



HAL
open science

A bias-variance perspective of data-driven control

Kévin Colin, Yue Ju, Xavier Bombois, Cristian Rojas, Håkan Hjalmarsson

► **To cite this version:**

Kévin Colin, Yue Ju, Xavier Bombois, Cristian Rojas, Håkan Hjalmarsson. A bias-variance perspective of data-driven control. 2023. hal-04338871

HAL Id: hal-04338871

<https://hal.science/hal-04338871>

Preprint submitted on 12 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A bias-variance perspective of data-driven control

Kévin Colin ^{*,**}, Yue Ju ^{*,**}, Xavier Bombois ^{***,****},
Cristian Rojas ^{**}, Håkan Hjalmarsson ^{*,**},

** Centre for Advanced Bio Production, KTH*

*** Division of Decision and Control Systems, KTH Royal Institute of Technology, Sweden*

**** Laboratoire Ampère, UMR CNRS 5005, Ecole Centrale de Lyon, Université de Lyon, France*

***** Centre National de la Recherche Scientifique (CNRS), France*

Abstract: Data-driven control, the task of designing a controller based process data, find application in a wide range of disciplines and the topic has been subject to intense research over more than half a decade. The main purpose of this contribution is to elucidate on the commonalities between data-driven control and parameter estimation. In particular, we discuss the bias-variance trade-off, i.e. rather than aiming for the optimal controller one should aim for a constrained version, that, for example, may be characterized by tunable parameters, corresponding to hyperparameters in parameter estimation. To keep technicalities at a minimum, still capturing the essential features of the problem, we consider a quadratic open loop control problem applied to a finite impulse response system. We consider the problem of minimizing the expected control cost and first show that, for our setting, model-based data-driven control outperforms other methods. We then make a systematic study of 3 different controller structures and two different (model-based) data-driven tuning techniques and illustrate their behaviours numerically.

Keywords: Data-driven control, Bayes control, Regularization and Kernel Methods

1. INTRODUCTION

A typical situation in control design is that physical knowledge of the system needs to be complemented with experimental data. Standard practice is to identify a model in the best way possible and then adjust the control objective to account for model errors (often in a crude way). Throughout the years a substantial body of research, including robust and adaptive control, has been built up on how to do this in a systematic way. Identification for control emerged in the end of the 1980s, first focusing on model-based control using models of restricted complexity, but also leading to methods for direct tuning of the controller from data [Hjalmarsson et al. 1998, Hjalmarsson 2005, Gevers 2005, Bazanella et al. 2023, Formentin et al. 2014].

Recently, there has been a new surge of interest in the topic. Spurred by advances in regularized parameter estimation [Pillonetto et al. 2022], model kernels accounting for the expected degradation of the control cost have been proposed [Formentin and Chiuso 2021]. Another approach to account for noisy data in the control design has been to use the average risk framework, leading to what could

be called Bayes control methods [Scampicchio et al. 2019, Formentin and Chiuso 2021, Ferizbegovic et al. 2021]. In data-driven predictive control, Willems' behavioural framework [Willems et al. 2005] is leveraged [Coulson et al. 2019, Berberich et al. 2021, van Waarde et al. 2022]. The control designs employ collected input-output trajectories in constraints together with regularization terms in the control objective. An attempt to unify these techniques can be found in Breschi et al. [2023]. We point to Carè et al. [2023] for an exposé of how regularization has been employed in data-driven control. The double special issue on data-driven control [Dörfler 2023a;b] contain many interesting contributions to the field.

A topic that has been discussed for almost half a century has been whether model based or direct control design is to be preferred, see, e.g., the recent editorials in [Dörfler 2023a;b]. However, one may notice that such a classification is hard to make precise. For example, in model reference control, the model may be parametrized directly in terms of the desired controller Hjalmarsson [2005], and very recently it has been shown that techniques based on the behavioural framework are closer to being model based than previously believed [Chiuso 2023]. The objective of this contribution is to switch focus to the bias/variance trade-off which controls the long-run performance of any data-driven control method. While many of the mentioned contributions deal with this issue, we believe its role and techniques that can be employed warrant to be highlighted and will help further developments in the field. This is a

