Stochastic Data-Driven Predictive Control: Regularization, Estimation, and Constraint Tightening

Mingzhou Yin * **Andrea Iannelli** ** **Roy S. Smith** *

* *Automatic Control Laboratory, ETH Zürich, Switzerland (e-mail: {myin,rsmith}@control.ee.ethz.ch)*
** *Institute for Systems Theory and Automatic Control, University of Stuttgart, Germany (e-mail: andrea.iannelli@ist.uni-stuttgart.de)*

Abstract: Data-driven predictive control methods based on the Willems' fundamental lemma have shown great success in recent years. These approaches use receding horizon predictive control with nonparametric data-driven predictors instead of model-based predictors. This study addresses three problems of applying such algorithms under unbounded stochastic uncertainties: 1) tuning-free regularizer design, 2) initial condition estimation, and 3) reliable constraint satisfaction, by using stochastic prediction error quantification. The regularizer is designed by leveraging the expected output cost. An initial condition estimator is proposed by filtering the measurements with the one-step-ahead stochastic data-driven prediction. A novel constraint-tightening method, using second-order cone constraints, is presented to ensure high-probability chance constraint satisfaction. Numerical results demonstrate that the proposed methods lead to satisfactory control performance in terms of both control cost and constraint satisfaction, with significantly improved initial condition estimation.

Keywords: data-driven control, model predictive control, estimation and filtering, stochastic optimal control, stochastic system identification

1. INTRODUCTION

Classical model-based control enables simple but powerful control design by considering typically parametric mathematical abstractions of system behaviors, known as models. However, this comes at the cost of additional modeling and identification efforts, which constitutes the majority of the budget in model-based control design. In this regard, the concept of data-driven control provides an appealing alternative that designs the controller directly from data without parametric identification.

In this work, we focus on data-driven predictive control (DDPC), the data-driven counterpart to model predictive control (MPC). Similar to MPC, it solves a finite-horizon optimal control problem in a receding horizon fashion, but with nonparametric data-driven predictors instead of model-based predictors. Data-driven predictors for linear systems can be constructed by using the so-called Willems' fundamental lemma (Willems et al., 2005; Markovsky and Dörfler, 2023), which characterizes all possible system behaviors with finite data.

While these predictors work well with deterministic data, showing equivalence to model-based design, they become ill-defined with stochastic data. Multiple works have stressed this issue by introducing an inner problem that finds the 'optimal' predictor under some statistical principle, possibly with instrumental variables to reduce the effect of noise, e.g., Fiedler and Lucia (2021); Yin et al. (2023, 2022); Breschi et al. (2023); Dinkla et al. (2024). This idea is known as indirect DDPC (Dörfler et al., 2023).

These stochastic data-driven predictors can be applied to DDPC design similar to the noise-free case. However, they suffer from the following problems with this certainty-equivalent implementation: 1) the control cost does not account for the prediction error, 2) the initial condition measurements suffer from noise which cannot be improved by collecting more data, and 3) the constraint satisfaction is not guaranteed. Aspects of these problems have been analyzed under the robust control framework where a bounded uncertainty set is prescribed (Coulson et al., 2019b; Berberich et al., 2021; Huang et al., 2023; Klöppelt et al., 2022). On the other hand, the situation is less clear under the stochastic control framework, where existing works adopt restrictive assumptions, such as noise-free offline data (Kerz et al., 2023) and exact polynomial chaos expansions of stochastic measurements (Pan et al., 2023). This paper addresses these problems under general unbounded stochastic uncertainties, utilizing the prediction error quantification provided in Yin et al. (2022).

In particular, three modifications are made to the certainty-equivalent DDPC algorithm. 1) The nominal control cost is replaced by the expected control cost. This introduces an additional uncertainty term whose weight is specified using a statistical approach without requiring tuning in a stochastic framework. Similar regularizers have been used in DDPC problems, but without a systematic approach to tune them before deployment. 2) The output initial condition is estimated by Kalman-filtering the output measurements with one-step-ahead predictions from the previous timestep. This provides optimal initial condition estimation and significantly reduces the prediction error. 3) Output constraints are formulated as chance constraints and guaranteed by second-order cone (SOC) constraints. The effectiveness of these modifications is tested in a numerical example, where satisfactory control performance in terms of both control cost and constraint satisfaction is observed with significantly improved initial condition estimation.

Notation. The expected value, standard deviation, and covariance are denoted by $\mathbb{E}[\cdot]$, $\mathrm{std}(\cdot)$, and $\mathrm{cov}(\cdot)$, respectively. The symbol $\Pr(\cdot)$ indicates the probability of a random event. For a sequence of matrices $X_1, \ldots, X_n$, we denote the row-wise and diagonal-wise concatenation by $\mathrm{col}(X_1, \ldots, X_n)$ and $\mathrm{blkdiag}(X_1, \ldots, X_n)$, respectively. $\mathrm{diag}(\cdot)$ denotes the vector of diagonal elements of a square matrix. Given a signal $x : \mathbb{Z} \to \mathbb{R}^n$, its trajectory from k to $k+N-1$ is indicated by $(x_i)_{i=k}^{k+N-1} = \mathrm{col}(x_k, \ldots, x_{k+N-1})$. For a vector x, $\|x\|_P$ denotes the weighted l_2-norm $(x^\top P x)^{\frac{1}{2}}$. The rank and trace of a matrix are indicated by $\mathrm{rank}(\cdot)$ and $\mathrm{tr}(\cdot)$, respectively.

