

A Unified Bayesian Framework for Data-Driven Smoothing, Prediction, and Control ^{*}

Mingzhou Yin ^{*} Andrea Iannelli ^{**} Seyed Ali Nazari ^{*}
Matthias A. Müller ^{*}

^{*} *Institute of Automatic Control, Leibniz University Hannover, 30167
Hannover, Germany (e-mail: yin@irt.uni-hannover.de,*

ali.nazari@stud.uni-hannover.de, mueller@irt.uni-hannover.de).

^{**} *Institute for Systems Theory and Automatic Control, University of
Stuttgart, 70569 Stuttgart, Germany (e-mail:
andrea.iannelli@ist.uni-stuttgart.de)*

Abstract: Extending data-driven algorithms based on Willems’ fundamental lemma to stochastic data often requires empirical and customized workarounds. This work presents a unified Bayesian framework for linear systems that provides a systematic and general method for handling stochastic data-driven tasks, including smoothing, prediction, and control, via maximum a posteriori estimation. This framework formulates a unified trajectory estimation problem and solves a Bayesian problem that optimally combines trajectory knowledge with trajectory characterization from offline data. This problem generalizes existing data-driven prediction and control algorithms. Numerical examples demonstrate the performance of the unified approach for all three tasks against other data-driven and system identification approaches.

Keywords: data-driven control, data-driven prediction, Bayesian methods, filtering and smoothing, stochastic system identification

1. INTRODUCTION

Data-driven analysis and control of linear systems, utilizing Willems’ fundamental lemma (Willems et al., 2005), have been a topic of considerable interest in recent years. When deterministic data are available, behavioral systems theory provides a powerful framework that allows for treating various data-driven problems in a unified manner, such as simulation and output matching (Markovsky and Dörfler, 2022). Since unique input-output behaviors can be identified with sufficiently informative data, equivalence can be shown for various data-driven formulations in the deterministic case (Coulson et al., 2019; Fiedler and Lucia, 2021).

Approaches start to diverge when stochastic uncertainties arise. One path is to recover the low-dimensional structure in the deterministic case via projection (Breschi et al., 2023) or low-rank matrix approximation (Markovsky and Dörfler, 2022). This is a practical choice, dating back to subspace identification (Van Overschee and De Moor, 2012), but it requires knowledge of the model dimension and fails to encode specific stochastic properties for finite-sample optimality. Others apply specific techniques to particular data-driven tasks, such as distributionally robust optimization (Coulson et al., 2022), empirical regularizers (Dörfler et al., 2023; Coulson et al., 2019), polynomial chaos expansions (Pan et al., 2023), Kalman filtering (Smith et al., 2024; Yin et al., 2024), kernel methods with truncated ARX models (Chiuso et al., 2025), and maximum likelihood estimation (Yin et al., 2023). Yet,

^{*} This work was supported by the Lower Saxony Ministry for Science and Culture within the program *zukunft.niedersachsen*.

these techniques lack a coherent framework for general data-driven tasks with stochastic data.

This work introduces a unified Bayesian framework for linear systems that addresses stochastic data-driven problems through maximum a posteriori (MAP) estimation. This framework can be seen as an extension of the behavioral systems theory to stochastic data. A wide class of stochastic uncertainties is considered, which allows for input errors, correlated noise, and elliptical distributions. Specifically, when a stochastic signal matrix is available, the following three data-driven tasks are considered.

- (**T1**) Smoothing: denoising measured input-output trajectories.
- (**T2**) Prediction: estimating future outputs given the immediate past input-output trajectory and future inputs.
- (**T3**) Optimal control: designing future inputs and outputs given the immediate past input-output trajectory and input-output references.

All these tasks can be posed as a trajectory estimation problem by specifying different types of observations and design requirements. An MAP estimator is then proposed that optimally combines the trajectory knowledge with a linear combination of offline data, following the characterization from Willems’ fundamental lemma. The linear combination vector is treated as a hyperparameter and estimated using the maximum marginal likelihood method. An iterative solution to the hyperparameter estimation problem is derived using sequential quadratic programming (SQP). This demonstrates connections to existing algorithms in data-driven prediction and control under special conditions. Numerical examples show that

the proposed framework achieves competitive performance against existing data-driven algorithms and classical approaches based on the system identification paradigm for all three tasks.

Notation. The row-wise and diagonal-wise concatenation and vectorization of matrices are denoted by $\text{col}(\cdot, \dots, \cdot)$, $\text{blkdiag}(\cdot, \dots, \cdot)$, and $\text{vec}(\cdot)$, respectively. The symbols $\mathbb{N}_{[m,n]}$, \mathbb{S}_+^n , and \mathbb{S}_{++}^n represent the set of natural numbers between m and n , positive semidefinite, and positive definite matrices of dimension $n \times n$, respectively.

2. BACKGROUND AND PROBLEM FORMULATION

Consider a discrete-time linear time-invariant (LTI) system $y_t^0 = Gu_t^0$ with n_x underlying states, where $u_t^0 \in \mathbb{R}^{n_u}$ and $y_t^0 \in \mathbb{R}^{n_y}$ are the true inputs and outputs, respectively. Suppose the inputs and outputs are subject to uncertainties with $u_t = u_t^0 + w_t$ and $y_t = y_t^0 + v_t$, where w_t and v_t denote the actuation and measurement uncertainties, respectively. All the perturbed signals at time t are denoted by $z_t := \text{col}(u_t, y_t) = z_t^0 + \nu_t \in \mathbb{R}^n$, $n := n_u + n_y$, where $z_t^0 := \text{col}(u_t^0, y_t^0)$ is the unknown true signals and $\nu_t := \text{col}(w_t, v_t)$ denotes the uncertainties.

This paper considers the problem of obtaining the optimal estimate of a length- L signal trajectory from possibly partial information of the corrupted trajectory and historical data of signal trajectories. In particular, assume that we have obtained M full signal trajectories of length L : $\mathbf{z}_i^d := \text{col}(z_{t_i+1}^d, z_{t_i+2}^d, \dots, z_{t_i+L}^d) \in \mathbb{R}^{nL}$, for $i = 1, 2, \dots, M$, where the superscript d denoted historical data. Concatenating these trajectories column-wise formulates the signal matrix: $H := [\mathbf{z}_1^d \ \mathbf{z}_2^d \ \dots \ \mathbf{z}_M^d] \in \mathbb{R}^{nL \times M}$, which is referred to as the offline data in what follows. For analysis purposes, we also define the true signal matrix H^0 and the uncertainty matrix Δ_H : $H = H^0 + \Delta_H$.

