\mathbf{y} \in \mathcal{Y}. \quad (16c)$$

In contrast to (1), we now treat $(\xi, \mathbf{u})$ as parameters and optimize over $(\mathbf{a}, \mathbf{y})$. Hence, additional set constraints $\mathbf{u} \in \mathcal{U}$ can be dropped while $\mathbf{y} \in \mathcal{Y}$ still need to be considered. However, similar to (1), note that (16) is an inner optimization problem to (6) and thus its parametric solution $\hat{\mathbf{y}}_{\text{DPC}}(\xi, \mathbf{u})$ naturally satisfies Definition 2.

In the presence of the rank deficiency (5), $\hat{\mathbf{y}}_{\text{DPC}}(\xi, \mathbf{u})$ is fully determined by the image representation (16b). Furthermore, it reveals the unique linear predictor mapping implicitly within the data matrix $\mathbf{D}$, which can be specified as $\hat{\mathbf{y}}_{\text{DPC}}(\xi, \mathbf{u}) = \hat{\mathbf{y}}_{\text{LS}}(\xi, \mathbf{u})$. Note that this is precisely the setting and results considered in [28]. That is, in the deterministic LTI case [28, Thm. 1], the rank deficiency (5) naturally occurs, as discussed in Section 2.1. Notably, in the non-deterministic setting, [28, Thm. 2], considers a data matrix $\mathbf{D}$ which has so few columns ℓ that the rank deficiency (5) is also satisfied, i.e.,

$$\ell = \text{rank}(\mathbf{Z}) \Rightarrow \text{rank}(\mathbf{D}) = \text{rank}(\mathbf{Z}). \quad (17)$$

However, in these two settings, the predictor $\hat{\mathbf{y}}_{\text{LS}}(\xi, \mathbf{u})$ somewhat loses its meaning as a least squares estimate (13a),

as it consistently yields zero residual ($\mathbf{E}_y = \mathbf{0}$). That is, in the presence of rank deficiency (5), the linear system of equations (7) can be solved algebraically for an exact and unique linear mapping $\hat{\mathbf{y}}(\xi, \mathbf{u})$, and using the least squares solution is merely one way of doing so. In other words [28], is often cited when arguing for equivalence between SPC and DPC. However, this equivalence only occurs if there is an exact linear output predictor mapping within the data $\mathcal{D}$, since both methods manage to find it. Any alternative method that successfully identifies this exact linear mapping and employs it for predictive control would also be considered equivalent.

In the absence of the rank-deficiency (5), i.e., in a realistic (non-deterministic) setting with more data columns than in (17), the additional degrees of freedom lead to output predictions deviating from $\hat{\mathbf{y}}_{\text{LS}}(\xi, \mathbf{u})$. Although it was shown in [20] that the optimal solution of DPC with projection-based quadratic regularization $h(\mathbf{a}) = \lambda \|\Pi_{\perp} \mathbf{a}\|_2^2$ is consistent with SPC for $\lambda \rightarrow \infty$ (and therefore $\hat{\mathbf{y}}_{\text{DPC}}(\xi, \mathbf{u}) = \hat{\mathbf{y}}_{\text{LS}}(\xi, \mathbf{u})$ in that case), a similar characterization of predictive behavior for $h(\mathbf{a}) = \lambda \|\mathbf{a}\|_2^2$ or finite λ is not discussed in [20]. Precisely this kind of characterization of predictive behavior for finite λ is possible using the concept of implicit predictors and results for (projection-based) quadratic regularization are given in [12], [14], which will be summarized in the following section. Although we focus on DPC with these two kinds of quadratic regularization, we want to emphasize the generality of Definition 2, which is suited for any kind of predictive control scheme.

An implicit predictor for DPC with quadratic regularization

In the following, we assume that the control objective is a quadratic output-tracking formulation

$$J(\xi, \mathbf{u}, \mathbf{y}) = \|\mathbf{y} - \mathbf{y}_{\text{ref}}\|_{\mathbf{Q}}^2 + J_{\mathbf{u}}(\xi, \mathbf{u}) \quad (18)$$

with reference $\mathbf{y}_{\text{ref}}$, positive semidefinite weighing matrix $\mathbf{Q}$, and arbitrary input control objective $J_{\mathbf{u}}(\xi, \mathbf{u})$. We further assume that no additional output constraints are present, i.e., $\mathcal{Y} = \mathbb{R}^{pN_f}$. For some general observations on the effect of output constraints, see [12, Section III.C], and specific results regarding (terminal) equality constraints and general feasibility are given in Sections 6.4 and 6.5, respectively. In [12, Thm 3, 4], results for an implicit predictor with reference $\mathbf{y}_{\text{ref}} = \mathbf{0}$ are given, which were later extended to possibly nonzero reference in [14, Thm. 2]. Under Assumption 1, solving (16) comes down to solving the unconstrained optimization problem

$$\hat{\mathbf{y}}_{\text{DPC}}(\xi, \mathbf{u}) = \arg \min_{\mathbf{y}} \|\mathbf{y} - \mathbf{y}_{\text{ref}}\|_{\mathbf{Q}}^2 + h^*(\xi, \mathbf{u}, \mathbf{y}).$$

Furthermore, since the trajectory-specific cost of $h(\mathbf{a}) = \lambda \|\mathbf{a}\|_2^2$ and $h(\mathbf{a}) = \lambda \|\Pi_{\perp} \mathbf{a}\|_2^2$ only differ in terms independent of $\mathbf{y}$ (see (10) and (15)), it turns out that the predictive behavior of both DPC schemes can be characterized by the same implicit predictor

$$\begin{aligned} \hat{\mathbf{y}}_{\text{DPC}}(\xi, \mathbf{u}) &= (\lambda \mathbf{Q}_{\text{reg}} + \mathbf{Q})^{-1} \lambda \mathbf{Q}_{\text{reg}} \hat{\mathbf{y}}_{\text{LS}}(\xi, \mathbf{u}) \\ &\quad + (\lambda \mathbf{Q}_{\text{reg}} + \mathbf{Q})^{-1} \mathbf{Q} \mathbf{y}_{\text{ref}}. \end{aligned} \quad (19)$$

Note that the involved inverse exists because $\mathbf{Q}_{\text{reg}}$ is positive definite (under Assumption 1) and $\mathbf{Q}$ is positive semidefinite. Structurally, (19) can be seen as a (matrix-)weighed sum, shifting output predictions towards $\mathbf{y}_{\text{ref}}$ and $\hat{\mathbf{y}}_{\text{LS}}(\xi, \mathbf{u})$, depending on the weights $\mathbf{Q}$ and $\lambda \mathbf{Q}_{\text{reg}}$, respectively. This structure is visualized in Figure 3 for the low-dimensional state-space example discussed in Section 2.3. The predictions blindly follow $\mathbf{y}_{\text{ref}}$ (independent of $\xi, \mathbf{u}$) for $\lambda = 0$ and tend towards $\hat{\mathbf{y}}_{\text{LS}}(\xi, \mathbf{u})$ for $\lambda \rightarrow \infty$. Furthermore, $\hat{\mathbf{y}}_{\text{DPC}}(\xi, \mathbf{u})$ is generally non-causal and time-variant (which cannot be observed in Figure 3 due to the limited prediction horizon $N_f = 1$). For finite λ and $\mathbf{y}_{\text{ref}} \neq \mathbf{0}$, it is also an affine predictor (otherwise linear), regardless of the true system class generating the data. This last point highlights the interpretation of implicit predictors being “the predictive behavior implicitly attributed to the data-generating system by the DPC scheme” [12], which may not necessarily match the true system if the scheme is poorly chosen.