* This work was supported by VINNOVA Competence Center AdBIOPRO, contract [2016-05181] and by the Swedish Research Council through the research environment NewLEADS (New Directions in Learning Dynamical Systems), contract [2016-06079].
E-mail addresses: Kévin Colin: kcolin@kth.se, Yue Ju: yuej@kth.se, Xavier Bombois: xavier.bombois@ec-lyon.fr, Cristian Rojas: crro@kth.se, Håkan Hjalmarsson: hjalmar@kth.se.

well known topic in parameter estimation Ljung [1999a], Pillonetto et al. [2022], and the main purpose of this paper is to connect closer these, to an outsider, disparate topics. The purpose is thus not to propose any new data-driven control method (although we use three controller structures that do not seem to have been considered before), but rather to showcase the key elements of the problem and the tools that exist. We have deliberately chosen a very simple open-loop quadratic control problem to propel our message. In similar spirit we consider the simplest possible data-generating mechanism, namely a finite-impulse response model subject to white Gaussian noise.

The paper unfolds by first describing the considered setting in Section 2, followed by the main section, Section 3, where the essential ingredients of data-driven control are discussed. Sections 4-5 outline the specifics of the resulting tuning methods, which are illustrated numerically in Section 8. In Section 9 we provide some concluding remarks.

Notation. The set of real-valued matrices of dimension $n \times m$ will be denoted $\mathbb{R}^{n \times m}$. We denote by j the complex number which satisfies $j^2 = -1$. For any complex number x , its conjugate is denoted by \bar{x} and its real part by $\Re(x)$. The identity matrix of dimension $n \times n$ will be denoted by \mathbf{I}_n . \mathbf{A}^\top denotes the transpose of any matrix \mathbf{A} . The notation $\mathbf{X} \sim \mathcal{N}(\mathbf{a}, \mathbf{B})$ denotes that X is a normally distributed random variable with mean vector \mathbf{a} and covariance matrix \mathbf{B} and the probability density function of X is denoted by $\mathcal{N}(x; \mathbf{a}, \mathbf{B})$.

2. CONSIDERED FRAMEWORK

The purpose of this paper is to showcase the parallels between data-driven control and parameter estimation. In order to present them and preserve clarity, we will consider a very basic feed-forward linear control problem.

2.1 System

Consider a single-input single-output (SISO) discrete-time linear time-invariant (LTI) system \mathcal{S} given by

$$y(t) = G(z)u(t) + e(t) \quad (1)$$

where $y(t) \in \mathbb{R}$, $u(t) \in \mathbb{R}$ and $e(t) \in \mathbb{R}$ are the output, input and measurement noise at the time instant t respectively. Moreover, we will assume that the noise e is zero-mean, white and Gaussian distributed with variance σ_e^2 . For simplicity, we will consider that $G(z)$ is a finite impulse response (FIR) of order n with impulse response $G(z) = \sum_{i=0}^n \theta_i z^{-i}$. We will introduce the following notation $G(z, \boldsymbol{\theta}) = \boldsymbol{\Gamma}(z)\boldsymbol{\theta}$, $\boldsymbol{\Gamma}(z) = [1, z^{-1}, \dots, z^{-n}]$ and $\boldsymbol{\theta} = [\theta_0, \theta_1, \dots, \theta_n]^\top$.

2.2 A simple feedforward control problem

Given a filtered white noise reference signal $r(t)$ with power spectrum density $\Phi_{rr}(\omega)$ which is non-zero at all frequencies $\omega \in]-\pi, \pi]$, we wish to compute the control input u which minimizes the following infinite horizon cost

$$\lim_{T \rightarrow +\infty} \frac{1}{T} \sum_{t=1}^T \mathbb{E}_r [(r(t) - G(z)u(t))^2 + \lambda u(t)^2]$$

where $\lambda \geq 0$ is an user-defined penalization coefficient. We will consider a feedforward control $u(t) = F(z)r(t)$

where $F(z)$ is the feedforward controller. Using Parseval's theorem and injecting the latter in the control cost, we get

$$\mathcal{J}(F, G) = \frac{1}{2\pi} \int_{-\pi}^{\pi} \mathcal{Q}_{F,G}(e^{j\omega}) \Phi_{rr}(\omega) d\omega$$

$$\mathcal{Q}_{F,G}(e^{j\omega}) = |1 - F(e^{j\omega})G(e^{j\omega})|^2 + \lambda |F(e^{j\omega})|^2. \quad (2)$$