The objective is to estimate another signal trajectory $\mathbf{z}^0 := \text{col}(z_1^0, z_2^0, \dots, z_L^0) \in \mathbb{R}^{nL}$ from offline data (instead of the model G) and existing knowledge of the trajectory. This trajectory is referred to as the online trajectory. The knowledge is expressed by

$$\zeta = \Phi \mathbf{z}^0 + \epsilon, \quad (1)$$

where $\zeta \in \mathbb{R}^p$ encodes direct observations and possibly design requirements of the trajectory \mathbf{z}^0 , $\epsilon \in \mathbb{R}^p$ represents the uncertainty of the knowledge, and $\Phi \in \mathbb{R}^{p \times nL}$ is a known transformation matrix.

By specifying different types of existing knowledge, (1) encompasses a broad range of data-driven tasks in systems and control as follows.

(T1) Smoothing. The complete trajectory is available, but it is corrupted. In this case, we have $p = nL$,

$$\Phi = \mathbb{I}_{nL}, \quad \zeta = \text{col}(z_1, z_2, \dots, z_L), \quad \epsilon = \text{col}(\nu_1, \nu_2, \dots, \nu_L).$$

(T2) Prediction. The complete input sequence and the initial output sequence of length $L_0 \geq l$ are available, where l is the lag of G . Let $L' := L - L_0$. In this case, we have $p = nL_0 + n_u L'$,

$$\begin{aligned} \Phi &= \text{blkdiag}(\mathbb{I}_{nL_0}, \mathbb{I}_{L'} \otimes [\mathbb{I}_{n_u} \ \mathbf{0}_{n_u \times n_y}]), \\ \zeta &= \text{col}(z_1, \dots, z_{L_0}, u_{L_0+1}, \dots, u_L), \\ \epsilon &= \text{col}(\nu_1, \dots, \nu_{L_0}, w_{L_0+1}, \dots, w_L). \end{aligned}$$

When $L_0 \geq l$, from (14) of Iannelli et al. (2021), there exists K such that $H^0 = K\Phi H^0$. This guarantees that $\text{rank}(H^0) = \text{rank}(\Phi H^0)$ in Assumption 4 below.

(T3) Optimal control. An initial input-output sequence of length $L_0 \geq l$ is available. We would like to design an input-output sequence $(z_t^0)_{t=L_0+1}^L$ such that the quadratic control cost $\sum_{t=L_0+1}^L (\|u_t^0 - u_t^{\text{ref}}\|_R^2 + \|y_t^0 - y_t^{\text{ref}}\|_Q^2)$ is minimized, where $u_t^{\text{ref}} \in \mathbb{R}^{n_u}$, $y_t^{\text{ref}} \in \mathbb{R}^{n_y}$ are the input and output references and $R \in \mathbb{S}_{++}^{n_u}$, $Q \in \mathbb{S}_{++}^{n_y}$ are the input and output cost matrices, respectively. Let $\zeta_t^{\text{ctr}} := \text{col}(u_t^{\text{ref}}, y_t^{\text{ref}}) \in \mathbb{R}^n$. In this case, we have $p = nL$

$$\begin{aligned} \Phi &= \mathbb{I}_{nL}, \quad \zeta = \text{col}(z_1, \dots, z_{L_0}, \zeta_{L_0+1}^{\text{ctr}}, \dots, \zeta_L^{\text{ctr}}), \\ \epsilon &= \text{col}(\nu_1, \dots, \nu_{L_0}, \epsilon^{\text{ctr}}), \end{aligned}$$

where ϵ^{ctr} denotes the errors from the references.

Then, we are ready to state the main problem of this work.

Problem 1. Given the offline data H and the online trajectory observation ζ , obtain an optimal estimate $\hat{\mathbf{z}}$ of \mathbf{z}^0 .

The offline data H will be incorporated into estimation as the model surrogate through an MAP framework, which will be detailed in Section 3.

For tractability, elliptical distributions and stationary elliptical processes are considered in this work. An m -dimensional elliptical distribution, denoted by $\mathcal{E}_f(\mu, \Sigma)$, has a probability density function in the form of $p(x; \mu, \Sigma) = \det(\Sigma)^{-\frac{1}{2}} f(\|x - \mu\|_{\Sigma^{-1}}^2)$, where $f(\cdot) \in [0, +\infty) \rightarrow [0, +\infty)$ with $\lim_{x \rightarrow +\infty} f(x) = 0$ and $\mu \in \mathbb{R}^m$ and $\Sigma \in \mathbb{S}_{++}^m$ are known as the location-scale parameters. Elliptical distributions are preserved under linear transformation, i.e., $x \sim \mathcal{E}_f(\mu, \Sigma)$ implies $Ax + b \sim \mathcal{E}_f(A\mu + b, A\Sigma A^\top)$ (Hult and Lindskog, 2002). Such distributions include the multivariate normal distribution ($f(x) \propto \exp(-\frac{1}{2}x)$) and the multivariate Student's t -distribution ($f(x) \propto (1 + x/\xi)^{-(\xi+m)/2}$ with parameter ξ). Stationary elliptical processes are defined as stationary processes with elliptical distributions (B ankestad et al., 2023).

Here, we slightly generalize the definition of elliptical distributions by allowing positive semidefinite Σ to account for the case where the uncertainty is zero in some directions. When $\text{rank}(\Sigma) = r < m$, consider the decomposition (U_1, U_2, Σ_1) of Σ , where

$$\Sigma =: [U_1 \ U_2] \begin{bmatrix} \Sigma_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} U_1^\top \\ U_2^\top \end{bmatrix}, \quad \begin{bmatrix} U_1^\top \\ U_2^\top \end{bmatrix} [U_1 \ U_2] = \mathbb{I}_m, \quad (2)$$

$U_1 \in \mathbb{R}^{m \times r}$, $U_2 \in \mathbb{R}^{m \times (m-r)}$, and $\Sigma_1 \in \mathbb{S}_{++}^r$. Then, the probability density function for singular Σ is given by

$$p(x; \mu, \Sigma) = \begin{cases} \det(\Sigma_1)^{-\frac{1}{2}} f(\|U_1^\top \delta\|_{\Sigma_1^{-1}}^2), & U_2^\top \delta = \mathbf{0}_{m-r}, \\ 0, & U_2^\top \delta \neq \mathbf{0}_{m-r}, \end{cases}$$

where $\delta := x - \mu$.