2. PROBLEM FORMULATION

Consider the observable part of a stable discrete-time linear time-invariant (LTI) dynamical system with disturbance and output noise, given by

$$\begin{cases} x_{t+1} &= Ax_t + Bu_t + Ew_t, \\ y_t &= Cx_t + Du_t + v_t, \end{cases} \tag{1}$$

where $x_t \in \mathbb{R}^{n_x}$, $u_t \in \mathbb{R}^{n_u}$, $y_t \in \mathbb{R}^{n_y}$, $w_t \in \mathbb{R}^{n_w}$, $v_t \in \mathbb{R}^{n_y}$ are the states, inputs, outputs, disturbance, and output noise, respectively. The noise-free output is denoted by y_t^0. The noise v_t is considered to be zero-mean i.i.d. with an arbitrary distribution and covariance $\mathrm{cov}(v_t) = \sigma^2 \mathbb{I}_{n_y}$; the disturbance w_t is assumed known in the collected offline trajectory and estimated online as detailed in Section 3.2.

The model parameters (A, B, C, D, E) are unknown, but a matrix of input-disturbance-output trajectory data $Z = \begin{bmatrix} z_0^d & \cdots & z_{M-1}^d \end{bmatrix}$ has been collected, where each column

$$z_i^d := \mathrm{col}\left(u_{t_i}^d, \ldots, u_{t_i+L-1}^d, w_{t_i}^d, \ldots, w_{t_i+L-1}^d, y_{t_i}^d, \ldots, y_{t_i+L-1}^d\right) \tag{2}$$

is a length-L trajectory of the system, where the superscript d denotes collected offline data. The availability of offline disturbance trajectories retroactively is commonly assumed in DDPC algorithms, e.g., Pan et al. (2023), and is practical in many applications. The noise in each column of outputs is assumed to be independent. This assumption holds exactly when the columns are separate trajectories or truncated from a longer trajectory with $t_{i+1} = t_i + L$, known as the Page construction (Iannelli et al., 2021). Another common construction of Z is by choosing $t_{i+1} = t_i + 1$, forming a block Hankel matrix. The matrix Z is dubbed the signal matrix (Yin et al., 2023).

We are interested in designing a receding horizon control algorithm with a predictor derived from the signal matrix Z instead of the model parameters. Such algorithms are known as indirect DDPC (Dörfler et al., 2023). Let $L = L_0 + L'$, and L_0 be no smaller than the observability index of the system. In this paper, we consider the stochastic optimal tracking problem within a horizon of L' that minimizes the following expected control cost

$$\min_{(\hat{u}_k^t)_{k=0}^{L'-1}} J_t := \sum_{k=0}^{L'-1} \left\|\hat{u}_k^t\right\|_R^2 + \mathbb{E}\left[\sum_{k=0}^{L'-1} \left\|\hat{y}_k^t - r_{t+k}\right\|_Q^2\right] \tag{3}$$

at time t, where $\hat{u}_k^t$ is the designed input at time $(t+k)$, $\hat{y}_k^t$ is a random variable that predicts the noise-free future output y_{t+k}^0, r_t denotes the reference trajectory, and Q, R are the output and the input cost matrices respectively.

It is also desired to constrain the outputs within a polytopic set $\mathscr{Y}_t := \{y_t \mid H^t y_t \le q^t\}$ at time t, where $H^t := \begin{bmatrix} h_1^t & \ldots & h_{n_c}^t \end{bmatrix}^\top \in \mathbb{R}^{n_c \times n_y}$ and $q^t := \mathrm{col}\left(q_1^t, \ldots, q_{n_c}^t\right) \in \mathbb{R}^{n_c}$. However, due to the existence of unbounded noise, the output constraints can only be satisfied with high probability as chance constraints, which will be detailed in later sections. The input is constrained to be in the set $\mathscr{U}^t$ at time t, i.e.,

$$\hat{u}_k^t \in \mathscr{U}_{t+k}, \ \forall\, k = 0, \ldots, L'-1. \tag{4}$$

Then, the first element in the optimal input sequence is applied to the system, i.e., $u_t := \hat{u}_0^t$ and the noisy output $y_t = y_t^0 + v_t$ is measured.

There are multiple aspects to consider when solving this problem, which will be discussed in the following sections.

1) There are uncertainties in both the prediction conditions and the signal matrix. An accurate predictor is desired under these uncertainties.
2) The expected control cost needs to be formulated as a tractable objective function.
3) Tractable formulations of the output constraints need to be derived.

3. STOCHASTIC DATA-DRIVEN PREDICTOR

3.1 Deterministic Prediction

With sufficiently persistently exciting inputs, the range space of Z contains all possible trajectories of the system in the noise-free case, i.e., $v_t = \mathbf{0}_{n_y}$. The following result derived from the Willems' fundamental lemma (Willems et al., 2005) provides a deterministic prediction by considering the augmented inputs $\psi_t := \mathrm{col}(u_t, w_t)$.

Proposition 1. Define a partition of Z as

$$Z := \mathrm{col}(U, W, Y_p, Y_f) := \mathrm{col}(\Psi, Y_p, Y_f), \tag{5}$$

where $U \in \mathbb{R}^{n_u L \times M}$, $W \in \mathbb{R}^{n_w L \times M}$, $Y_p \in \mathbb{R}^{n_y L_0 \times M}$, $Y_f \in \mathbb{R}^{n_y L' \times M}$, and $\Psi \in \mathbb{R}^{(n_u+n_w)L \times M}$. If $v_t = \mathbf{0}_{n_y}$ and $\mathrm{rank}(Z) = (n_u + n_w)L + n_x$, we have

$$\hat{\mathbf{y}}^t = Y_f\, g_{\mathrm{pinv}}^t,\ g_{\mathrm{pinv}}^t = \mathrm{col}(\Psi, Y_p)^\dagger\, \mathrm{col}\left(\mathbf{u}_{\mathrm{ini}}^t, \hat{\mathbf{u}}^t, \mathbf{w}^t, \mathbf{y}_{\mathrm{ini}}^t\right), \tag{6}$$

where $\hat{\mathbf{u}}^t := (\hat{u}_k^t)_{k=0}^{L'-1}$ denotes the input sequence within the control horizon, $\mathbf{u}_{\mathrm{ini}}^t := (u_k)_{k=t-L_0}^{t-1}$, $\mathbf{y}_{\mathrm{ini}}^t := (y_k)_{k=t-L_0}^{t-1}$ denote the immediate past input and output sequences of length L_0, and $\mathbf{w}^t := (w_k)_{k=t-L_0}^{t+L'-1}$ denotes the immediate past and future disturbance sequence of length L. The output sequence $\hat{\mathbf{y}}^t := (\hat{y}_k^t)_{k=0}^{L'-1}$ provides a deterministic output prediction within the control horizon.