6 Effects of modifications in DPC

While the previous section introduced the tools, namely trajectory-specific effect of regularization and implicit predictors, and exemplified them via results from [12]–[14], we now extend these results towards some common modifications in DPC. We analyze the extension to affine DPC in Section 6.1, the inclusion of an offset in the regularization in Section 6.2, the inclusion of slack variables in Section 6.3, and the inclusion of additional (terminal) equality constraints in Section 6.4. Finally, we also give two brief results on (recursive) feasibility for regularized DPC in Section 6.5.

6.1 DPC for affine systems

While standard (linear) DPC can yield exact predictions for deterministic LTI systems (see the discussion in Section 2.1), exact extensions to particular classes of nonlinear systems have been proposed, e.g., in [6], [7], [22]. Among these, we

want to briefly discuss the case of affine time-invariant (ATI) systems

$$\mathbf{x}(k+1) = \mathbf{A}\mathbf{x}(k) + \mathbf{B}\mathbf{u}(k) + \mathbf{e} \quad (20a)$$

$$\mathbf{y}(k) = \mathbf{C}\mathbf{x}(k) + \mathbf{D}\mathbf{u}(k) + \mathbf{r} \quad (20b)$$

proposed in [22], which is also used for other nonlinear systems with continuously updated trajectory data in order to approximate a local (affine) linearization of the nonlinear system for predictions as in [10], [22]. Similarly to how trajectories of LTI systems (2) can be generated by linear combinations, trajectories of ATI systems (20) can be generated by affine combinations of trajectory data. That is, in addition to (1), generated trajectories must also satisfy

$$\sum_{i=1}^{\ell} \mathbf{a}_i = \mathbf{1}_{\ell}^{\top} \mathbf{a} = 1. \quad (21)$$

Intuitively, this condition can be explained by noting that the effect of $\mathbf{e}, \mathbf{r}$ is present exactly once in each data trajectory and, accordingly, should be present exactly once in the generated trajectory. Assuming exact data generated by an ATI system, the affine hull $\text{aff}(\mathbf{D})$ of data columns in $\mathbf{D}$ is equivalent to the set of all possible system trajectories iff [29]

$$\text{rank} \begin{pmatrix} \mathbf{1}_{\ell}^{\top} \\ \mathbf{D} \end{pmatrix} = Lm + n + 1. \quad (22)$$

However, in the presence of noise and (other) nonlinearities, the same discussions as in Section 2.1 apply, i.e., the unique and exact predictions are no longer possible. Accordingly, our analysis is not confined to affine DPC applied to data from ATI systems (20), but rather extends to affine DPC with data generated by any system, including the nonlinear tracking case in [10]. To understand the features of affine DPC in the presence of such realistic data, we extend our results from the linear DPC case. As discussed in [19, Section II] and, in particular, [19, Rem. 4], many analysis results for linear DPC also apply to such nonlinear systems, which are linear in known (nonlinear) transformations of the state, input, and output. This also applies to the affine DPC scheme at hand, where we can simply consider $\check{\xi} := \begin{pmatrix} 1 & \xi^{\top} \end{pmatrix}^{\top}$ to be a nonlinear transformation of ξ , and $\check{\mathbf{W}} := \begin{pmatrix} \mathbf{1}_{\ell} & \mathbf{W}^{\top} \end{pmatrix}^{\top}$ contains the corresponding data. Similarly, we define

$$\check{\mathbf{z}} := \begin{pmatrix} \check{\xi} \\ \mathbf{u} \end{pmatrix} \quad \text{and} \quad \check{\mathbf{Z}} := \begin{pmatrix} \check{\mathbf{W}} \\ \mathbf{U} \end{pmatrix}$$

for the transformed state-input data. Instead of the linear least-squares estimates $\hat{\mathbf{y}}_{\text{LS}}(\xi, \mathbf{u}), \hat{\mathbf{u}}_{\text{LS}}(\xi)$, consider their affine counterparts

$$\hat{\mathbf{y}}_{\text{ALS}}(\xi, \mathbf{u}) := \check{\mathbf{G}}_{\text{ALS}} \check{\mathbf{z}} = \mathbf{G}_{\text{ALS}} \mathbf{z} + \mathbf{g}_{\text{ALS}}$$

$$\hat{\mathbf{u}}_{\text{ALS}}(\xi) := \check{\mathbf{K}}_{\text{ALS}} \check{\xi} = \mathbf{K}_{\text{ALS}} \xi + \mathbf{k}_{\text{ALS}}$$

with

$$\check{\mathbf{G}}_{\text{ALS}} := \arg \min_{\mathbf{G}} \|\mathbf{Y} - \mathbf{G}\check{\mathbf{Z}}\|_{\text{F}}^2 = \mathbf{Y}\check{\mathbf{Z}}^+ = \begin{pmatrix} \mathbf{g}_{\text{ALS}} & \mathbf{G}_{\text{ALS}} \end{pmatrix}$$

$$\check{\mathbf{K}}_{\text{ALS}} := \arg \min_{\mathbf{K}} \|\mathbf{U} - \mathbf{K}\check{\mathbf{W}}\|_{\text{F}}^2 = \mathbf{U}\check{\mathbf{W}}^+ = \begin{pmatrix} \mathbf{k}_{\text{ALS}} & \mathbf{K}_{\text{ALS}} \end{pmatrix},$$

and the corresponding residual matrices

$$\check{\mathbf{E}}_{\mathbf{y}} := \mathbf{Y} - \check{\mathbf{G}}_{\text{ALS}} \check{\mathbf{Z}} \quad \text{and} \quad \check{\mathbf{E}}_{\mathbf{u}} := \mathbf{U} - \check{\mathbf{K}}_{\text{ALS}} \check{\mathbf{W}}.$$

The analysis of affine DPC with regularizations in terms of their trajectory-specific effect and implicit predictors then follows accordingly. In the following, we present the case of (projection-based) quadratic regularization.

Proposition 1. For affine DPC, the trajectory-specific effect of $h(\mathbf{a}) = \lambda \|\mathbf{a}\|_2^2$ is given by

$$\lambda \|\mathbf{a}^*(\xi, \mathbf{u}, \mathbf{y})\|_2^2 = \lambda \|\mathbf{y} - \hat{\mathbf{y}}_{\text{ALS}}(\xi, \mathbf{u})\|_{\check{\mathbf{Q}}_{\text{reg}}}^2 \quad (23a)$$

$$+ \lambda \left\| \mathbf{u} - \hat{\mathbf{u}}_{\text{ALS}}(\xi) \right\|_{\check{\mathbf{R}}_{\text{reg}}}^2 \quad (23b)$$

$$+ \lambda \left\| \check{\xi} \right\|_{(\check{\mathbf{W}}\check{\mathbf{W}}^{\top})^{-1}}^2 \quad (23c)$$

with weighing matrices

$$\check{\mathbf{Q}}_{\text{reg}} := (\check{\mathbf{E}}_{\mathbf{y}} \check{\mathbf{E}}_{\mathbf{y}}^{\top})^{-1} \quad \text{and} \quad \mathbf{R}_{\text{reg}} := (\check{\mathbf{E}}_{\mathbf{u}} \check{\mathbf{E}}_{\mathbf{u}}^{\top})^{-1}.$$