This work considers the following assumption.

Assumption 2. The random variable ϵ is subject to zero-mean elliptical distributions: $\mathcal{E}_f(\mathbf{0}_p, \Sigma_\epsilon)$, where $\Sigma_\epsilon \in \mathbb{S}_+^p$. The random variables ν_t and ν_t^d are realizations of zero-mean stationary elliptical processes. Define

$$\begin{bmatrix} \nu_t \\ \nu_{t+\tau} \end{bmatrix} \sim \mathcal{E}_f\left(\mathbf{0}_{2n}, \begin{bmatrix} \Sigma_\nu(0) & \Sigma_\nu(\tau) \\ \Sigma_\nu(\tau) & \Sigma_\nu(0) \end{bmatrix}\right), \quad (3)$$

for $\tau \in \mathbb{N}$, where $\Sigma_\nu(\tau) =: \text{blkdiag}(\Sigma_w(\tau), \Sigma_v(\tau)) \in \mathbb{S}_+^n$, $\Sigma_w(\tau) \in \mathbb{S}_+^{n_u}$, and $\Sigma_v(\tau) \in \mathbb{S}_+^{n_y}$. Similarly, $\Sigma_\nu^d(\tau) := \text{blkdiag}(\Sigma_w^d(\tau), \Sigma_v^d(\tau)) \in \mathbb{S}_+^n$ is defined for ν_t^d .

We allow for different levels of uncertainty in the offline data and the online trajectory, as the system may operate under different conditions. Depending on the data-driven tasks, we have $\Sigma_\epsilon = \Sigma_z$ for (T1) and $\Sigma_\epsilon = \Phi \Sigma_z \Phi^\top$ for (T2), where the (i, j) -th $n \times n$ -block element of Σ_z is $\Sigma_\nu(|i - j|)$. For (T3), it is natural to encode the control objective as a design requirement that z_t^0 is subject to an elliptical distribution: $z_t^0 \sim \mathcal{E}_f(\zeta_t^{\text{ctr}}, \Sigma^{\text{ctr}})$ for $t = L_0 + 1, \dots, L$, where $\Sigma^{\text{ctr}} := \text{blkdiag}(R^{-1}, Q^{-1}) \in \mathbb{S}_{++}^n$. Note that R and Q can be uniformly scaled to trade off against data uncertainties. In this case, we have $\Sigma_\epsilon = \text{blkdiag}(\Gamma \Sigma_z \Gamma^\top, \mathbb{I}_{L'} \otimes \Sigma^{\text{ctr}})$, where $\Gamma = [\mathbb{I}_{n_{L_0}} \ \mathbf{0}_{n_{L'}}]$.

2.1 Willems' Fundamental Lemma and Deterministic Case

We would like to estimate \mathbf{z}^0 by combining both direct knowledge of the trajectory through the observation ζ and the behavior of the LTI system through the signal matrix H . To characterize \mathbf{z}^0 with the signal matrix H , we utilize Willems' fundamental lemma as the foundation.

Theorem 3. (Markovsky and Dörfler, 2023, Corollary 19) Suppose $\text{rank}(H^0) = n_u L + n_x$. The input-output sequence \mathbf{z}^0 is a trajectory of G iff there exists $\mathbf{g} \in \mathbb{R}^M$, such that $\mathbf{z}^0 = H^0 \mathbf{g}$.

Theorem 3 extends the original Willems' lemma (Willems et al., 2005) by allowing for non-Hankel H and providing a necessary and sufficient rank condition. This theorem solves Problem 1 in the deterministic case, i.e., when $\epsilon = \mathbf{0}_p$ and $H = H^0$, which is given in the following corollary.

Assumption 4. $\text{rank}(H^0) = \text{rank}(\Phi H^0) = n_u L + n_x$.

Corollary 5. Suppose $\epsilon = \mathbf{0}_p$, $H = H^0$, and Assumption 4 is satisfied. The online trajectory is determined uniquely by $\mathbf{z}^0 = H \mathbf{g}$, where $\mathbf{g} \in \mathbb{R}^M$ satisfies $\zeta = \Phi H \mathbf{g}$.

This corollary generalizes Proposition 9 in Markovsky and Dörfler (2022). The proof remains valid with a linear transformation of H instead of selected rows of H .

Assumption 4 provides the identifiability condition for Problem 1. When uncertainties arise and H^0 is not available, the rank condition of H^0 can be verified by the sufficient condition in Theorem 2 of Van Waarde et al. (2020), i.e., G is controllable and the inputs are collectively persistently exciting of order $(L + n_x)$. The rank condition of ΦH^0 is satisfied for the problems considered in this work, as shown in the following subsection.

3. THE MAXIMUM A POSTERIORI FRAMEWORK

To solve Problem 1, this work considers an MAP framework, where ζ is treated as the observation and the Willems' fundamental lemma characterization is used as prior information.