Because the integrand is non-negative and $\Phi_{rr}(\omega) > 0$ at all ω , the optimal feedforward controller $F_G(z)$ is such that its frequency response minimizes $\mathcal{Q}_{F,G}$ frequency wise which implies

$$F_G(e^{j\omega}) = \frac{G^{-1}(e^{j\omega})}{1 + \lambda/|G(e^{j\omega})|^2} \quad (3)$$

which is to be interpreted as the frequency function of a stable, possibly non-causal, filter.

The twist of the problem is now that G is unknown. Instead we have N_{id} samples from (1) at our disposal. Next, we will introduce the identification framework.

2.3 The maximum likelihood estimate

The dataset will be denoted by \mathcal{D} and is given by

$$\mathcal{D} = \{(u(t), y(t)) \mid t = 1, \dots, N_{id}\},$$

where u is generated in open loop. For convenience, let $\mathbf{y} = [y(1) y(2) \dots y(N_{id})]^\top$ and $\boldsymbol{\Phi}$ be the lower triangular Toeplitz matrix whose first column is given by $[u(1), \dots, u(N_{id})]^\top$. Under the given assumptions, the maximum likelihood of $\boldsymbol{\theta}$ is given¹ by

$$\hat{\boldsymbol{\theta}} = (\boldsymbol{\Phi}^\top \boldsymbol{\Phi})^{-1} \boldsymbol{\Phi}^\top \mathbf{y}. \quad (4)$$

The white Gaussian noise assumption of the noise e implies that $\hat{\boldsymbol{\theta}} \sim \mathcal{N}(\boldsymbol{\theta}, \mathbf{P})$, where $\mathbf{P} = \sigma_e^2 (\boldsymbol{\Phi}^\top \boldsymbol{\Phi})^{-1}$.

3. DATA-DRIVEN CONTROL

3.1 Expected control cost and a bias/variance perspective

Consider now any data-driven feedforward controller $F_{\mathcal{D}}(z)$. This will be a random quantity due to the presence of noise in the collected data \mathcal{D} . The long-run performance of the controller, i.e. the average cost when the process of collecting data \mathcal{D} , computing $F_{\mathcal{D}}$, and applying the controller to the same system (1), is given by the expected control cost, defined by

$$\mathcal{V}(F_{\mathcal{D}}, G) = \mathbb{E}[\mathcal{J}(F_{\mathcal{D}}, G)] = \frac{1}{2\pi} \int_{-\pi}^{\pi} \mathbb{E}[\mathcal{Q}_{F_{\mathcal{D}},G}(e^{j\omega})] \Phi_{rr}(\omega) d\omega \quad (5)$$

where \mathbb{E} denotes expectation with respect to the data \mathcal{D} . Using that $\mathcal{Q}_{F,G}$ is quadratic in F , we have²

$$\mathbb{E}[\mathcal{Q}_{F_{\mathcal{D}},G}] - \mathcal{Q}_{F_G,G} = \underbrace{\mathcal{Q}_{\mathbb{E}[F_{\mathcal{D}}],G} - \mathcal{Q}_{F_G,G}}_{\text{Bias increase} \geq 0} + \underbrace{(|G|^2 + \lambda) \text{Var}[F_{\mathcal{D}}]}_{\text{Variance increase} \geq 0},$$

where $\text{Var}[F_{\mathcal{D}}] = \mathbb{E}[(F_{\mathcal{D}} - \mathbb{E}[F_{\mathcal{D}}])^2]$. Above, the bias increase and the variance increase terms represent the increase in the control cost due to the systematic error and the variability of the data driven controller, respectively. Exactly as in parameter estimation, the minimal control cost is obtained by balancing these two terms. There is thus a bias-variance trade-off also in data-driven control. The main point of this contribution is to showcase that the same techniques as used in parameter estimation can be used to this end.

¹ This implies that we know the order n of $G(z)$. The order can be verified with, e.g., residual tests Ljung [1999b].

² We drop the frequency argument for brevity and clarity.