Indirect DDPC with deterministic predictor (6) is known as subspace predictive control (Favoreel et al., 1999).

3.2 Stochastic Prediction

In this work, two sources of uncertainties are considered. 1) Noise in output measurements. This induces noise in the output part of the signal matrix Y_p, Y_f, and the past output sequence $\mathbf{y}_{\text{ini}}^t$ with $\mathbb{E}[\mathbf{y}_{\text{ini}}^t] = \bar{\mathbf{y}}_{\text{ini}}^t$, $\text{cov}(\mathbf{y}_{\text{ini}}^t) = P_t$. 2) Uncertainties in the online disturbance sequence $\mathbf{w}^t$ with $\mathbb{E}[\mathbf{w}^t] = \bar{\mathbf{w}}^t$, $\text{cov}(\mathbf{w}^t) = \Sigma_w$. Statistics $\bar{\mathbf{w}}^t$ and Σ_w can come from online measurements and predictions or prior knowledge.

Multiple algorithms have been developed to extend the deterministic prediction to the stochastic case. Such algorithms typically involve solving the following quadratic program:

$$\begin{aligned} g^t = \operatorname*{argmin}_g \ & \left\|Y_p g - \bar{\mathbf{y}}_{\text{ini}}^t\right\|_S^2 + \lambda \left\|g\right\|_2^2 \\ \text{s.t.} \quad & \Psi g = \text{col}\left(\mathbf{u}_{\text{ini}}^t, \hat{\mathbf{u}}^t, \bar{\mathbf{w}}^t\right), \end{aligned} \tag{7}$$

where λ and S are design parameters. Different choices of λ and S have been proposed, including the subspace predictor (Fiedler and Lucia, 2021), Wasserstein distance minimization (Lian et al., 2023), the signal matrix model (Yin et al., 2023, 2021), and the minimum mean-squared error predictor (Yin et al., 2022). See Yin et al. (2022) for a comparison between these choices. The quadratic program (7) admits the following closed-form solution:

$$g^t = \left[R_1 \ R_2 \ R_3 \ R_4\right] \text{col}\left(\mathbf{u}_{\text{ini}}^t, \mathbf{u}^t, \bar{\mathbf{w}}^t, \bar{\mathbf{y}}_{\text{ini}}^t\right), \tag{8}$$

where

$$[R_1 \ R_2 \ R_3] := F^{-1}\Psi^\top(\Psi F^{-1}\Psi^\top)^{-1}, \tag{9}$$

$$R_4 := \left(F^{-1} - F^{-1}\Psi^\top(\Psi F^{-1}\Psi^\top)^{-1}\Psi F^{-1}\right) Y_p^\top S, \tag{10}$$

and $F := \lambda \mathbb{I}_M + Y_p^\top S Y_p$.

The stochastic predictor can be constructed based on the solution g^t with the following lemma.

Lemma 2. The stochastic output sequence within the control horizon is given by

$$\mathbb{E}\left[\hat{\mathbf{y}}^t | \, g^t\right] = \bar{\mathbf{y}}^t, \ \text{cov}\left(\hat{\mathbf{y}}^t | \, g^t\right) = \Sigma^t, \tag{11}$$

where

$$\bar{\mathbf{y}}^t := Y_f g^t - \Gamma(Y_p g^t - \bar{\mathbf{y}}_{\text{ini}}^t), \tag{12}$$

$$\Sigma^t := \Gamma P_t \Gamma^\top + \Gamma_w \Sigma_w \Gamma_w^\top + \left\|g^t\right\|_2^2 T, \tag{13}$$

$$T := \sigma^2\left(\Gamma\Gamma^\top + \mathbb{I}_{n_y L'}\right), \quad \Gamma_w := (Y_f - \Gamma Y_p)R_3, \tag{14}$$

$$\Gamma := \text{col}\left(CA^{L_0}, \ldots, CA^{L-1}\right) \text{col}\left(C, \ldots, CA^{L_0-1}\right)^\dagger \tag{15}$$

is the autonomous transformation matrix from $\mathbf{y}_{\text{ini}}^t$ to $\hat{\mathbf{y}}^t$.

Proof. This comes directly from the proof of Theorem 1 in Yin et al. (2022) by considering the augmented inputs ψ_t. □

Unfortunately, Γ cannot be obtained exactly since A and C are unknown. However, this transformation matrix can also be estimated using a data-driven approach. Note that the true output prediction is given by $\hat{\mathbf{y}}_0^t = \Gamma \bar{\mathbf{y}}_{\text{ini}}^t$ if $\text{col}(\mathbf{u}_{\text{ini}}^t, \hat{\mathbf{u}}^t, \mathbf{w}^t) = \mathbf{0}$ and $P_t = \mathbf{0}$. Using the certainty equivalence principle, an estimate $\hat{\Gamma}_Z$ can be found by replacing $\hat{\mathbf{y}}_0^t$ with $\bar{\mathbf{y}}^t$. Then we have

$$\bar{\mathbf{y}}^t = Y_f R_4 \bar{\mathbf{y}}_{\text{ini}}^t - \hat{\Gamma}_Z(Y_p R_4 \bar{\mathbf{y}}_{\text{ini}}^t - \bar{\mathbf{y}}_{\text{ini}}^t) := \hat{\Gamma}_Z \, \bar{\mathbf{y}}_{\text{ini}}^t, \tag{16}$$

which leads to

$$\hat{\Gamma}_Z = Y_f R_4 \left(Y_p R_4\right)^{-1}. \tag{17}$$

In what follows, it is assumed that $\Gamma = \hat{\Gamma}_Z$. This estimate is correct in the noise-free case and consistent under mild conditions as shown in the following propositions.