In detail, \mathbf{z}^0 is estimated by solving

$$\hat{\mathbf{z}} = \underset{\mathbf{z}^0}{\text{argmax}} p(\mathbf{z}^0 | \zeta) = \underset{\mathbf{z}^0}{\text{argmax}} p(\zeta | \mathbf{z}^0) p(\mathbf{z}^0). \quad (4)$$

The conditional distribution of ζ given \mathbf{z}^0 is given by $\zeta | \mathbf{z}^0 \sim \mathcal{E}_f(\Phi \mathbf{z}^0, \Sigma_\epsilon)$. The prior distribution of \mathbf{z}^0 can be obtained from Theorem 3 using the equation $\mathbf{z}^0 = H \mathbf{g} - \Delta_H \mathbf{g}$, where \mathbf{g} is treated as a hyperparameter. To quantify the uncertainty of $\Delta_H \mathbf{g}$ given \mathbf{g} , we follow a similar derivation as in Yin et al. (2023) by vectorizing Δ_H : $\Delta_H \mathbf{g} = (\mathbf{g}^\top \otimes \mathbb{I}_{nL}) \text{vec}(\Delta_H)$. Let $\text{vec}(\Delta_H) \sim \mathcal{E}_f(\mathbf{0}_{nLM}, \Sigma_d)$, where $\Sigma_d \in \mathbb{S}_+^{nLM}$. For $i, j \in \mathbb{N}_{[1, ML]}$, define $i =: (i_1 - 1)L + i_2$, $j =: (j_1 - 1)L + j_2$, where $i_1, j_1 \in \mathbb{N}_{[1, M]}$ and $i_2, j_2 \in \mathbb{N}_{[1, L]}$. Then, Σ_d is given by

$$(\Sigma_d)_{i,j} = \Sigma_\nu^d(|t_{i_1} + i_2 - (t_{j_1} + j_2)|), \quad (5)$$

where $(\Sigma_d)_{i,j}$ is the (i, j) -th $n \times n$ -block element of Σ_d .

This leads to $\Delta_H \mathbf{g} | \mathbf{g} \sim \mathcal{E}_f(\mathbf{0}_{nL}, \Sigma_g(\mathbf{g}))$, where

$$\Sigma_g(\mathbf{g}) := (\mathbf{g}^\top \otimes \mathbb{I}_{nL}) \Sigma_d (\mathbf{g} \otimes \mathbb{I}_{nL}) \in \mathbb{S}_+^{nL}, \quad (6)$$

and thus the prior distribution: $\mathbf{z}^0 | \mathbf{g} \sim \mathcal{E}_f(H \mathbf{g}, \Sigma_g(\mathbf{g}))$.

The following lemma presents an efficient approach to calculate $\Sigma_g(\mathbf{g})$. In particular, we discuss two common constructions of H , namely the Page construction (Damen et al., 1982) with $t_{i+1} = t_i + L$ and the Hankel construction (Van Overschie and De Moor, 2012) with $t_{i+1} = t_i + 1$. Page construction has less noise correlation than the Hankel construction at the cost of fewer data columns in the signal matrix.

Lemma 6. Let $(\Sigma_g(\mathbf{g}))_{i,j}$ be the (i, j) -th $n \times n$ -block element of $\Sigma_g(\mathbf{g})$. The prior scale parameter $\Sigma_g(\mathbf{g})$ is given by

$$(\Sigma_g(\mathbf{g}))_{i,j} = \sum_{k=1}^M \sum_{m=1}^M g_k g_m \Sigma_\nu^d(|t_k + i - (t_m + j)|). \quad (7)$$

For Page and Hankel constructions of H , (7) is equivalent to

$$(\Sigma_g(\mathbf{g}))_{i,j} = \sum_{\tau=1-M}^{M-1} \Sigma_\nu^d(|\tau l + i - j|) K(\tau) \quad (8)$$

with $l = L$ for the Page construction and $l = 1$ for the Hankel construction, where $K(\tau) := \sum_{k=1}^{M-|\tau|} g_k g_{k+|\tau|}$, $\tau \in \mathbb{N}_{[1-M, M-1]}$ is the unnormalized empirical autocorrelation function of \mathbf{g} .

Proof. Let $\mathbf{g}^\top \otimes \mathbb{I}_{nL} =: [\gamma_1 \ \dots \ \gamma_{LM}]$, where $\gamma_i \in \mathbb{R}^{nL \times n}$. From (6), we have $\Sigma_g(\mathbf{g}) = \sum_{i=1}^{LM} \sum_{j=1}^{LM} \gamma_i (\Sigma_d)_{i,j} \gamma_j^\top$.

According to the structure of $(\mathbf{g}^\top \otimes \mathbb{I})$, we have $\gamma_i = \mathbf{e}_{i_2} \otimes g_{i_1}$, $\gamma_j = \mathbf{e}_{j_2} \otimes g_{j_1}$, where $\mathbf{e}_i \in \mathbb{R}^L$ is the unit vector with i -th non-zero entry. This leads to

$$\Sigma_g(\mathbf{g}) = \sum_{i=1}^{LM} \sum_{j=1}^{LM} g_{i_1} g_{j_1} (\mathbf{e}_{i_2} \otimes \mathbb{I}_n) (\Sigma_d)_{i,j} (\mathbf{e}_{j_2}^\top \otimes \mathbb{I}_n).$$

From (5), we have

$$(\Sigma_g(\mathbf{g}))_{i_2, j_2} = \sum_{i_1=1}^M \sum_{j_1=1}^M g_{i_1} g_{j_1} \Sigma_\nu^d(|t_{i_1} + i_2 - (t_{j_1} + j_2)|),$$

which is equivalent to (7). Consider variable transformations $\tau := i_1 - j_1$ and $k = \min(i_1, j_1)$. We note $g_{i_1} g_{j_1} = g_k g_{k+|\tau|}$, $t_{i_1} + i_2 - (t_{j_1} + j_2) = \tau L + i - j$ for the Page construction, and $t_{i_1} + i_2 - (t_{j_1} + j_2) = \tau + i - j$ for the Hankel construction. These directly lead to (8). \square

Then, we are ready to state the MAP estimator. When both Σ_ϵ and $\Sigma_g(\mathbf{g})$ are positive definite, (4) becomes

$$\hat{\mathbf{z}}(\mathbf{g}) = \operatorname{argmax}_{\mathbf{z}^0} f\left(\|\delta_\epsilon\|_{\Sigma_\epsilon^{-1}}^2\right) f\left(\|\delta_g\|_{\Sigma_g^{-1}(\mathbf{g})}^2\right), \quad (9)$$

where $\delta_\epsilon = \zeta - \Phi\mathbf{z}^0$ and $\delta_g := \mathbf{z}^0 - H\mathbf{g}$. Recall that $f(\cdot)$ defines the probability density function of \mathcal{E}_f . The scaling factors $\det(\Sigma_\epsilon)^{-\frac{1}{2}}$ and $\det(\Sigma_g)^{-\frac{1}{2}}$ are omitted since they do not depend on \mathbf{z}^0 .