Proof. The proof follows analogously to the linear DPC case. $\square$

The interpretation of these cost terms also directly follows from the discussion below (10). Importantly, note that instead of the least-squares estimates for a linear predictor/controller $\hat{\mathbf{y}}_{\text{LS}}(\xi, \mathbf{u}), \hat{\mathbf{u}}_{\text{LS}}(\xi)$ the regularization now favors the least-squares estimates for the affine predictor/controller $\hat{\mathbf{y}}_{\text{ALS}}(\xi, \mathbf{u}), \hat{\mathbf{u}}_{\text{ALS}}(\xi)$ and shifts the predicted trajectory towards them. Similarly to (15), the effect of the cost terms (23b) and (23c) can be eliminated by considering a projection-based quadratic regularization $h(\mathbf{a}) = \lambda \|\check{\mathbf{\Pi}}_{\perp} \mathbf{a}\|_2^2$ with

$$\check{\mathbf{\Pi}} := \check{\mathbf{Z}}^+ \check{\mathbf{Z}} \quad \text{and} \quad \check{\mathbf{\Pi}}_{\perp} := \mathbf{I} - \check{\mathbf{\Pi}}.$$

Finally, the analysis of predictive behavior via implicit predictors follows accordingly.

Proposition 2. Consider affine DPC with quadratic output-tracking objective (18), (projection-based) quadratic regularization $h(\mathbf{a}) = \lambda \|\mathbf{a}\|_2^2$ or $h(\mathbf{a}) = \lambda \|\tilde{\mathbf{\Pi}}_{\perp} \mathbf{a}\|_2^2$, and without additional output constraints. Under Assumption 1,

$$\hat{\mathbf{y}}_{\text{ADPC}}(\xi, \mathbf{u}) = \left(\mathbf{Q} + \lambda \check{\mathbf{Q}}_{\text{reg}} \right)^{-1} \lambda \check{\mathbf{Q}}_{\text{reg}} \hat{\mathbf{y}}_{\text{ALS}}(\xi, \mathbf{u}) + \left(\mathbf{Q} + \lambda \check{\mathbf{Q}}_{\text{reg}} \right)^{-1} \mathbf{Q} \mathbf{y}_{\text{ref}} \quad (24)$$

is an implicit predictor for this problem.

Proof. The proof follows analogously to the linear DPC case from the trajectory-specific effect of regularization in (23). $\square$

Again, the interpretation of this predictor follows from the linear DPC case discussed below (19).

6.2 Regularization with offset

Some DPC schemes, in which the tracking of a non-zero equilibrium is desired, modify the regularization by including an offset, i.e., $\tilde{h}(\mathbf{a}) = h(\mathbf{a} - \bar{\mathbf{a}})$ [10], [30]. Although this adjustment is based on the assumption that the original regularization shifts predicted trajectories toward zero (see our discussion in Section 3), it introduces some interesting features, which can be analyzed via its trajectory-specific effect. Although [30] considers various norm choices, we focus on the quadratic regularization $h(\mathbf{a}) = \lambda \|\mathbf{a} - \bar{\mathbf{a}}\|_2^2$ used in both [10], [30]. Additionally, while the results in [10] are framed for an affine DPC setting, our analysis will focus on linear DPC as in [30]. Note that these results can be readily extended to the affine case, as discussed in Section 6.1. Analogously to how $\mathbf{a}$ generates the trajectory $\mathbf{w} = \mathbf{D}\mathbf{a}$, we denote the trajectory generated by $\bar{\mathbf{a}}$ as $\bar{\mathbf{w}} := \mathbf{D}\bar{\mathbf{a}}$, and similarly with $\bar{\mathbf{z}}, \bar{\mathbf{u}}, \bar{\mathbf{y}}, \bar{\xi}$. Furthermore, we denote the difference between the two as

$$\Delta \mathbf{w} := \mathbf{D}\mathbf{a} - \mathbf{D}\bar{\mathbf{a}} = \mathbf{w} - \bar{\mathbf{w}},$$

and similarly with $\Delta \mathbf{z}, \Delta \mathbf{u}, \Delta \mathbf{y}, \Delta \xi$.

Theorem 1. Consider the DPC problem (6) with quadratic offset-regularization $h(\mathbf{a}) = \lambda \|\mathbf{a} - \bar{\mathbf{a}}\|_2^2$. Under Assumption 1, the trajectory-specific effect of this regularization is given by

$$\lambda \|\mathbf{a}^*(\xi, \mathbf{u}, \mathbf{y}) - \bar{\mathbf{a}}\|_2^2 = \lambda \|\Delta \mathbf{y} - \hat{\mathbf{y}}_{\text{LS}}(\Delta \xi, \Delta \mathbf{u})\|_{\mathbf{Q}_{\text{reg}}}^2 \quad (25a)$$

$$+ \lambda \|\Delta \mathbf{u} - \hat{\mathbf{u}}_{\text{LS}}(\Delta \xi)\|_{\mathbf{R}_{\text{reg}}}^2 \quad (25b)$$

$$+ \lambda \|\Delta \xi\|_{\left(\mathbf{w} \mathbf{w}^T\right)^{-1}}^2 \quad (25c)$$

with $\mathbf{Q}_{\text{reg}}$ and $\mathbf{R}_{\text{reg}}$ as defined in (14).

Proof. For a given tuple $(\xi, \mathbf{u}, \mathbf{y}, \bar{\mathbf{a}})$, one can check that the parametric optimal solution to (1) is given by

$$\mathbf{a}^*(\xi, \mathbf{u}, \mathbf{y}, \bar{\mathbf{a}}) = \mathbf{D}^+ \mathbf{w} + (\mathbf{I} - \mathbf{D}^+ \mathbf{D}) \bar{\mathbf{a}}$$

$$\lambda \|\mathbf{a}^*(\xi, \mathbf{u}, \mathbf{y}, \bar{\mathbf{a}}) - \bar{\mathbf{a}}\|_2^2 = \lambda \Delta \mathbf{w}^T (\mathbf{D} \mathbf{D}^T)^{-1} \Delta \mathbf{w}.$$

Except for the translation into Δ -coordinates, this expression is equivalent to the one obtained in [12, Section III.A], and the same block-matrix decomposition steps of $(\mathbf{D} \mathbf{D}^T)^{-1}$ lead to (25). $\square$

Hence, an offset by $\bar{\mathbf{a}}$ in the regularization simply translates to the same trajectory-specific effect with an offset $\bar{\mathbf{w}}$ for the considered trajectory. Note that (25) can also be equivalently expressed as

$$\begin{aligned} \lambda \|\mathbf{a}^*(\xi, \mathbf{u}, \mathbf{y}, \bar{\mathbf{a}}) - \bar{\mathbf{a}}\|_2^2 &= \lambda \|\mathbf{y} - \hat{\mathbf{y}}_{\text{LS}}(\xi, \mathbf{u}) - (\bar{\mathbf{y}} - \hat{\mathbf{y}}_{\text{LS}}(\bar{\xi}, \bar{\mathbf{u}}))\|_{\mathbf{Q}_{\text{reg}}}^2 \\ &\quad + \lambda \|\mathbf{u} - \hat{\mathbf{u}}_{\text{LS}}(\xi) - (\bar{\mathbf{u}} - \hat{\mathbf{u}}_{\text{LS}}(\bar{\xi}))\|_{\mathbf{R}_{\text{reg}}}^2 \\ &\quad + \lambda \|\xi - \bar{\xi}\|_{\left(\mathbf{w}_p \mathbf{w}_p^T\right)^{-1}}^2. \end{aligned} \quad (26)$$