Proposition 3. If $v_t = \mathbf{0}_{n_y}$, we have $\hat{\Gamma}_Z = \Gamma$.

Proof. If $v_t = \mathbf{0}_{n_y}$, we have $\sigma^2 = 0$. All designs of λ and S are equivalent to the subspace predictor, under which case R_4 is the last $n_y L_0$ columns of $\text{col}\left(\Psi, Y_p\right)^\dagger$ and thus $Y_p R_4 = \mathbb{I}_{n_y L_0}$. Then Lemma 2 in Yin et al. (2022) directly leads to $\hat{\Gamma}_Z = \Gamma$. □

Proposition 4. Let the singular values of $\text{col}\left(\Psi, Y_p\right)$ be $\sigma_1, \ldots, \sigma_{L_\sigma}$ in descending order, where $L_\sigma := (n_u + n_w)L + n_y L_0$. Then as $M \to \infty$, $\hat{\Gamma}_Z \to \Gamma$ w.p. 1, if $\sigma_{L_\sigma} \to \infty$.

Proof. Let $\text{col}\left(\Psi, Y_p\right) := \Omega S V^\top$ be the singular value decomposition, where $\Omega, S \in \mathbb{R}^{L_\sigma \times L_\sigma}$ and $V \in \mathbb{R}^{M \times L_\sigma}$. Then, $g_{\text{pinv}}^t = V S^{-1} \Omega^\top \omega^t$, where $\omega^t := \text{col}\left(\mathbf{u}_{\text{ini}}^t, \mathbf{u}^t, \bar{\mathbf{w}}^t, \bar{\mathbf{y}}_{\text{ini}}^t\right)$, and $\left\|g_{\text{pinv}}^t\right\|_2^2 \leq \left\|V\right\|_2^2 \left\|S^{-1}\right\|_2^2 \left\|\Omega\right\|_2^2 \left\|\omega^t\right\|_2^2 = \left\|\omega^t\right\|_2^2 / \sigma_{L_\sigma}^2$. Note that g_{pinv}^t is also the least-norm solution to the linear system $\text{col}\left(\Psi, Y_p\right) g = \omega^t$, so we have $\left\|g^t\right\|_2^2 \leq \left\|g_{\text{pinv}}^t\right\|_2^2$. Therefore, if $\sigma_{L_\sigma} \to \infty$, $\left\|g^t\right\|_2^2 \to 0$ and $\Sigma^t \to \mathbf{0}$ since $P_t = \mathbf{0}$. This directly leads to the convergence of $\hat{\Gamma}_Z$ to Γ. □

Remark 5. The singular value condition $\sigma_{L_\sigma} \to \infty$ requires that the columns of $\text{col}\left(\Psi, Y_p\right)$ activate all directions persistently as $M \to \infty$. This is satisfied for, for example, independent random or repeated full-rank inputs and disturbances.

4. STOCHASTIC INDIRECT DATA-DRIVEN PREDICTIVE CONTROL

Based on the stochastic predictor (11), the stochastic indirect DDPC algorithm can be proposed. In the following subsections, the stochastic control cost is first formulated as a quadratic objective. Then, the prediction accuracy is improved by filtering the output initial condition measurements with a Kalman filter. Finally, the satisfaction of chance constraints is guaranteed by formulating tightened SOC constraints.

4.1 Stochastic Control Cost

The stochastic control cost J_t is formulated as a quadratic function in the following lemma.

Lemma 6. The expected control cost is given by

$$J_t = \left\|\hat{\mathbf{u}}^t\right\|_{\bar{R}}^2 + \left\|\bar{\mathbf{y}}^t - \mathbf{r}^t\right\|_{\bar{Q}}^2 + \text{tr}\left(\bar{Q}T\right) \left\|g^t\right\|_2^2 + \text{constant}, \tag{18}$$

where $\bar{R} := \mathbb{I}_{L'} \otimes R$, $\bar{Q} := \mathbb{I}_{L'} \otimes Q$, $\mathbf{r}^t := (r_{t+k})_{k=0}^{L'-1}$, and $T := \sigma^2\left(\Gamma\Gamma^\top + \mathbb{I}_{n_y L'}\right)$. The cost is quadratic with respect to the optimization variable $\hat{\mathbf{u}}^t$.

Proof. The expected output cost is calculated as:

$$\begin{aligned} &\mathbb{E}\left[\left\|\hat{\mathbf{y}}^t - \mathbf{r}^t\right\|_{\bar{Q}}^2\right] = \mathbb{E}\left[\left(\bar{\mathbf{y}}^t + \mathbf{e}^t - \mathbf{r}^t\right)^\top \bar{Q}\left(\bar{\mathbf{y}}^t + \mathbf{e}^t - \mathbf{r}^t\right)\right] \\ &= \left\|\bar{\mathbf{y}}^t - \mathbf{r}^t\right\|_{\bar{Q}}^2 + \mathbb{E}\left[\left(\mathbf{e}^t\right)^\top \bar{Q}\, \mathbf{e}^t\right] = \left\|\bar{\mathbf{y}}^t - \mathbf{r}^t\right\|_{\bar{Q}}^2 + \text{tr}\left(\bar{Q}\Sigma^t\right) \\ &= \left\|\bar{\mathbf{y}}^t - \mathbf{r}^t\right\|_{\bar{Q}}^2 + \text{tr}\left(\bar{Q}T\right) \left\|g^t\right\|_2^2 + \text{const}, \end{aligned}$$

where $\mathbf{e}^t$: $\mathbb{E}\left[\mathbf{e}^t\right]=\mathbf{0}$, $\operatorname{cov}\left(\mathbf{e}^t\right)=\Sigma^t$ is the prediction error. The second to last equality is due to the cyclic property of the trace function. This cost is quadratic with respect to $\hat{\mathbf{u}}^t$ since both g^t and $\bar{\mathbf{y}}^t$ are linear with respect to $\hat{\mathbf{u}}^t$.□