When Σ_ϵ and/or $\Sigma_g(\mathbf{g})$ are singular, define $(U_{\epsilon,1}, U_{\epsilon,2}, \Sigma_{\epsilon,1})$ and $(U_{g,1}, U_{g,2}, \Sigma_{g,1}(\mathbf{g}))$ by applying the decomposition (2) on Σ_ϵ and $\Sigma_g(\mathbf{g})$, respectively. Then, the MAP estimator is given by the following constrained problem:

$$\begin{aligned} \hat{\mathbf{z}}(\mathbf{g}) = \operatorname{argmax}_{\mathbf{z}^0} & f\left(\|U_{\epsilon,1}^\top \delta_\epsilon\|_{\Sigma_{\epsilon,1}^{-1}}^2\right) f\left(\|U_{g,1}^\top \delta_g\|_{\Sigma_{g,1}^{-1}(\mathbf{g})}^2\right) \\ \text{s.t. } & U_{\epsilon,2}^\top \delta_\epsilon = \mathbf{0}, \quad U_{g,2}^\top \delta_g = \mathbf{0}. \end{aligned} \quad (10)$$

When the uncertainties are Gaussian, i.e., $f(x) \propto \exp(-\frac{1}{2}x)$, (9) is equivalent to a quadratic program:

$$\hat{\mathbf{z}}(\mathbf{g}) = \operatorname{argmin}_{\mathbf{z}^0} \|\delta_\epsilon\|_{\Sigma_\epsilon^{-1}}^2 + \|\delta_g\|_{\Sigma_g^{-1}(\mathbf{g})}^2 \quad (11)$$

with the closed-form solution:

$$\hat{\mathbf{z}}(\mathbf{g}) = (\Phi^\top \Sigma_\epsilon^{-1} \Phi + \Sigma_g^{-1}(\mathbf{g}))^{-1} (\Phi^\top \Sigma_\epsilon^{-1} \zeta + \Sigma_g^{-1}(\mathbf{g}) H \mathbf{g}). \quad (12)$$

Similarly, (10) also admits a closed-form solution under Gaussian uncertainties (Yin et al., 2023).

3.1 Hyperparameter Estimation

To apply the MAP estimator (9) or (10), the hyperparameter $\mathbf{g} \in \mathbb{R}^M$ needs to be estimated as well. In this regard, we adopt the maximum marginal likelihood method. Consider the equation $\zeta = \Phi H \mathbf{g} - \Phi \Delta_H \mathbf{g} + \epsilon$. The hyperparameter \mathbf{g} is estimated by maximizing the marginal likelihood: $\hat{\mathbf{g}} = \operatorname{argmax}_{\mathbf{g}} p(\zeta|\mathbf{g})$, where $\zeta|\mathbf{g} \sim \mathcal{E}_f(\Phi H \mathbf{g}, \Phi \Sigma_g(\mathbf{g}) \Phi^\top + \Sigma_\epsilon)$.

Let $\Psi(\mathbf{g}) := \Phi \Sigma_g(\mathbf{g}) \Phi^\top + \Sigma_\epsilon$ and $\delta_\zeta(\mathbf{g}) = \zeta - \Phi H \mathbf{g}$. If there exists \mathbf{g} such that $\Psi(\mathbf{g})$ is positive definite, the hyperparameter estimation problem is given by

$$\hat{\mathbf{g}} = \operatorname{argmax}_{\mathbf{g}} \det(\Psi(\mathbf{g}))^{-\frac{1}{2}} f\left(\|\delta_\zeta(\mathbf{g})\|_{\Psi^{-1}(\mathbf{g})}^2\right). \quad (13)$$

When $\Psi(\mathbf{g})$ is singular for all \mathbf{g} , similar to (10), let $(U_{\Psi,1}, U_{\Psi,2}, \Psi_1(\mathbf{g}))$ be the decomposition of $\Psi(\mathbf{g})$ according to (2). We have

$$\begin{aligned} \hat{\mathbf{g}} = \operatorname{argmax}_{\mathbf{g}} & \det(\Psi_1(\mathbf{g}))^{-\frac{1}{2}} f\left(\|U_{\Psi,1}^\top \delta_\zeta(\mathbf{g})\|_{\Psi_1^{-1}(\mathbf{g})}^2\right) \\ \text{s.t. } & U_{\Psi,2}^\top \delta_\zeta(\mathbf{g}) = \mathbf{0}. \end{aligned} \quad (14)$$

The estimated hyperparameter $\hat{\mathbf{g}}$ is used in (9) or (10) to obtain the MAP estimate. This assumes certainty equivalence without considering the hyperparameter estimation error, which is common for Bayesian methods with hyperparameters. The algorithm is summarized in Algorithm 1.

3.2 An Iterative Solution

Problems (13) and (14) are non-convex even for Gaussian uncertainties. This subsection derives an SQP algorithm to solve (13) and (14) under the following assumption.

Algorithm 1 MAP estimator for data-driven smoothing, prediction, and control

- 1: **Given:** $H, \zeta, \Sigma_\nu^d(\tau), \Sigma_\epsilon, \Phi, f(\cdot)$
 - 2: Solve (13) or (14) for $\hat{\mathbf{g}}$, where $\Sigma_g(\mathbf{g})$ is given by (7) or (8).
 - 3: Solve (9) or (10) for $\hat{\mathbf{z}}(\hat{\mathbf{g}})$.
 - 4: **Output:** $\hat{\mathbf{z}}(\hat{\mathbf{g}})$
-

Assumption 7. 1) The uncertainties are Gaussian, 2) the Page construction is used, and 3) $\Sigma_\nu^d(\tau) = \mathbf{0}_{n \times n}$ for all $\tau > 0$.

Assumption 7.1 specifies a particular $f(\cdot)$, whereas Assumptions 7.2 and 7.3 lead to a block diagonal $\Sigma_g(\mathbf{g})$. This iterative algorithm can serve as a computationally efficient approximate solution to the problem in practice, even when Assumption 7 is not satisfied. It also demonstrates the relations between the MAP framework and existing algorithms for data-driven prediction and control.