While the third cost term is always irrelevant, we can see that the first two terms also may have no additional effect (compared to the usual quadratic regularization) if the trajectory $\bar{\mathbf{w}}$ adheres to

$$\bar{\mathbf{y}} = \hat{\mathbf{y}}_{\text{LS}}(\bar{\xi}, \bar{\mathbf{u}}), \quad \text{or} \quad \bar{\mathbf{u}} = \hat{\mathbf{u}}_{\text{LS}}(\bar{\xi}),$$

respectively. Similarly to (15), the effect of the cost terms (25b) and (25c) can be eliminated by considering a projection-based quadratic regularization. On that note, we briefly remark that both

$$\begin{aligned} h(\mathbf{a}) &= \lambda \|\mathbf{\Pi}_{\perp} (\mathbf{a} - \bar{\mathbf{a}})\|_2^2 = \lambda (\mathbf{a}^T \mathbf{\Pi}_{\perp} \mathbf{a} - 2\bar{\mathbf{a}}^T \mathbf{\Pi}_{\perp} \mathbf{a} + \bar{\mathbf{a}}^T \mathbf{\Pi}_{\perp} \bar{\mathbf{a}}) \\ \text{or } h(\mathbf{a}) &= \lambda \|\mathbf{\Pi}_{\perp} \mathbf{a} - \bar{\mathbf{a}}\|_2^2 = \lambda (\mathbf{a}^T \mathbf{\Pi}_{\perp} \mathbf{a} - 2\bar{\mathbf{a}}^T \mathbf{\Pi}_{\perp} \mathbf{a} + \bar{\mathbf{a}}^T \bar{\mathbf{a}}), \end{aligned}$$

yield the same effect, since they only differ in a constant term unrelated to $\mathbf{a}$. Furthermore, using the first term of the alternative cost expression (26), we can also state an implicit predictor as follows.

Proposition 3. Consider the DPC scheme (6) with quadratic output-tracking objective (18), (projection-based) quadratic offset-regularization $h(\mathbf{a}) = \lambda \|\mathbf{a} - \bar{\mathbf{a}}\|_2^2$ or $h(\mathbf{a}) = \lambda \|\mathbf{\Pi}_{\perp} \mathbf{a} - \bar{\mathbf{a}}\|_2^2$, and without additional output constraints. Under Assumption 1,

$$\hat{\mathbf{y}}_{\text{DPC,off}}(\xi, \mathbf{u}) = (\lambda \mathbf{Q}_{\text{reg}} + \mathbf{Q})^{-1} \lambda \mathbf{Q}_{\text{reg}} \hat{\mathbf{y}}_{\text{LS,off}}(\xi, \mathbf{u}) \\ + (\lambda \mathbf{Q}_{\text{reg}} + \mathbf{Q})^{-1} \mathbf{Q} \mathbf{y}_{\text{ref}}.$$

with

$$\hat{\mathbf{y}}_{\text{LS,off}}(\xi, \mathbf{u}) = \hat{\mathbf{y}}_{\text{LS}}(\xi, \mathbf{u}) + \bar{\mathbf{y}} - \hat{\mathbf{y}}_{\text{LS}}(\bar{\xi}, \bar{\mathbf{u}}) \quad (27)$$

is an implicit predictor for this problem.

Proof. The proof follows analogously to the linear DPC case from the trajectory-specific effect of regularization in (26). $\square$

We want to highlight that $\hat{\mathbf{y}}_{\text{LS,off}}(\xi, \mathbf{u})$ in (27) is an affine predictor mapping, which simply results from shifting the linear least-squares predictor $\hat{\mathbf{y}}_{\text{LS}}(\xi, \mathbf{u})$ by a constant term $\bar{\mathbf{y}} - \hat{\mathbf{y}}_{\text{LS}}(\bar{\xi}, \bar{\mathbf{u}})$, which depends on the discrepancy between said offset trajectory $\bar{\mathbf{w}}$ and the least-squares predictor. Other than that, the interpretation of this predictor follows from the linear DPC case discussed below (19).

6.3 Effect of slack variables in DPC

A common modification to DPC, first introduced in [2], is the inclusion of a slack variable as follows

$$\min_{\mathbf{u}, \mathbf{y}, \mathbf{a}} J(\xi, \mathbf{u}, \mathbf{y}) + h(\mathbf{a}) + h_{\sigma}(\sigma) \quad (28a)$$

$$\text{s.t.} \quad \begin{pmatrix} \mathbf{u}_p \\ \mathbf{y}_p \\ \mathbf{u} \\ \mathbf{y} \end{pmatrix} + \begin{pmatrix} \mathbf{0} \\ \sigma \\ \mathbf{0} \\ \mathbf{0} \end{pmatrix} = \begin{pmatrix} \mathbf{U}_p \\ \mathbf{Y}_p \\ \mathbf{U} \\ \mathbf{Y} \end{pmatrix} \mathbf{a}, \quad (28b)$$

$$(\mathbf{u}, \mathbf{y}) \in \mathcal{U} \times \mathcal{Y}. \quad (28c)$$

Here, we briefly decompose the notation from $\xi, \mathbf{W}$ back to $\mathbf{u}_p, \mathbf{y}_p, \mathbf{U}_p, \mathbf{Y}_p$ to show exactly which part the slack variable σ is acting on. The use of slack variables is common to avoid infeasibility of the initial condition $\mathbf{W}\mathbf{a} = \xi$ in situations where $\mathbf{W}$ does not have full row-rank. That is, while in the deterministic LTI case $\xi \in \mathcal{R}(\mathbf{W})$ is guaranteed by the system dynamics, even a tiny amount of noise or nonlinearity may cause a measured initial condition ξ to become inconsistent with $\mathcal{R}(\mathbf{W})$. Note that some works like [10] extend the use of slack variables to all output-related parts of the image representation, i.e., they add another slack variable to the last block-row of (28b), and our analyses will naturally extend to those cases as well. In order to keep the use of σ small, another regularization term $h_{\sigma}(\sigma)$ is added to the objective function, where typical choices

are $h_{\sigma}(\sigma) = \lambda_{\sigma} \|\sigma\|_1$ [2] or $h_{\sigma}(\sigma) = \lambda_{\sigma} \|\sigma\|_2^2$ [10], and we will focus on the latter in combination with quadratic regularization $h(\mathbf{a}) = \lambda \|\mathbf{a}\|_2^2$. While adding slack variables is not strictly necessary in cases where Assumption 1 applies, we will show an alternative interpretation of slack variables, which actually helps further justifying Assumption 1, instead. Consider the re-scaling

$$\tilde{\sigma} := \sqrt{\frac{\lambda_{\sigma}}{\lambda}} \sigma$$

for which we have $\lambda_{\sigma} \|\sigma\|_2^2 = \lambda \|\tilde{\sigma}\|_2^2$. Furthermore, we can rearrange (28b) as

$$\begin{pmatrix} \mathbf{u}_p \\ \mathbf{y}_p \\ \mathbf{u} \\ \mathbf{y} \end{pmatrix} = \begin{pmatrix} \mathbf{U}_p & \mathbf{0} \\ \mathbf{Y}_p & -\sqrt{\frac{\lambda}{\lambda_{\sigma}}} \mathbf{I}_{p \times N_p} \\ \mathbf{U} & \mathbf{0} \\ \mathbf{Y} & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{a} \\ \tilde{\sigma} \end{pmatrix} = \begin{pmatrix} \tilde{\mathbf{W}} \\ \mathbf{U} \\ \mathbf{Y} \end{pmatrix} \tilde{\mathbf{a}} =: \tilde{\mathbf{D}} \tilde{\mathbf{a}}$$

and sum up the regularization terms

$$\lambda_{\sigma} \|\sigma\|_2^2 + \lambda \|\mathbf{a}\|_2^2 = \lambda \|\tilde{\mathbf{a}}\|_2^2 =: h(\tilde{\mathbf{a}}).$$