3.2 First observations

The first observation to be made is that, under the given assumptions, $\hat{\boldsymbol{\theta}}$ in (4) is a sufficient statistic and that it follows from this that it is sufficient to, in the class of data-driven controllers, consider those that are function of $\hat{\boldsymbol{\theta}}$. This follows from the following result.

Theorem 1. Suppose that S is a sufficient statistic for \mathcal{D} . Let $F_{\mathcal{D}}$ be a data-driven controller and define $F_S = \mathbb{E}[F_{\mathcal{D}}|S]$. Then

$$\mathcal{V}(F_S, G) \leq \mathcal{V}(F_{\mathcal{D}}, G)$$

with equality if and only if $F_{\mathcal{D}} = F_S$ with probability 1.

Proof: This is essentially the Rao-Blackwell theorem (Theorem 1.7.8 Lehmann and Casella [1998]). Using that \mathcal{J} is strictly convex in F , Jensen's inequality gives

$$\mathcal{J}(F_S, G) \leq \mathbb{E}[\mathcal{J}(F_{\mathcal{D}}, G)|S]$$

and taking the expectation gives the result. \blacksquare

The interpretation of Theorem 1 is that it is sufficient to consider only model-based data-driven controllers. Therefore, in the following, we will denote by $\hat{G}(z) = \mathbf{\Gamma}(z)\hat{\boldsymbol{\theta}}$, and index a data-driven controller with subscript \hat{G} , e.g. $F_{\hat{G}}$, suggesting that in terms of data it is a function of $\hat{\boldsymbol{\theta}}$ only (albeit the relationship may be involved).

Our next observation is that there is no uniformly optimal data-driven controller. Paralleling the parameter estimation case, this is seen by taking $F_{\hat{G}} = F_{\tilde{G}}$ for a particular system \tilde{G} . This controller, which ignores any information in the data set \mathcal{D} , will have $\mathcal{V}(F_{\hat{G}}, \tilde{G}) = \mathcal{J}(F_{\tilde{G}}, \tilde{G})$, i.e. the minimum possible cost when the true system happens to be \tilde{G} , but will of course give poor performance when the true system G is far from \tilde{G} .

3.3 Tuning methods

So suppose that we have a parametrized family of data-driven controllers³ $F_{\hat{G}}(z, \eta)$, $\eta \in D_{\eta}$. So how can we tune η to leverage the bias-variance trade-off? Again drawing on the parameter estimation problem, one may use risk minimization techniques or take a Bayes approach.

Risk minimization. The idea in risk minimization is to minimize an estimate of $\mathcal{V}(F_{\hat{G}}(\eta), G)$. An obvious possibility is to use the ML-estimate $\mathcal{V}(F_{\hat{G}}(\eta), \hat{G})$, giving

$$\hat{\eta}_{\text{MLRT}} = \arg \min_{\eta \in D_{\eta}} \mathcal{V}(F_{\hat{G}}(\eta), \hat{G}),$$

which we will call ML-risk tuning, but also other estimates may be considered. For example the bias in the ML-estimate may be removed as in SURE [Stein 1981]. To apply risk minimization we need to derive an expression for $\mathcal{V}(F_{\hat{G}}(\eta), \hat{G})$. To this end, we will assume that the identification experiment is such that as the sample size N_{id} grows, $\hat{\boldsymbol{\theta}} \rightarrow \boldsymbol{\theta}$, and hence $F_{\hat{G}}(\eta) \rightarrow F_G(\eta)$. It is straightforward to show that

$$\mathbb{E}[\mathcal{Q}_{F_{\hat{G}}(\eta), G}] = 1 + \alpha_{\eta} + \beta_{\eta} + \bar{\beta}_{\eta}, \quad (6)$$

where

$$\alpha_{\eta} = (|G|^2 + \lambda) \frac{(|F_G(\eta)|^2 + \Delta_{G, \eta, V} + F_G(\eta)\overline{\Delta}_{F_{\hat{G}}(\eta), E} + \overline{F_G(\eta)}\Delta_{G, \eta, E})}{|G|^2 + \lambda}$$

³ For brevity of notation we will use the notation $F_{\hat{G}}(\eta)$ for $F_{\hat{G}}(\cdot, \eta)$.