Under Assumption 7, (13) is equivalent to

$\hat{\mathbf{g}} = \operatorname{argmin}_{\mathbf{g}} J(\mathbf{g})$, $J(\mathbf{g}) := \log \det(\Psi(\mathbf{g})) + \|\delta_\zeta(\mathbf{g})\|_{\Psi^{-1}(\mathbf{g})}^2$, where $\Psi(\mathbf{g}) = \|\mathbf{g}\|_2^2 \Phi (\mathbb{I}_L \otimes \Sigma_\nu^d(0)) \Phi^\top + \Sigma_\epsilon$ from (8). Let $\Lambda := \Phi (\mathbb{I}_L \otimes \Sigma_\nu^d(0)) \Phi^\top$ and $\hat{\mathbf{g}}_k$ be the k -th iterate of the iterative algorithm. We approximate $J(\mathbf{g})$ locally at $\hat{\mathbf{g}}_k$ with a quadratic function as

$$\begin{aligned} J(\mathbf{g}) \approx & J(\hat{\mathbf{g}}_k) + \operatorname{tr}(\Lambda \Psi^{-1}(\hat{\mathbf{g}}_k)) \left(\|\mathbf{g}\|_2^2 - \|\hat{\mathbf{g}}_k\|_2^2 \right) \\ & - \|\Psi^{-1}(\hat{\mathbf{g}}_k) \delta_\zeta(\hat{\mathbf{g}}_k)\|_\Lambda^2 \left(\|\mathbf{g}\|_2^2 - \|\hat{\mathbf{g}}_k\|_2^2 \right) \\ & + \left(\|\delta_\zeta(\mathbf{g})\|_{\Psi^{-1}(\hat{\mathbf{g}}_k)}^2 - \|\delta_\zeta(\hat{\mathbf{g}}_k)\|_{\Psi^{-1}(\hat{\mathbf{g}}_k)}^2 \right). \end{aligned}$$

This leads to the SQP algorithm:

$$\hat{\mathbf{g}}_{k+1} = \operatorname{argmin}_{\mathbf{g}} c(\hat{\mathbf{g}}_k) \|\mathbf{g}\|_2^2 + \|\delta_\zeta(\mathbf{g})\|_{\Psi^{-1}(\hat{\mathbf{g}}_k)}^2, \quad (15)$$

where $c(\hat{\mathbf{g}}_k) := \operatorname{tr}(\Lambda \Psi^{-1}(\hat{\mathbf{g}}_k)) - \|\Psi^{-1}(\hat{\mathbf{g}}_k) \delta_\zeta(\hat{\mathbf{g}}_k)\|_\Lambda^2$.

The SQP algorithm for (14) can be obtained similarly:

$$\begin{aligned} \hat{\mathbf{g}}_{k+1} = \operatorname{argmin}_{\mathbf{g}} & \bar{c}(\hat{\mathbf{g}}_k) \|\mathbf{g}\|_2^2 + \|U_{\Psi,1}^\top \delta_\zeta(\mathbf{g})\|_{\Psi_1^{-1}(\hat{\mathbf{g}}_k)}^2 \\ \text{s.t. } & U_{\Psi,2}^\top \delta_\zeta(\mathbf{g}) = \mathbf{0}, \end{aligned} \quad (16)$$

where $\bar{c}(\hat{\mathbf{g}}_k) := \operatorname{tr}(\bar{\Lambda} \Psi_1^{-1}(\hat{\mathbf{g}}_k)) - \|\Psi_1^{-1}(\hat{\mathbf{g}}_k) U_{\Psi,1}^\top \delta_\zeta(\hat{\mathbf{g}}_k)\|_{\bar{\Lambda}}^2$ and $\bar{\Lambda} = U_{\Psi,1}^\top \Lambda U_{\Psi,1}$.

Problem (15) can be non-convex when the Hessian $c(\hat{\mathbf{g}}_k) \mathbb{I}_M + H^\top \Phi^\top \Psi^{-1}(\hat{\mathbf{g}}_k) \Phi H$ is not positive semidefinite, as commonly observed in SQP. Standard techniques in sequential convex programming can be applied to address this issue, such as projecting the Hessian onto the space of positive semidefinite matrices and utilizing the trust region or cubic regularization method (Nesterov and Polyak, 2006).

If closed-form solutions are desired, one can further approximate (15) or (16) by only running one iteration from the pseudoinverse solution to $\zeta = \Phi H \mathbf{g}$: $\hat{\mathbf{g}}_0 = \hat{\mathbf{g}}_{\text{pinv}} = H^\top \Phi^\top (\Phi H H^\top \Phi^\top)^{-1} \zeta$. In this case, $\delta_\zeta(\hat{\mathbf{g}}_0) = \mathbf{0}_p$ and $c(\hat{\mathbf{g}}_k), \bar{c}(\hat{\mathbf{g}}_k)$ are positive.¹ So, (15) and (16) are guar-

¹ Note that for $A, B \in \mathbb{S}_{++}^n$, there exists D such that $B = DD^\top$, and $\operatorname{tr}(AB) = \operatorname{tr}(ADD^\top) = \operatorname{tr}(D^\top AD) = \sum_{i=1}^n d_i^\top A d_i > 0$, where d_i is the i -th column of D .

anted to be convex. Similar to (12), convex quadratic programs with only equality conditions admit closed-form solutions, which we omit here for space constraints. A similar strategy is shown to perform well numerically for the prediction problem (Yin et al., 2021).

3.3 Relations to Existing Algorithms

The SQP algorithm is connected to existing algorithms in data-driven prediction and control under the following additional assumption.

Assumption 8. 1) $\Sigma_w(0) = \Sigma_w^d(0) = \mathbf{0}_{n_y \times n_u}$, 2) $\Sigma_\nu(\tau) = \mathbf{0}_{n \times n}$ for all $\tau > 0$, 3) $\Sigma_v(0) = \sigma^2 \mathbb{I}_{n_y}$, $\Sigma_v^d(0) = \sigma_d^2 \mathbb{I}_{n_y}$, and for the control task, 4) $Q = q \mathbb{I}_{n_y}$ and $R = r \mathbb{I}_{n_y}$.