Using this re-parameterization, the slack variables can be interpreted as adding artificial trajectory data columns in the new augmented data matrix $\tilde{\mathbf{D}}$, where $\tilde{\mathbf{W}}$ is now guaranteed to be full row-rank under the usual input persistency of excitation conditions (e.g., the one used in [1, Thm. 1]), even if $\mathbf{W}$ is not. If an additional slack variable is used for the last block-row of future outputs $\mathbf{y}$ (like in [1, [9]]), the same input persistency of excitation conditions guarantee $\tilde{\mathbf{D}}$ to be full row-rank, i.e., Assumption 1 to be satisfied, even if $\mathbf{D}$ is not. Hence, the consideration of slack variables further increases the scope of cases, where the (already commonly made) Assumption 1 is justified. Other than that, all analyses made using the trajectory-specific effect of regularization and implicit predictors for DPC with quadratic regularization $h(\mathbf{a}) = \lambda \|\mathbf{a}\|_2^2$ carry over by simply considering $\tilde{\mathbf{D}}$ instead of $\mathbf{D}$. The same does not apply to projection-based regularization $h(\mathbf{a}) = \lambda \|\Pi_{\perp} \mathbf{a}\|_2^2$, since the combined weight matrix $\mathbf{S} = \text{diag}(\Pi_{\perp}, \mathbf{I})$ for the resulting $h(\tilde{\mathbf{a}}) = \lambda \|\mathbf{S} \tilde{\mathbf{a}}\|_2^2$ has not yet been analyzed in the literature. For future work, we consider a more general analysis of quadratic regularization $h(\mathbf{a}) = \lambda \|\mathbf{S} \mathbf{a}\|_2^2$ with arbitrary weight matrix $\mathbf{S}$ an insightful endeavor.

Regarding the tuning of λ_{σ} with respect to λ , how prominently the artificial trajectories are used in the resulting DPC predictions mainly depends on the ratio $\lambda \lambda_{\sigma}^{-1}$ and its ratio w.r.t. the signal power of the data contained in $\mathbf{Y}_p$. Empirically, one can also observe an improved conditioning of the matrix $\tilde{\mathbf{W}} \tilde{\mathbf{W}}^{\top}$ compared to $\mathbf{W} \mathbf{W}^{\top}$ for increasing $\lambda \lambda_{\sigma}^{-1}$,

the quantification of which might also be interesting for future work. An interesting (and perhaps unexpected) link shows up in [31], where the authors iteratively optimize the weight λ for $\lambda_\sigma = 1$ (see [31, Eq. (33)]) in order to compute an approximation to the maximum-likelihood predictor for data generated by an LTI system with additive output noise.

6.4 Predictive behavior with (terminal) equality constraints

In order to provide closed-loop stability guarantees, classical MPC typically makes use of terminal ingredients (see, e.g., [32]). Similarly, practical stability for some DPC schemes like [1], [9], [10] has been proven by employing (among other modifications) terminal equality constraints for the last steps of the predicted I/O sequence (see [11] for a tutorial). While we view terminal constraints as the main use case of our following analysis, the results naturally expand to other kinds of equality constraints. When analyzing the effect of additional terminal equality constraints in the DPC scheme (6), first, let us briefly recall that the trajectory-specific effect of regularization introduced in Section 4 is universally unaffected, since it applies to any chosen triple $(\xi, \mathbf{u}, \mathbf{y})$ satisfying (8b), and thus also the ones satisfying additional (terminal equality) constraints. Regarding the characterization of predictive behavior via implicit predictors, similar considerations apply for the effect of additional constraints on the input sequence $\mathbf{u}$, as already discussed below (16).

However, the predictive behavior of DPC is significantly influenced by any kind of output constraints and often contradicts the unconstrained behavior as observed in [12, Section III.C]. To simplify upcoming notation, we assume




Figure 4: Implicit predictor $\hat{\mathbf{x}}_{\text{DPC}}(\mathbf{x}_0, \mathbf{u})$ for the DPC example discussed in Section 2.3 with prediction horizon $N_f = 2$ and reference. (a) No terminal constraints. (b) Terminal constraints. $\mathbf{x}_{\text{ref}} = (0.5 \ 0.5)^\top$. To deal with the higher dimensionality, the first (green) and second (orange) prediction step of $\hat{\mathbf{x}}_{\text{DPC}}(\mathbf{x}_0, \mathbf{u})$ are visualized individually. Furthermore, we constrain the second input to $\mathbf{u}_2 = 0$ and only visualize the first input $\mathbf{u}_1$.

that the terminal output constraints require the predicted output $\mathbf{y}$ to match the reference $\mathbf{y}_{\text{ref}}$ over the final n (or any other amount of) steps of the prediction. This assumption is made without loss of generality; if the original reference $\mathbf{y}_{\text{ref}}$ does not naturally satisfy this condition, we can simply define a modified reference $\tilde{\mathbf{y}}_{\text{ref}}$, where the last n steps are adjusted to incorporate the terminal constraints. Once again, we make use of Assumption 1, which is further justified by the use of slack variables in [1], [9], as discussed in Section 6.3. Furthermore, we assume

$$\mathbf{Q} := \text{diag}(\mathbf{Q}, \dots, \mathbf{Q}) \text{ with positive definite } \mathbf{Q} \in \mathbb{R}^{p \times p}, \quad (29)$$

which aligns with the setup in [1]. Given these assumptions, the following theorem precisely characterizes the effect of (terminal) output equality constraints on the predictive behavior of DPC with (projection-based) quadratic regularization.