The stochastic control cost adds a $\|g^t\|_2^2$-regularization term to the nominal cost. Such regularization is required in direct DDPC (Coulson et al., 2019a) for well-definedness and is proposed to enhance robustness in indirect DDPC (Yin et al., 2021). However, it was unclear how to tune the weighting factor for the regularizer other than trial and error. By considering the regularizer as the uncertainty term in the expected output cost, the weighting factor can be reliably selected as $\operatorname{tr}\left(\bar{Q}T\right)$, which depends on the output cost matrix and the noise level.

4.2 Initial Condition Estimation

In model-based output-feedback MPC, an estimator has to be designed to estimate the initial state of the predictor, which is not measurable. This is not required in DDPC since the output initial condition $\bar{\mathbf{y}}^t_{\text{ini}}$ can be directly measured. In fact, in most existing DDPC implementations with stochastic data, the output initial condition $\bar{\mathbf{y}}^t_{\text{ini}}$ comes from measurements as in the deterministic case, i.e., $\bar{\mathbf{y}}^t_{\text{ini}}:=(y_k)_{k=t-L_0}^{t-1}$. Thus, the associated covariance $P_t=\sigma^2\mathbb{I}$ is constant in (13). This source of uncertainty can be alleviated by choosing a larger L_0 (Yin et al., 2021). On the other hand, in the presence of stochastic uncertainties, a properly designed estimator can estimate the initial state with a diminishing covariance that is much smaller than the noise level in the measurements.

Therefore, although not required, it can be beneficial to improve the output initial condition measurements based on output predictions at previous time steps by designing an estimator, especially in cases where the online measurement error is large. In this subsection, a Kalman filter is designed as the estimator. In particular, we replace y_k with its Kalman-filtered counterpart for the output initial condition. This reduces the prediction errors by shrinking P_t as time progresses.

In detail, the same predictor (11) for predictive control design at time $(t-1)$ is used to filter the output at time t and update $\bar{\mathbf{y}}^t_{\text{ini}}$. The predictor can be considered as a non-minimal state-space "model" with "state"

$$\bar{x}_t:=\operatorname{col}\left(u_{t-L_0},\dots,u_{t-1},y^0_{t-L_0},\dots,y^0_{t-1}\right). \tag{19}$$

Let $\bar{y}^t_k$ and e^t_k denote the $(k+1)$-th block element of $\bar{\mathbf{y}}^t$ and $\mathbf{e}^t$, respectively, and Σ^t_k be the covariance of e^t_k, i.e., the $(k+1)$-th $n_y\times n_y$ block on the diagonal of Σ^t. The data-driven "model" is then given by

$$\begin{cases}\bar{x}_{t+1} &= \underbrace{\begin{bmatrix}\Lambda^{n_u} & \mathbf{0}\\ \mathbf{0} & \Lambda^{n_y}\end{bmatrix}}_{\bar{\Lambda}}\bar{x}_t+\begin{bmatrix}\mathbf{0}\\ \hat{u}^t_0\\ \mathbf{0}\\ \bar{y}^t_0+e^t_0\end{bmatrix},\\ \zeta_{t+1} &= \begin{bmatrix}\mathbf{0} & \mathbb{I}_{n_y}\end{bmatrix}\bar{x}_{t+1}+v_t=y^0_t+v_t=y_t,\end{cases} \tag{20}$$

where Λ^k denotes the k-step upper shift matrix with ones on the k-th superdiagonal. The covariances of the "process noise" e^t_0 and the measurement noise v_t are Σ^t_0 and $\sigma^2\mathbb{I}_{n_y}$, respectively. Then, a Kalman filter for (20) can be designed to estimate the initial condition $\bar{x}_t$. Let the state estimate and the output part of the state error covariance be $\bar{x}_{t,t}$ and $P_{t,t}$, respectively. Then, the initial conditions for the DDPC problem can be set as $\operatorname{col}\left(\mathbf{u}^t_{\text{ini}},\bar{\mathbf{y}}^t_{\text{ini}}\right):=\bar{x}_{t,t}$ and $P_t:=P_{t,t}$. The Kalman filtering algorithm is summarized in Algorithm 1.

Algorithm 1 Kalman filter in stochastic indirect DDPC

1: **Initialization:**

$$\bar{x}_{0,0}:=\operatorname{col}\left(u_{-L_0},\dots,u_{-1},y_{-L_0},\dots,y_{-1}\right), \tag{21}$$
$$P_{0,0}:=\mathbb{I}_{n_yL_0}. \tag{22}$$

2: **Prediction:**

$$\bar{x}_{t,t+1}:=\bar{\Lambda}\bar{x}_{t,t}+\operatorname{col}\left(\mathbf{0},\hat{u}^t_0,\mathbf{0},\bar{y}^t_0\right), \tag{23}$$
$$P_{t,t+1}:=\Lambda^{n_y}P_{t,t}\left(\Lambda^{n_y}\right)^\top+\Sigma^t_0. \tag{24}$$

3: **Update:**

$$K_{t+1}:=\Sigma^t_0\left(\Sigma^t_0+\sigma^2\mathbb{I}_{n_y}\right)^{-1}, \tag{25}$$
$$\bar{x}_{t+1,t+1}:=\bar{x}_{t,t+1}+\operatorname{col}\left(\mathbf{0},K_{t+1}\left(y_t-\bar{y}^t_0\right)\right), \tag{26}$$
$$P_{t+1,t+1}:=\left(\mathbb{I}_{n_y}-K_{t+1}\right)P_{t,t+1}. \tag{27}$$

Remark 7. A similar idea was proposed in Alpago et al. (2020) for a direct DDPC algorithm. However, no approach is provided to quantify the covariance of the prediction error required in the Kalman filter as there is no well-defined predictor in direct DDPC.