Assumption 8.1 assumes no actuation uncertainties as in most existing algorithms, whereas Assumptions 8.2 to 8.4 leads to a diagonal $\Psi(\mathbf{g})$ and thus scalar weights in the objective function.

For the prediction task (T2), define $\mathbf{y}_p := U_{\Psi,1}^\top \boldsymbol{\zeta}$, $H_{yp} := U_{\Psi,1}^\top \Phi H$, $\mathbf{u} := U_{\Psi,2}^\top \boldsymbol{\zeta}$, and $H_u := U_{\Psi,2}^\top \Phi H$. The first iterate of (16) is equivalent to

$$\begin{aligned} \hat{\mathbf{g}}_1 = \operatorname{argmin}_{\mathbf{g}} \lambda \|\mathbf{g}\|_2^2 + \|\mathbf{y}_p - H_{yp}\mathbf{g}\|_2^2 \\ \text{s.t. } \mathbf{u} = H_u \mathbf{g}, \end{aligned} \quad (17)$$

where $\lambda := n_y L_0 \sigma_d^2$, when $\boldsymbol{\delta}_\zeta(\hat{\mathbf{g}}_0) = \mathbf{0}_p$. This problem coincides with the data-driven predictor presented in Lian et al. (2023), which is based on minimizing the Wasserstein distance between \mathbf{y}_p and $H_{yp}\mathbf{g}$. Note that by choosing different λ , (17) covers a wide range of data-driven predictors; see Yin et al. (2022) for an overview.

For the optimal control task (T3), define $\operatorname{col}(\mathbf{y}_p, \mathbf{y}^{\text{ref}}, \mathbf{u}^{\text{ref}}) := U_{\Psi,1}^\top \boldsymbol{\zeta}$, $\operatorname{col}(H_{yp}, H_{yf}, H_{uf}) := U_{\Psi,1}^\top \Phi H$, $\mathbf{u}_p := U_{\Psi,2}^\top \boldsymbol{\zeta}$, and $H_{up} := U_{\Psi,2}^\top \Phi H$. The first iterate of (16) is equivalent to

$$\begin{aligned} \hat{\mathbf{g}}_1 = \operatorname{argmin}_{\mathbf{g}} r \|\mathbf{u}^{\text{ref}} - H_{uf}\mathbf{g}\|_2^2 + \lambda_1 \|\mathbf{y}^{\text{ref}} - H_{yf}\mathbf{g}\|_2^2 \\ + \lambda_2 \|\mathbf{y}_p - H_{yp}\mathbf{g}\|_2^2 + \lambda_3 \|\mathbf{g}\|_2^2 \\ \text{s.t. } \mathbf{u}_p = H_{up}\mathbf{g}, \end{aligned} \quad (18)$$

where

$$\lambda_1 := \left(\sigma_d^2 \|\hat{\mathbf{g}}_0\|_2^2 + 1/q \right)^{-1}, \quad \lambda_2 := \left(\sigma_d^2 \|\hat{\mathbf{g}}_0\|_2^2 + \sigma^2 \right)^{-1}, \\ \lambda_3 := n_y \sigma_d^2 (L' \lambda_1 + L_0 \lambda_2),$$

when $\boldsymbol{\delta}_\zeta(\hat{\mathbf{g}}_0) = \mathbf{0}_p$. This problem has the same form as regularized data-enabled predictive control (DeePC) schemes; see Coulson et al. (2019); Berberich et al. (2021) and others. Selecting the regularization parameters in regularized DeePC is a non-trivial problem and significantly affects the performance (Dörfler et al., 2023). In this regard, (18) provides a meaningful method for choosing the regularization parameters without trial and error.

Remark 9. If we further assume no data uncertainties, i.e., $\sigma^2 = \sigma_d^2 = 0$, it is easy to see that (18) becomes the unregularized DeePC problem: $\min_{\mathbf{g}} r \|\mathbf{u}^{\text{ref}} - H_{uf}\mathbf{g}\|_2^2 + q \|\mathbf{y}^{\text{ref}} - H_{yf}\mathbf{g}\|_2^2$, s.t. $\mathbf{u}_p = H_{up}\mathbf{g}$, $\mathbf{y}_p = H_{yp}\mathbf{g}$.

Despite the relations shown above, the proposed unified framework differs from existing algorithms in the following ways. 1) Existing algorithms typically solely adopt the




Fig. 1. Performance comparison for data-driven smoothing, prediction, and control (lower is better). *Bayes*: Algorithm 1, *Approx*: approximate solution with one iteration (Section 3.2), *N4SID*: system identification with N4SID, *Proj*: two-step approach with projected signal matrices.

prior estimate $H\hat{\mathbf{g}}$ as the trajectory estimate, whereas an additional MAP step is performed in this work. 2) This work provides more general results compared to existing algorithms. Note that multiple conditions (Assumptions 7 and 8) are required to demonstrate resemblance to existing algorithms. 3) This work unifies data-driven prediction and direct data-driven predictive control under the same MAP formulation.

4. NUMERICAL EXAMPLES

This section presents three numerical examples covering data-driven smoothing, prediction, and control.² For each task, we compare the following methods: 1) *Bayes*: the proposed Bayesian framework (Algorithm 1) implemented with a nonlinear solver, 2) *Approx*: the approximate convex solution by implementing one SQP iteration as described in Section 3.2, 3) *N4SID*: the system identification paradigm where the model is identified with subspace identification by MATLAB command `n4sid`, and 4) *Proj*: the framework presented in Breschi et al. (2023) where H_{yf} is first projected and then deterministic behavioral systems theory is applied (Markovsky and Dörfler, 2022). For *N4SID*, we assume that the true model order is known and apply a Kalman filter on the identified state-space model to obtain state estimates. For *Proj*, we apply indirect data-driven control with the derived predictor.

For all tasks, 100 Monte Carlo simulations are conducted with the following parameters: $n_x = 10$, $n_u = n_y = 1$, $L = 40$, $L_0 = 10$, and $M = 61$. Signal matrices are built using the Hankel construction with 100 data points. Random systems are generated by MATLAB command `drss` unless otherwise stated. The systems are normalized to have an \mathcal{H}_2 -norm of 1. Non-convex optimization problems are solved by MATLAB command `fmincon`; convex optimization problems are solved by MOSEK.