Theorem 2. Consider the DPC problem (6) with $h(\mathbf{a}) = \lambda \|\mathbf{a}\|_2^2$ or $h(\mathbf{a}) = \lambda \|\mathbf{\Pi}_\perp \mathbf{a}\|_2^2$, and the additional terminal equality constraint

$$\begin{pmatrix} \mathbf{0} & \mathbf{I}_n \end{pmatrix} \mathbf{y} = \begin{pmatrix} \mathbf{0} & \mathbf{I}_n \end{pmatrix} \mathbf{y}_{\text{ref}}. \quad (30)$$

Furthermore, consider the block partitioning

$$\mathbf{Q} = \begin{pmatrix} \mathbf{Q}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{Q}_2 \end{pmatrix}, \quad \mathbf{Q}_{\text{reg}} = \begin{pmatrix} \mathbf{Q}_{\text{reg}}^{(11)} & \mathbf{Q}_{\text{reg}}^{(12)} \\ \mathbf{Q}_{\text{reg}}^{(21)} & \mathbf{Q}_{\text{reg}}^{(22)} \end{pmatrix}$$

with $\mathbf{Q}_2, \mathbf{Q}_{\text{reg}}^{(22)} \in \mathbb{R}^{n \times n}$ aligning with the assumption made on the structure of $\mathbf{Q}$ in (29). Under Assumption 1,

$$\hat{\mathbf{y}}_{\text{DPC}}(\xi, \mathbf{u}) = \Lambda_{\text{ref}} \mathbf{y}_{\text{ref}} + \Lambda_{\text{reg}} \hat{\mathbf{y}}_{\text{LS}}(\xi, \mathbf{u}) \quad (31)$$

with the weights

$$\Lambda_{\text{ref}} := \begin{pmatrix} (\mathbf{Q}_1 + \lambda \mathbf{Q}_{\text{reg}}^{(11)})^{-1} \mathbf{Q}_1 & -(\mathbf{Q}_1 + \lambda \mathbf{Q}_{\text{reg}}^{(11)})^{-1} \lambda \mathbf{Q}_{\text{reg}}^{(12)} \\ \mathbf{0} & \mathbf{I} \end{pmatrix}$$

$$\Lambda_{\text{reg}} := \begin{pmatrix} (\mathbf{Q}_1 + \lambda \mathbf{Q}_{\text{reg}}^{(11)})^{-1} \lambda \mathbf{Q}_{\text{reg}}^{(11)} & (\mathbf{Q}_1 + \lambda \mathbf{Q}_{\text{reg}}^{(11)})^{-1} \lambda \mathbf{Q}_{\text{reg}}^{(12)} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}$$

is an implicit predictor for this problem.

Proof. The proof strategy lies in characterizing the effect of hard terminal constraints as soft constraints with costs tending to infinity. Note that this is an exact characterization that works due to Assumption 1 allowing for any triple $(\xi, \mathbf{u}, \mathbf{y})$ in (6b) and would be invalid for an indirect (i.e., model-based) scheme, where the set of feasible $(\xi, \mathbf{u}, \mathbf{y})$ is limited by a

prediction model enforced as a hard constraint. The result thus follows by considering $\mathcal{Q}_q := \text{diag}(\mathcal{Q}_1, q\mathbf{I})$ instead of $\mathcal{Q}$ in (19) and taking the limit $q \rightarrow \infty$. For brevity, we only show the limit computation

$$\Lambda_{\text{ref}} = \lim_{q \rightarrow \infty} (\lambda \mathcal{Q}_{\text{reg}} + \mathcal{Q}_q)^{-1} \mathcal{Q}_q,$$

and note that

$$\Lambda_{\text{reg}} = \lim_{q \rightarrow \infty} (\lambda \mathcal{Q}_{\text{reg}} + \mathcal{Q}_q)^{-1} \lambda \mathcal{Q}_{\text{reg}}$$

follows analogously. From the structure and positive definiteness of $\mathcal{Q}_q$ we have $\mathcal{Q}_q^{-1} = \text{diag}(\mathcal{Q}_1^{-1}, q^{-1}\mathbf{I})$, and therefore

$$\begin{aligned} (\lambda \mathcal{Q}_{\text{reg}} + \mathcal{Q}_q)^{-1} \mathcal{Q}_q &= \left(\lambda \mathcal{Q}_q^{-1} \mathcal{Q}_{\text{reg}} + \mathbf{I} \right)^{-1} \\ &= \begin{pmatrix} \mathbf{I} + \mathcal{Q}_1^{-1} \lambda \mathcal{Q}_{\text{reg}}^{(11)} & \mathcal{Q}_1^{-1} \lambda \mathcal{Q}_{\text{reg}}^{(12)} \\ q^{-1} \lambda \mathcal{Q}_{\text{reg}}^{(21)} & \mathbf{I} + q^{-1} \lambda \mathcal{Q}_{\text{reg}}^{(22)} \end{pmatrix}^{-1}. \end{aligned}$$

Computing the limit for this expression yields

$$\begin{aligned} \lim_{q \rightarrow \infty} (\lambda \mathcal{Q}_{\text{reg}} + \mathcal{Q}_q)^{-1} \mathcal{Q}_q &= \begin{pmatrix} \mathbf{I} + \mathcal{Q}_1^{-1} \lambda \mathcal{Q}_{\text{reg}}^{(11)} & \mathcal{Q}_1^{-1} \lambda \mathcal{Q}_{\text{reg}}^{(12)} \\ \mathbf{0} & \mathbf{I} \end{pmatrix}^{-1} \\ &= \begin{pmatrix} (\mathbf{I} + \mathcal{Q}_1^{-1} \lambda \mathcal{Q}_{\text{reg}}^{(11)})^{-1} & -(\mathbf{I} + \mathcal{Q}_1^{-1} \lambda \mathcal{Q}_{\text{reg}}^{(11)})^{-1} \mathcal{Q}_1^{-1} \lambda \mathcal{Q}_{\text{reg}}^{(12)} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \\ &= \begin{pmatrix} (\mathcal{Q}_1 + \lambda \mathcal{Q}_{\text{reg}}^{(11)})^{-1} \mathcal{Q}_1 & -(\mathcal{Q}_1 + \lambda \mathcal{Q}_{\text{reg}}^{(11)})^{-1} \lambda \mathcal{Q}_{\text{reg}}^{(12)} \\ \mathbf{0} & \mathbf{I} \end{pmatrix}, \end{aligned}$$

where we used a well known block-matrix inversion formula in the third step. The computation of Λ_{reg} follows analogously. $\square$

As expected, the last n steps of the predicted $\hat{\mathbf{y}}_{\text{DPC}}(\xi, \mathbf{u})$ fully align with the terminal equality constraints given by (30), which can be seen from the zero and identity blocks in the lower block-row of Λ_{ref} and Λ_{reg} . Furthermore, it is remarkable that the structure of a (matrix-) weighed sum observed in (19) is preserved in the upper blocks of Λ_{ref} and Λ_{reg} , characterizing how the trade-off between $\mathbf{y}_{\text{ref}}$ and $\hat{\mathbf{y}}_{\text{LS}}(\xi, \mathbf{u})$ in the first $N_f - n$ prediction steps is affected by enforcing equality to the reference for the last n steps. This trade-off is visualized in Figure 4 for the low-dimensional state-space example discussed in Section 2.3. This visualization also highlights the discrepancy between the effect of terminal constraints in regularized DPC compared to model-based schemes, where a predictor is enforced as hard

constraints $\mathbf{y} = \hat{\mathbf{y}}(\xi, \mathbf{u})$. In the latter, a contradiction between $\hat{\mathbf{y}}(\xi, \mathbf{u})$ and the constraint (30) would lead to infeasibility of the corresponding $(\xi, \mathbf{u})$. Instead, regularized DPC not only changes the last n steps of predictions to match (30), but alters the first $N_f - n$ steps as specified in Theorem 2. While this preserves the feasibility, which we next discuss in Section 6.5, it may be an unintuitive and potentially unwanted feature for practitioners who are unaware of it.