$$\beta_{\eta} = -G(\Delta_{G, \eta, E} + F_G(\eta))$$

where

$$\Delta_{G, \eta, V} := \mathbb{E}_{\hat{G}}[|F_{\hat{G}}(e^{j\omega}, \eta) - F_G(e^{j\omega}, \eta)|^2]$$

$$\Delta_{G, \eta, E} := \mathbb{E}_{\hat{G}}[F_{\hat{G}}(\eta)] - F_G(\eta)$$

Both $\Delta_{G, \eta, V}$ and $\Delta_{G, \eta, E}$ depend on the parametrization of $F_{\hat{G}}(\eta)$. These quantities are in general not easy to compute. One way to handle this is to use sampling techniques. Another approach is to use Gauss' approximation formula. Below we show how this is done. Introduce

$$G(e^{j\omega}) = R + jI, \quad \hat{G}(e^{j\omega}) = \hat{R} + j\hat{I}$$

$$\delta_G = [\delta_R \ \delta_I]^{\top} := [\hat{R} - R \ \hat{I} - I]^{\top}$$

Then, a second order Taylor approximation of $F_{\hat{G}}(\eta) - F_G(\eta)$ with respect to $[\hat{R} \ \hat{I}]^{\top}$ is given by

$$F_{\hat{G}}(\eta) - F_G(\eta) \approx F'_G(\eta)\delta_G + \frac{1}{2}\delta_G^{\top}F''_G(\eta)\delta_G$$

giving

$$\Delta_{G, \eta, V} \approx F'_G(\eta) \mathbb{E}[\delta_G \delta_G^{\top}] \overline{F'_G(\eta)}^{\top} \quad (7)$$

$$\Delta_{G, \eta, E} \approx \frac{1}{2} \text{Tr}(F''_G(\eta) \mathbb{E}[\delta_G \delta_G^{\top}]) \quad (8)$$

where the vector $F'_G(\eta) \in \mathbb{C}^2$ and the matrix $F''_G(\eta) \in \mathbb{C}^{2 \times 2}$ are, respectively, the gradient and the Hessian matrix of $F_G(\eta)$ with respect to $[R \ I]^{\top}$. Now, from the properties of $\hat{\boldsymbol{\theta}}$ and the definition of \hat{G} it follows

$$\mathbb{E}[\delta_G(e^{j\omega})\delta_G^{\top}(e^{j\omega})] = \mathbf{\Lambda}(e^{j\omega})\bar{\mathbf{P}}\mathbf{\Lambda}(e^{j\omega})^{\top}$$

where

$$\mathbf{\Lambda}(e^{j\omega}) = \begin{bmatrix} 1 & \cos(\omega) & \dots & \cos(n\omega) \\ 0 & \sin(\omega) & \dots & \sin(n\omega) \end{bmatrix}$$

Using the approximations (7)-(8) in the expressions for α_{η} and β_{η} , gives an approximation of $\mathbb{E}[\mathcal{Q}_{F_{\hat{G}}(\eta), G}(e^{j\omega})]$, that we denote by $\tilde{\mathcal{Q}}_{F_{\hat{G}}(\eta), G}(e^{j\omega})$, and integrating this quantity with respect to ω gives an approximation $\tilde{\mathcal{V}}(F_{\hat{G}}(\eta), G)$ to $\mathcal{V}(F_{\hat{G}}(\eta), G)$. The minimizer to $\tilde{\mathcal{V}}(F_{\hat{G}}(\eta), \hat{G})$ gives a controller that approximates ML-risk tuning.

Bayes control. One may arrive at Bayes tuning by taking a Bayesian perspective as in [Scampicchio et al. 2019], or from a Bayes perspective [Ferizbegovic et al. 2021]. Both these contributions consider model reference control and here we will follow the latter exposé adapted to the control problem considered here.