(T1) Smoothing. We assume that there are no actuation uncertainties. The measurement uncertainties are i.i.d.

² The codes are available at <https://doi.org/10.25835/5f7y68pm>.

and subject to Student's t -distributions with $\xi = 10$, $\Sigma_v^d(0) = 10^{-4}$, and $\Sigma_v(0) = 10^{-2}$. Figure 1(left) shows the boxplot of root mean squared errors for the smoothed outputs. Results show that *Approx* and *Proj* outperform *Bayes* and *N4SID*. The relatively poor performance of *Bayes* may result from the local minimum issue with non-convex optimization.

(T2) Prediction. Both actuation and measurement uncertainties are i.i.d. Gaussian with $\Sigma_w^d(0) = \Sigma_v^d(0) = 10^{-4}$ and $\Sigma_w(0) = \Sigma_v(0) = 10^{-2}$. Figure 1(middle) shows the boxplot of root mean squared errors for the predicted outputs. The Bayesian framework performs better than *Proj* and similar to *N4SID*. The approximate solution yields nearly identical results to the non-convex solution.

(T3) Optimal control. We consider a discretized 1-D diffusion model specified as $x_{t+1}^1 = (1 - \beta)x_t^1 + \alpha(x_t^2 - x_t^1) + u_t$, $x_{t+1}^i = (1 - \beta)x_t^i + \alpha(x_t^{i+1} + x_t^{i-1} - 2x_t^i)$ for $i \in \mathbb{N}_{[2,9]}$, $x_{t+1}^{10} = (1 - \beta)x_t^{10} + \alpha(x_t^9 - x_t^{10})$, and $y_t = x_t^1$, where x_t^i denotes the i -th component of x_t and we select $\alpha = 0.4$, $\beta = 0.3$. We assume that there are no actuation uncertainties. The measurement uncertainties are Gaussian with $\Sigma_v^d(\tau) = \Sigma_v(\tau) = 10^{-2} \times 0.95^\tau$. The control objective is set as $Q = 5$, $R = 0.5$, $\mathbf{y}^{\text{ref}} = \text{col}(\mathbf{1}_{10}, -\mathbf{1}_{10}, \mathbf{1}_{10})$, and $\mathbf{u}^{\text{ref}} = \mathbf{y}^{\text{ref}}/d$, where d denotes the DC gain of the system. Figure 1(right) shows the boxplot of the normalized root control cost for the true trajectory:

$$J_{\text{ctr}} = \sqrt{\frac{1}{L'} \sum_{t=L_0+1}^L \left(\frac{R}{Q} \|u_t^0 - u_t^{\text{ref}}\|_2^2 + \|y_t^0 - y_t^{\text{ref}}\|_2^2 \right)}.$$

Here, *Bayes* achieves marginally better results than the other methods.

Overall, the numerical examples illustrate that the proposed Bayesian framework achieves performance comparable to or slightly better than that of existing algorithms across all three tasks. It is worth noting that this is achieved through a unified direct data-driven formulation, whereas *Proj* relies on an indirect control formulation.

5. CONCLUSION

This work proposes solving the same Bayesian estimation problem for data-driven smoothing, prediction, and control with stochastic data. This approach first solves a hyperparameter estimation problem, followed by a maximum a posteriori step that combines the knowledge from offline data and the online trajectory. It unifies these three data-driven tasks and generalizes them to accommodate input errors, correlated uncertainties, and elliptical distributions. Further work may investigate uncertainty quantification, computational aspects, and applications to moving horizon estimation and receding horizon control.

REFERENCES

Bänkestad, M.M., Sjölund, J., Taghia, J., and Schön, T.B. (2023). Variational elliptical processes. *Transactions on Machine Learning Research*.

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.

Chiuso, A., Fabris, M., Breschi, V., and Formentin, S. (2025). Harnessing uncertainty for a separation principle in direct data-driven predictive control. *Automatica*, 173, 112070.

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

Coulson, J., Lygeros, J., and Dörfler, F. (2022). Distributionally robust chance constrained data-enabled predictive control. *IEEE Transactions on Automatic Control*, 67(7), 3289–3304.

Damen, A., Van den Hof, P., and Hajdasinski, A. (1982). Approximate realization based upon an alternative to the Hankel matrix: the Page matrix. *Systems & Control Letters*, 2(4), 202–208.

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

Fiedler, F. and Lucia, S. (2021). On the relationship between data-enabled predictive control and subspace predictive control. In *European Control Conference (ECC)*, 222–229.

Hult, H. and Lindskog, F. (2002). Multivariate extremes, aggregation and dependence in elliptical distributions. *Advances in Applied Probability*, 34(3), 587–608.

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.

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.

Markovskiy, I. and Dörfler, F. (2022). Data-driven dynamic interpolation and approximation. *Automatica*, 135, 110008.

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

Nesterov, Y. and Polyak, B.T. (2006). Cubic regularization of Newton method and its global performance. *Mathematical Programming*, 108(1), 177–205.

Pan, G., Ou, R., and Faulwasser, T. (2023). On a stochastic fundamental lemma and its use for data-driven optimal control. *IEEE Transactions on Automatic Control*, 68(10), 5922–5937.

Smith, R.S., Abdalmoaty, M., and Yin, M. (2024). Data-driven formulation of the Kalman filter and its application to predictive control. In *IEEE Conference on Decision and Control (CDC)*, 2633–2639.

Van Overschee, P. and De Moor, B. (2012). *Subspace identification for linear systems: Theory, implementation, applications*. Springer, New York, NY.

Van Waarde, H.J., De Persis, C., Camlibel, M.K., and Tesi, P. (2020). Willems' fundamental lemma for state-space systems and its extension to multiple datasets. *IEEE Control Systems Letters*, 4(3), 602–607.

Willems, J.C., Rapisarda, P., Markovskiy, 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.

Yin, M., Iannelli, A., and Smith, R.S. (2024). Stochastic data-driven predictive control: Regularization, estimation, and constraint tightening. *IFAC-PapersOnLine*, 58(15), 79–84.