6.5 (Recursive) feasibility in DPC

Recursive feasibility of the closed-loop is an important concept in stabilizing predictive control schemes, and a lot of focus is placed on either guaranteeing it *a priori* or certifying it for a given controller [33]. Correspondingly, practical stability results of DPC schemes also include results on recursive feasibility (see [1, Section IV.D], [11, Prop. IV.1], [9, Thm. 14]). Our work emphasizes the analysis of DPC schemes through the structure of the underlying OCP, without making assumptions about the class of systems generating the data. While we should not expect in-depth closed-loop analysis results without making such assumptions, our approach still reveals broad results on (recursive) feasibility based on the OCP structure, which seem to be currently overlooked in the literature. Although the DPC scheme (6) does not include constraints $\xi \in \mathcal{X}$ on the initial state, we will consider them here to state more nuanced and general results.

Proposition 4. *Consider the DPC scheme (6) with non-empty constraint sets $\mathcal{U}, \mathcal{Y}$, and additional state constraints $\xi \in \mathcal{X}$. Under Assumption 1, the set of feasible initial states ξ is equivalent to $\mathcal{X}$.*

Proof. Under Assumption 1, there exists an $\mathbf{a}$ satisfying (6b) for any triple $(\xi, \mathbf{u}, \mathbf{y})$. Hence, for any $\xi \in \mathcal{X}$ there also exist $(\mathbf{u}, \mathbf{y})$ satisfying the remaining constraints (6c). $\square$

The following is a simple consequence of this result.

Proposition 5. *Consider the DPC scheme (6) with non-empty constraint sets $\mathcal{U}, \mathcal{Y}$, and without additional state constraints, i.e., $\xi \in \mathbb{R}^{(m+p)N_p}$. Under Assumption 1, the OCP (6) is always feasible (and hence recursively feasible).*

Proof. Simply consider $\mathcal{X} = \mathbb{R}^{(m+p)N_p}$ in the proof of Proposition 4. $\square$

Although straightforward and already briefly discussed in [12, Section III.C], these results have important

implications. The DPC problem (6) is typically (i.e., under Assumption 1) always feasible. This feasibility is by design of the OCP itself, rather than coming from the closed-loop control or system dynamics. For the schemes in [1], [11], the use of slack variables on the output variables $\mathbf{y}_p$ and $\mathbf{y}$ ensures that Assumption 1 holds for the extended data matrix $\tilde{\mathbf{D}}$ (see our discussion in Section 6.3) under typical input persistency of excitation assumptions [1, Thm. 1]. Similarly, in [9], where slack variables are applied to the entire trajectory variable $\mathbf{w}$, Assumption 1 is also always satisfied for $\tilde{\mathbf{D}}$. Consequently, the recursive feasibility analysis in these schemes could be replaced by Proposition 5. Finally, we want to highlight the generality of these results. For example, the considered constraint sets $\mathcal{U}, \mathcal{Y}$ may also include terminal equality constraints like the ones discussed in Section 6.4.

7 Conclusions and outlook

This work discussed the use of trajectory-specific effects of regularizations (see Definition 1) and implicit predictors (see Definition 2) as analysis tools to improve explainability in regularized DPC. The former concretizes the effects of any regularization $h(\mathbf{a})$ by eliminating auxiliary variables and reformulating an equivalent cost $h^*(\xi, \mathbf{u}, \mathbf{y})$, which is specific to the trajectory variables $(\xi, \mathbf{u}, \mathbf{y})$, instead. The latter is a predictor mapping $\hat{\mathbf{y}}(\xi, \mathbf{u})$, which coincides with the predictions generated while applying regularized DPC and thus provides an indirect (i.e., model-based) view on this direct scheme, as visualized in Figure 1. We first demonstrated the use of these tools by summarizing results from [12]–[14] on trajectory-specific effects and implicit predictors for (projection-based) quadratic regularization, and provided intuitive interpretations of the uncovered effects in Sections 4 and 5. We then extended these results towards common modifications in DPC, including DPC for affine systems, offset regularizations, slack variables, and terminal constraints. Notably, we showed that slack variables can be interpreted as introducing artificial data samples in an augmented data matrix $\tilde{\mathbf{D}}$, which can be analyzed just like before and is even more likely to satisfy the already common Assumption 1. For each modification, we computed the trajectory-specific cost or presented an implicit predictor if the modification led to a change in the regularization effect or predictive behavior. A key feature of our results is modularity, since they can be freely combined. For example, the nonlinear tracking scheme [10] uses an affine formulation, offset regularization, slack variables, and terminal equality constraints. Hence, a combination of results from Sections 6.1–6.4 can be used to explain the interactions

between control objective, constraints, and regularization in this scheme. Finally, we gave two brief results on (recursive) feasibility, stating that common DPC formulations are often feasible by design, hence eliminating the need for a deeper analysis of recursive feasibility in those cases.

Although this work primarily focused on DPC with (projection-based) quadratic regularizations, we emphasize the broadness of our proposed analysis tools, which are (in principle) applicable to any choice of regularization $h(\mathbf{a})$. Therefore, similar analyses for more general quadratic regularization $h(\mathbf{a}) = \lambda \|\mathbf{S}\mathbf{a}\|_2^2$, 1-norm regularization $h(\mathbf{a}) = \lambda \|\mathbf{a}\|_1$ (which is sometimes used for nonlinear systems [2]), or general p-norms $h(\mathbf{a}) = \lambda \|\mathbf{a}\|_p$ should yield similarly insightful characterizations. Furthermore, similarly to the relation with γ -DDPC presented in Section 4.2, the trajectory-specific effect of regularizations may uncover links to other frameworks, such as generalized DPC [34]. Finally, combining the generality of the presented analysis tools with more specific assumptions and prior knowledge about the true underlying system could help establish new robustness and stability results.

Research ethics: Not applicable.

Informed consent: Not applicable.

Author contributions: All authors have accepted responsibility for the entire content of this manuscript and approved its submission.

Use of Large Language Models, AI and Machine Learning Tools: None declared.

Conflict of interest: The author states no conflict of interest.

Research funding: None declared.


Data availability: Not applicable.

References

- [1] J. Berberich, J. Köhler, M. Müller, and F. Allgöwer, “Data-driven model predictive control with stability and robustness guarantees,” *IEEE Trans. Autom. Control*, vol. 66, no. 4, pp. 1702–1717, 2021.
- [2] J. Coulson, J. Lygeros, and F. Dörfler, “Data-enabled predictive control: in the shallows of the DeePC,” in *18th European Control Conference*, 2019, pp. 307–312.
- [3] H. Yang and S. Li, “A new method of direct data-driven predictive controller design,” in *2013 9th Asian Control Conference*, 2013, pp. 1–6.
- [4] J. B. Rawlings, D. Q. Mayne, and M. M. Diehl, *Model Predictive Control: Theory, Computation, and Design*, 2nd ed., Santa Barbara, California, Nob Hill Publishing, 2017.
- [5] J. C. Willems, P. Rapisarda, I. Markovsky, and B. L. M. De Moor, “A note on persistency of excitation,” *Syst. Control Lett.*, vol. 54, no. 4, pp. 325–329, 2005.