Let us introduce a non-negative weighting function $p(\boldsymbol{\theta})$ encoding how important it is to obtain a good control performance should the system really be $\boldsymbol{\theta}$. We then form the weighted criterion

$$W(F_{\hat{G}}) = \int \mathcal{V}(F_{\hat{G}}, G(\boldsymbol{\theta}))p(\boldsymbol{\theta})d\boldsymbol{\theta}$$

With $p_{\boldsymbol{\theta}}(\mathcal{D})$ denoting the probability density function for the data when $\boldsymbol{\theta}$ corresponds to the true system, we can, after interchanging the order of the integrals, write

$$W(F_{\hat{G}}) = \frac{1}{2\pi} \int_{-\pi}^{\pi} \int \int \mathcal{Q}_{F_{\hat{G}}, G(\boldsymbol{\theta})}p_{\boldsymbol{\theta}}(\mathcal{D})p(\boldsymbol{\theta})d\boldsymbol{\theta} d\mathcal{D} d\omega$$

As the integrand is non-negative and the inner integral is independent of $\boldsymbol{\theta}$, the data-driven controller minimizing the weighted criterion W is given frequency-by-frequency by

$$F_{B, \mathcal{D}}(e^{j\omega}) := \arg \min_{F_{\hat{G}}} \tilde{\mathcal{Q}}_{F_{\hat{G}}, \mathcal{D}}(e^{j\omega}) \quad (9)$$

where $\tilde{Q}_{F,\mathcal{D}}(e^{j\omega}) = \int Q_{F,G(\boldsymbol{\theta})}(e^{j\omega}) p_{\boldsymbol{\theta}}(\mathcal{D}) p(\boldsymbol{\theta}) d\boldsymbol{\theta}$. For the case of a parametrized data-driven controller $F_{\hat{G}}(\eta)$, the minimizing η is given by

$$\eta_{\text{BT}}(\mathcal{D}) = \arg \min_{\eta \in D_{\eta}} \frac{1}{2\pi} \int_{-\pi}^{\pi} \tilde{Q}_{F_{\hat{G}}(\eta),\mathcal{D}}(e^{j\omega}) d\omega$$

We call this Bayes tuning (BT) of η .

3.4 Some novel leads for balancing bias vs variance

The certainty equivalence (CE) controller consists of replacing the unknown G in (3) by the ML-estimate \hat{G} .

For risk minimization, again paralleling the parameter estimation problem, a meaningful trade-off can only be achieved by keeping in rein the degrees of freedom in $F_{\hat{G}}$ as otherwise the variance term will dominate. To see this, consider (6) and let us assume that we let the parametrization of $F_{\hat{G}}(\eta)$ completely free, we call this parametrization *unstructured*, and that we use ML-risk tuning. Then taking CE-controller as $F_{\hat{G}}$ gives $\mathcal{V}(F_{\hat{G}}(\eta), \hat{G})$ equal to the minimum of $\mathcal{Q}_{F,\hat{G}}$, i.e. the minimum cost that can be achieved, implying that it is the CE controller that will be obtained in ML-risk tuning (MLRT). In Section 4 we will see that Bayes tuning (BT) is able to handle unstructured controller parametrizations in a meaningful way.

To illustrate the concepts of this paper we will use the following alternatives, which, in line with Theorem 1, all are model-based:

- shrunked certainty equivalence:

$$F_{sh}(e^{j\omega}, \eta) = \eta \frac{\hat{G}(e^{j\omega})^{-1}}{1 + \lambda/|\hat{G}(e^{j\omega})|^2} \quad (10)$$

- kernel-based certainty equivalence

$$F_{ke}(e^{j\omega}, \boldsymbol{\eta}) = \frac{\hat{G}(e^{j\omega}, \boldsymbol{\eta})^{-1}}{1 + \lambda/|\hat{G}(e^{j\omega}, \boldsymbol{\eta})|^2} \quad (11)$$

where $\hat{G}(e^{j\omega}, \boldsymbol{\eta})$ is a kernel estimate of G , employing, for example, the stable-spline (SS) kernel [Pillonetto et al. 2011] and $\boldsymbol{\eta}$ the vector gathering the kernel hyperparameters.

- input penalized certainty equivalence:

$$F_{ip}(e^{j\omega}, \eta(\omega)) = \frac{\hat{G}(e^{j\omega})^{-1}}{1 + \eta(\omega)/|\hat{G}(e^{j\omega})|^2} \quad (12)$$

where $\omega \rightarrow \eta(\omega)$ is a non-negative function to be tuned. We will explain later the intuition behind this choice of controller.