4.3 Chance Constraint Satisfaction

As mentioned in Section 2, the output constraints $y_t\in\mathscr{Y}_t$ cannot be guaranteed robustly under unbounded noise. Instead, high-probability chance constraints are considered, either element-wise as

$$\Pr\left(h_i^{t+k\top}\hat{y}^t_k\le q_i^{t+k}\right)\ge p,\ \forall\, i=1,\dots,n_c, k=0,\dots,L'-1, \tag{28}$$

or set-wise as

$$\Pr\left(\hat{y}^t_k\in\mathscr{Y}_{t+k}\right)\ge p,\ \forall\, k=0,\dots,L'-1, \tag{29}$$

where p is the targeted probability. These chance constraints are typically guaranteed by tightening the nominal constraints to account for prediction uncertainties. However, unlike standard model-based predictors with additive uncertainties, the prediction error covariance of the data-driven predictor (11) depends on the particular inputs and initial conditions via g^t. So the amount of constraint tightening cannot be calculated offline. Define the augmented linear constraints by $\bar{\mathscr{Y}}_t=\left\{\mathbf{y}\;\middle|\;\bar{H}^t\mathbf{y}\le\bar{q}^t\right\}$, where

$$\bar{H}^t:=\begin{bmatrix}\bar{h}^t_1 & \dots & \bar{h}^t_{L'n_c}\end{bmatrix}^\top:=\operatorname{blkdiag}\left(H^t,\dots,H^{t+L'-1}\right),$$
$$\bar{q}^t:=\operatorname{col}\left(\bar{q}^t_1,\dots,\bar{q}^t_{L'n_c}\right):=\operatorname{col}\left(q^t,\dots,q^{t+L'-1}\right).$$

The following lemma guarantees chance constraint satisfaction by constraint tightening.

Lemma 8. The constraint

$$\bar{q}^t-\bar{H}^t\bar{\mathbf{y}}^t\ge\mu\sqrt{\operatorname{diag}\left(\bar{H}^t\Sigma^t\bar{H}^{t\top}\right)} \tag{30}$$

guarantees the satisfaction of the chance constraints (28) if $\mu\ge\sqrt{\frac{1}{1-p}-1}$ and (29) if $\mu\ge\sqrt{\frac{n_y}{1-p}}$.

Proof. Applying the one-sided Chebyshev's inequality, we have

$$\Pr\left(\bar{h}^t_i\hat{\mathbf{y}}^t-\bar{h}^t_i\bar{\mathbf{y}}^t\le\sqrt{\tfrac{1}{1-p}-1}\cdot\operatorname{std}\left(\bar{h}^t_i\hat{\mathbf{y}}^t\right)\right)\ge p,\ \forall\, i, \tag{31}$$

Algorithm 2 Stochastic indirect DDPC

```
Select a data-driven predictor and calculate predictor parameters from (9), (10), (14), and (17).
Initialize the Kalman filter from Algorithm 1.
for t ← 0, 1, ... do
    col(u_ini^t, y_ini^t) ← x̄_{t,t}, P_t ← P_{t,t}
    û^t ← argmin_{û^t} (18) s.t. (8), (12), (4), (36).
    Apply u_t = û_0^t to the system and measure y_t.
    Run the Kalman filter from Algorithm 1.
end for
```

where $\operatorname{std}\left(\bar{h}_i^t \hat{\mathbf{y}}^t\right) = \sqrt{\bar{h}_i^t \Sigma^t \bar{h}_i^{t\top}}$. From (30), we have

$$\bar{q}_i^t - \bar{h}_i^t \bar{\mathbf{y}}^t \geq \mu \sqrt{\bar{h}_i^t \Sigma^t \bar{h}_i^{t\top}}. \tag{32}$$

Equations (31) and (32) lead to (28) for $\mu \geq \sqrt{\frac{1}{1-p} - 1}$.

From the multi-dimensional Chebyshev's inequality, the ellipsoidal set

$$\mathscr{E}_k^t := \left\{ e_k^t \,\middle|\, e_k^{t\top} \left(\Sigma_k^t\right)^{-1} e_k^t \leq \frac{n_y}{1-p} \right\} \tag{33}$$

is a confidence region of prediction error e_k^t with at least probability p. Then, the chance constraint is satisfied if

$$\bar{y}_k^t \in \mathscr{Y}_{t+k} \ominus \mathscr{E}_k^t := \left\{ y \,\middle|\, y + e \in \mathscr{Y}_{t+k}, \forall\, e \in \mathscr{E}_k^t \right\}, \tag{34}$$

where $\ominus$ denotes the Pontryagin difference. For polytope $\mathscr{Y}_{t+k}$ and ellipsoid $\mathscr{E}_k^t$, we have

$$\mathscr{Y}_{t+k} \ominus \mathscr{E}_k^t = \left\{ y \,\middle|\, h_i^{t+k\top} y \leq q_i^{t+k} - \eta_{\mathscr{E}_k^t}\left(h_i^{t+k}\right), i = 1, \ldots, n_c \right\}, \tag{35}$$

where $\eta_{\mathscr{E}_k^t}(h) := \sqrt{\frac{n_y}{1-p} h^\top \Sigma^t h}$ is the support function of $\mathscr{E}_k^t$. Aggregating the constraints for all i leads to (29) for $\mu \geq \sqrt{\frac{n_y}{1-p}}$. □

Remark 9. The lemma can be tightened if both v_t and $\mathbf{w}^t$ are Gaussian, by using the cumulative distribution functions of the χ^2 and Gaussian distributions.