- [6] J. Berberich and F. Allgöwer, “A trajectory-based framework for data-driven system analysis and control,” in *2020 European Control Conference*, 2020, pp. 1365–1370.
- [7] M. Alsalti, J. Berberich, V. G. Lopez, F. Allgöwer, and M. A. Müller, “Data-based system analysis and control of flat nonlinear systems,” in *2021 60th IEEE Conference on Decision and Control*, 2021, pp. 1484–1489.
- [8] J. Coulson, J. Lygeros, and F. Dörfler, “Regularized and distributionally robust data-enabled predictive control,” in *2019 IEEE 58th Conference on Decision and Control*, 2019, pp. 2696–2701.
- [9] M. Alsalti, M. Barkey, V. G. Lopez, and M. A. Müller, “Robust and efficient data-driven predictive control,” digital preprint, *arXiv:2409.18867*, 2024.
- [10] J. Berberich, J. Köhler, M. Müller, and F. Allgöwer, “Linear tracking MPC for nonlinear systems—part II: the data-driven case,” *IEEE Trans. Autom. Control*, vol. 67, no. 9, pp. 4406–4421, 2021.
- [11] J. Berberich, J. Köhler, M. A. Müller, and F. Allgöwer, “Stability in data-driven MPC: an inherent robustness perspective,” in *2022 IEEE 61st Conference on Decision and Control*, 2022, pp. 1105–1110.
- [12] M. Klädtke and M. Schulze Darup, “Implicit predictors in regularized data-driven predictive control,” *IEEE Control Syst. Lett.*, vol. 7, pp. 2479–2484, 2023.
- [13] M. Klädtke and M. Schulze Darup, “Towards a unifying framework for data-driven predictive control with quadratic regularization,” digital preprint, *arXiv:2404.02721*, 2024.
- [14] M. Klädtke, M. Schulze Darup, and D. E. Quevedo, “Extending direct data-driven predictive control towards systems with finite control sets,” in *2024 European Control Conference*, 2024, pp. 3345–3350.
- [15] I. Markovsky and F. Dörfler, “Identifiability in the behavioral setting,” *IEEE Trans. Autom. Control*, vol. 68, no. 3, pp. 1667–1677, 2023.
- [16] C. De Persis and P. Tesi, “Formulas for data-driven control: stabilization, optimality, and robustness,” *IEEE Trans. Autom. Control*, vol. 65, no. 3, pp. 909–924, 2020.
- [17] I. Markovsky, “Structured low-rank approximation and its applications,” *Automatica*, vol. 44, no. 4, pp. 891–909, 2008.
- [18] V. Breschi, A. Chiuso, and S. Formentin, “Data-driven predictive control in a stochastic setting: a unified framework,” *Automatica*, vol. 152, 2023, Art. no. 110961. <https://doi.org/10.1016/j.automatica.2023.110961>.
- [19] P. Mattsson, F. Bonassi, V. Breschi, and T. B. Schön, “On the equivalence of direct and indirect data-driven predictive control approaches,” *IEEE Control Syst. Lett.*, vol. 8, pp. 796–801, 2024.
- [20] F. Dörfler, J. Coulson, and I. Markovsky, “Bridging direct and indirect data-driven control formulations via regularizations and relaxations,” *IEEE Trans. Autom. Control*, vol. 68, no. 2, pp. 883–897, 2023.
- [21] L. Huang, J. Zhen, J. Lygeros, and F. Dörfler, “Quadratic regularization of data-enabled predictive control: theory and application to power converter experiments,” *IFAC-PapersOnLine*, vol. 54, no. 7, pp. 192–197, 2021.
- [22] J. Berberich, J. Köhler, M. A. Müller, and F. Allgöwer, “Data-driven model predictive control: closed-loop guarantees and experimental results,” *at — Automatisierungstechnik*, vol. 69, no. 7, pp. 608–618, 2021.
- [23] V. Breschi, A. Chiuso, M. Fabris, and S. Formentin, “On the impact of regularization in data-driven predictive control,” in *2023 62nd IEEE Conference on Decision and Control*, 2023, pp. 3061–3066.
- [24] W. Favoreel, B. De Moor, and M. Gevers, “SPC: subspace predictive control,” *IFAC Proc. Vol.*, vol. 32, no. 2, pp. 4004–4009, 1999.
- [25] M. Sader, Y. Wang, D. Huang, C. Shang, and B. Huang, “Causality-informed Data-Driven Predictive Control,” digital preprint, *arXiv:2311.09545*, 2023.
- [26] V. Breschi, M. Fabris, S. Formentin, and A. Chiuso, “Uncertainty-aware data-driven predictive control in a stochastic setting,” *IFAC-PapersOnLine*, vol. 56, no. 2, pp. 10083–10088, 2023.
- [27] A. Chiuso, M. Fabris, V. Breschi, and S. Formentin, “Harnessing uncertainty for a separation principle in direct data-driven predictive control,” digital preprint, *arXiv:2312.14788*, 2024.
- [28] F. Fiedler and S. Lucia, “On the relationship between data-enabled predictive control and subspace predictive control,” in *2021 European Control Conference*, 2021, pp. 222–229.
- [29] A. Padoan, F. Dörfler, and J. Lygeros, “Data-driven representations of conical, convex, and affine behaviors,” in *2023 62nd IEEE Conference on Decision and Control*, 2023, pp. 596–601.
- [30] E. Elokda, J. Coulson, P. N. Beuchat, J. Lygeros, and F. Dörfler, “Data-enabled predictive control for quadcopters,” *Int. J. Robust Nonlinear Control*, vol. 31, no. 18, pp. 8916–8936, 2021.
- [31] M. Yin, A. Iannelli, and R. S. Smith, “Maximum likelihood estimation in data-driven modeling and control,” *IEEE Trans. Autom. Control*, vol. 68, no. 1, pp. 317–328, 2023.
- [32] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, “Constrained model predictive control: stability and optimality,” *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.
- [33] J. Löfberg, “Oops! I cannot do it again: testing for recursive feasibility in MPC,” *Automatica*, vol. 48, no. 3, pp. 550–555, 2012.
- [34] M. Lazar and P. C. N. Verheijen, “Generalized data-driven predictive control: merging subspace and hankel predictors,” *Mathematics*, vol. 11, no. 9, p. 2216, 2023.

Bionotes



Manuel Klädtke

Department of Mechanical Engineering,
Control and Cyberphysical Systems Group,
TU Dortmund University, Leonhard-Euler-
Straße 2, 44227 Dortmund, Germany
manuel.klaedtke@tu-dortmund.de

Manuel Klädtke received a B.Sc. and a M.Sc. in Electrical Engineering from Paderborn University in 2019 and 2021, respectively. He is currently pursuing a Ph.D. in Control and Cyberphysical Systems Group at TU Dortmund University. His primary research focus is on data-driven predictive control schemes.

**Moritz Schulze Darup**

Department of Mechanical Engineering,
Control and Cyberphysical Systems Group,
TU Dortmund University, Leonhard-Euler-
Straße 2, 44227 Dortmund, Germany
moritz.schulzedarup@tu-dortmund.de

Moritz Schulze Darup received a Diploma degree in Mechanical Engineering, a B.Sc. in Physics, and a Ph.D. in Control Engineering from the Ruhr-Universität Bochum in 2008, 2010, and 2014, respectively. He became Assistant Professor and leader of an Emmy Noether group for encrypted control in 2019 at Paderborn University. Since 2020, he has been Full Professor for Control and Cyberphysical Systems at TU Dortmund University. His research interests include secure, predictive, and data-driven control.