Unfortunately, the tightened constraint (30) is not convex. The following corollary provides a convex surrogate for (30).

Corollary 10. The SOC constraint

$$\bar{q}^t - \bar{H}^t \bar{\mathbf{y}}^t \geq \mu \left(\mathbf{c}_1 + \mathbf{c}_2 \left\|g^t\right\|_2\right), \tag{36}$$

where

$$\mathbf{c}_1 := \sqrt{\operatorname{diag}\left(\bar{H}^t \left(\Gamma P_t \Gamma^\top + \Gamma_w \Sigma_w \Gamma_w^\top\right) \bar{H}^{t\top}\right)}, \tag{37}$$

$$\mathbf{c}_2 := \sqrt{\operatorname{diag}\left(\bar{H}^t T \bar{H}^{t\top}\right)}, \tag{38}$$

guarantees the satisfaction of (30).

Proof. Since $\sqrt{\sum_i a_i} \leq \sum_i \sqrt{a_i}$, we have

$$\mathbf{c}_1 + \mathbf{c}_2 \left\|g^t\right\|_2 \geq \sqrt{\mathbf{c}_1^2 + \mathbf{c}_2^2 \left\|g^t\right\|_2^2} = \sqrt{\operatorname{diag}\left(\bar{H}^t \Sigma^t \bar{H}^{t\top}\right)}. \quad \square$$

The proposed stochastic indirect DDPC algorithm is summarized in Algorithm 2.

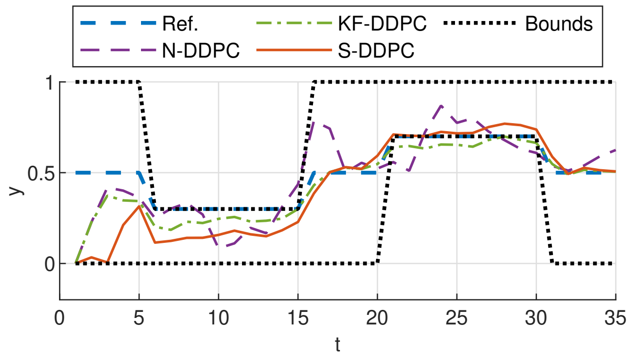

Fig. 1. Closed-loop trajectories of DDPC algorithms.

5. NUMERICAL EXAMPLE

In this section, we compare the performance of nominal DDPC (*N-DDPC*), DDPC with initial condition estimation in Algorithm 1 (*KF-DDPC*), and stochastic DDPC in Algorithm 2 (*S-DDPC*)[1]. Consider the following fourth-order dynamics:

$$\left[\begin{array}{c|c|c} A & B & E \\ \hline C & 0 & 0 \end{array}\right] = \left[\begin{array}{cccc|c|c} 0.36 & 0.64 & 0.07 & 0.02 & 0.29 & 0.03 \\ 0.42 & 0.58 & 0.02 & 0.07 & 0.03 & 0.20 \\ -9.34 & 9.34 & 0.23 & 0.58 & 4.90 & 1.07 \\ 5.88 & -5.88 & 0.39 & -0.39 & 1.07 & 3.48 \\ \hline 1 & 0 & 0 & 0 & 0 & 0 \end{array}\right].$$

The following parameters are used in the example: $L_0 = 4$, $L' = 10$, $Q = 20$, $R = 1$, $\sigma^2 = 0.01$, $p = 0.95$. An offline trajectory of length 500 is collected with unit Gaussian inputs and the signal matrix is constructed with a Hankel structure, which leads to $M = 487$. The statistics of the disturbance are given by $\bar{\mathbf{w}}^t = \mathbf{0}$ and $\Sigma_w = 0.001 \cdot \mathbb{I}$. Both the noise and the disturbance are Gaussian. The elementwise chance constraints (28) are used. The same online noise and disturbance sequences are used to compare the three algorithms. The minimum mean-squared error predictor in Yin et al. (2022) is employed as the predictor. No input constraint is considered in this example, i.e., $\mathscr{U}_t = \mathbb{R}$. Upper and lower output bounds are specified as the output constraints.

The closed-loop trajectories of the algorithms are presented in Figure 1, alongside the reference trajectory and the output bounds. As observed in Figure 1, *KF-DDPC* outperforms *N-DDPC* by introducing the initial condition estimator, although constraint violations are still evident. *S-DDPC* further eliminates the constraint violations experienced in *KF-DDPC* by constraint tightening. To underscore the effectiveness of the Kalman filter, Figure 2 showcases the comparison between the filtered output initial conditions and the measured ones for *S-DDPC*. The filtered trajectory is notably closer to the true trajectory compared to the measured trajectory.

The performance is further evaluated quantitatively by 50 Monte Carlo simulations with different noise and disturbance realizations. Figure 3 shows the boxplots of the true total control cost and the total amount of constraint violations, calculated as $\sum_t \max\left(H^t y_t - q^t, 0\right)$. The results validate our observations from Figure 1 that our proposed algorithm *S-DDPC* performs much better than the nom-

[1] The MATLAB scripts for reproducing the results are available at https://doi.org/10.3929/ethz-b-000672607.

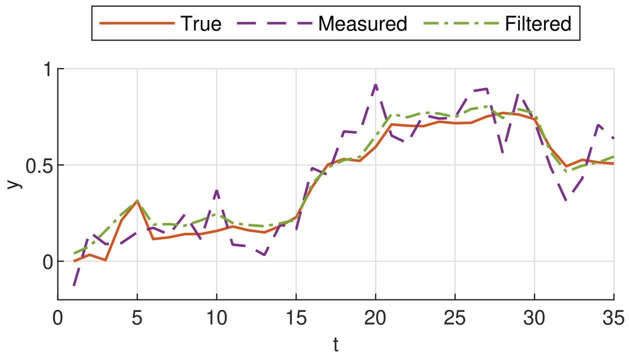

Fig. 2. Comparison of filtered and measured output trajectories.

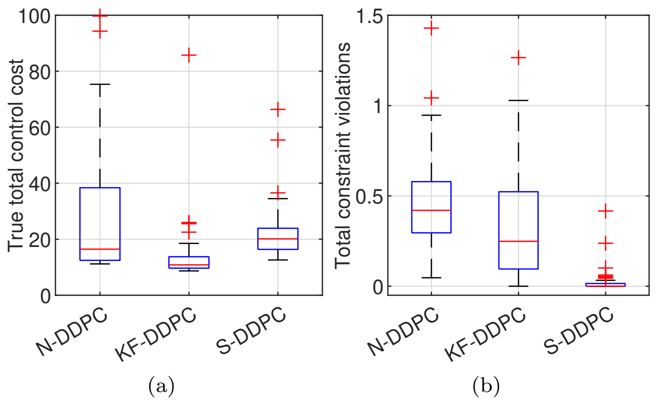

Fig. 3. Boxplots of (a) the true total control cost and (b) the total amount of constraint violations.

inal algorithm with almost no constraint violation. *KF-DDPC* obtains lower control costs than *S-DDPC*, but at the expense of larger constraint violations.

6. CONCLUSION

This work discusses several modifications in stochastic data-driven predictive control (DDPC) algorithms. They provide a tuning-free regularizer design in the control cost, improved initial condition estimation, and reliable constraint satisfaction. These are achieved by evaluating the expected cost, designing a Kalman filter, and formulating convex constraint tightening terms, respectively. These modifications pave the way for providing theoretical guarantees for DDPC algorithms under general unbounded stochasticity.

REFERENCES

Alpago, D., Dörfler, F., and Lygeros, J. (2020). An extended Kalman filter for data-enabled predictive control. *IEEE Control Systems Letters*, 4(4), 994–999.

Berberich, J., Köhler, J., Müller, M.A., and Allgöwer, F. (2021). Data-driven model predictive control with stability and robustness guarantees. *IEEE Transactions on Automatic Control*, 66(4), 1702–1717.

Breschi, V., Chiuso, A., and Formentin, S. (2023). Data-driven predictive control in a stochastic setting: a unified framework. *Automatica*, 152, 110961.

Coulson, J., Lygeros, J., and Dörfler, F. (2019a). Data-enabled predictive control: In the shallows of the DeePC. In *European Control Conference (ECC)*, 307–312.

Coulson, J., Lygeros, J., and Dörfler, F. (2019b). Regularized and distributionally robust data-enabled predictive control. In *IEEE Conference on Decision and Control (CDC)*, 2696–2701.

Dinkla, R., Mulders, S., Oomen, T., and van Wingerden, J.W. (2024). Closed-loop data-enabled predictive control and its equivalence with closed-loop subspace predictive control. *arXiv preprint arXiv:2402.14374*.

Dörfler, F., Coulson, J., and Markovsky, I. (2023). Bridging direct & indirect data-driven control formulations via regularizations and relaxations. *IEEE Transactions on Automatic Control*, 68(2), 883–897.

Favoreel, W., Moor, B.D., and Gevers, M. (1999). SPC: Subspace predictive control. *IFAC Proceedings Volumes*, 32(2), 4004–4009.

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

Huang, L., Zhen, J., Lygeros, J., and Dörfler, F. (2023). Robust data-enabled predictive control: Tractable formulations and performance guarantees. *IEEE Transactions on Automatic Control*, 68(5), 3163–3170.

Iannelli, A., Yin, M., and Smith, R.S. (2021). Design of input for data-driven simulation with hankel and page matrices. In *IEEE Conference on Decision and Control (CDC)*, 139–145.

Kerz, S., Teutsch, J., Brüdigam, T., Leibold, M., and Wollherr, D. (2023). Data-driven tube-based stochastic predictive control. *IEEE Open Journal of Control Systems*, 2, 185–199.

Klöppelt, C., Berberich, J., Allgöwer, F., and Müller, M.A. (2022). A novel constraint-tightening approach for robust data-driven predictive control. *International Journal of Robust and Nonlinear Control*.

Lian, Y., Shi, J., Koch, M., and Jones, C.N. (2023). Adaptive robust data-driven building control via bilevel reformulation: An experimental result. *IEEE Transactions on Control Systems Technology*, 31(6), 2420–2436.

Markovsky, I. and Dörfler, F. (2023). Identifiability in the behavioral setting. *IEEE Transactions on Automatic Control*, 68(3), 1667–1677.

Pan, G., Ou, R., and Faulwasser, T. (2023). Data-driven stochastic output-feedback predictive control: Recursive feasibility through interpolated initial conditions. In *Proceedings of the 5th Annual Learning for Dynamics and Control Conference*, 980–992.

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

Yin, M., Iannelli, A., and Smith, R.S. (2021). Maximum likelihood signal matrix model for data-driven predictive control. In *Proceedings of the 3rd Conference on Learning for Dynamics and Control*, 1004–1014.

Yin, M., Iannelli, A., and Smith, R.S. (2022). Data-driven prediction with stochastic data: Confidence regions and minimum mean-squared error estimates. In *European Control Conference (ECC)*, 853–858.

Yin, M., Iannelli, A., and Smith, R.S. (2023). Maximum likelihood estimation in data-driven modeling and control. *IEEE Transactions on Automatic Control*, 68(1), 317–328.