4. BAYES CONTROL FOR FIR MODELS

In this section we will derive the Bayes controller for (1). For this model, $p_{\boldsymbol{\theta}}(\mathcal{D}) = \mathcal{N}(\mathbf{y}; \boldsymbol{\Phi}\boldsymbol{\theta}, \sigma_e^2 \mathbf{I}_{N_{id}})$ and taking the weighting function as $p(\boldsymbol{\theta}) = \mathcal{N}(\boldsymbol{\theta}; 0, \mathbf{P}_{\boldsymbol{\theta}})$ gives

$$\tilde{Q}_{F,\mathcal{D}} = \int \mathcal{Q}_{F,G} \mathcal{N}(\mathbf{y}; \boldsymbol{\Phi}\boldsymbol{\theta}, \sigma_e^2 \mathbf{I}_{N_{id}}) \mathcal{N}(\boldsymbol{\theta}; 0, \mathbf{P}_{\boldsymbol{\theta}}) d\boldsymbol{\theta}$$

It is straightforward to show that

$$\tilde{Q}_{F,\mathcal{D}} = c(\mathbf{y}) \left(|1 - F\hat{G}|^2 + (\lambda + \Delta_G)|F|^2 \right) \Phi_{rr} \quad (13)$$

where $\hat{G} = \boldsymbol{\Gamma}\hat{\boldsymbol{\theta}}$ and $\Delta_G(e^{j\omega}) := \boldsymbol{\Gamma}(e^{j\omega})\tilde{\mathbf{P}}\boldsymbol{\Gamma}^*(e^{j\omega})$, where in turn $\tilde{\mathbf{P}} := (\mathbf{P}_{\boldsymbol{\theta}}^{-1} + \mathbf{P}^{-1})^{-1}$, where $\mathbf{P} = \sigma_e^2(\boldsymbol{\Phi}^T\boldsymbol{\Phi})^{-1}$, and

where $\hat{\boldsymbol{\theta}} := \tilde{\mathbf{P}}\boldsymbol{\Phi}^T\mathbf{y}/\sigma_e^2$ ⁴, and where $c(\mathbf{y})$ is a constant of no consequence for the considerations. As (13) has the same format as (2), save that the input penalty λ is replaced by $\lambda + \Delta_G(e^{j\omega})$. We can interpret this as that the input penalty has been made frequency dependent with higher penalty at frequencies where Δ_G is large. It also directly follows that the Bayes controller (9) is given by

$$F_{B,\mathcal{D}}(e^{j\omega}) = \frac{\hat{G}^{-1}(e^{j\omega})}{1 + (\lambda + \Delta_G(e^{j\omega}))/|G(e^{j\omega})|^2} \quad (14)$$

Thus, if W is taken as design criterion, the overall optimal data-driven controller can be interpreted as model based (using the model \hat{G}) obtained from a criterion where the input penalty is frequency dependent.

Before proceeding, we remark that the weighting $p(\boldsymbol{\theta})$ does not have to be integrable. In the setting above, we can take $p(\boldsymbol{\theta}) = c$ where c is an arbitrary positive constant. Redoing the calculations gives that the Bayes controller still is given by (14), but with $\hat{\boldsymbol{\theta}}$ being the ML-estimate (4) and $\tilde{\mathbf{P}} = \mathbf{P}$, implying that Δ_G is the variance of the least-squares estimate in the frequency domain. According to [Chiuso 2023], in a predictive control context this corresponds to a version of the DeePC-controller for certain settings of the hyperparameters in that algorithm.

In the following sections we will detail how the "hyperparameter" η in the different controllers in Section 3.4 are tuned when MLRT and BT are employed. However, in order to apply BT, we need to perform similar approximation for $W(F_{\hat{G}})$ as we did for the MLRT-method.

5. SHRINKAGE OF THE CE CONTROLLER

Let us consider the shrinkage strategy $F_{sh}(\eta)$ in (10) which can be rewritten as $F_{sh}(\eta) = \eta F_{\hat{G}}$ with $F_{\hat{G}}$ is the CE controller, i.e., the controller (3) where G is replaced by \hat{G} . In this case, we only have one hyperparameter η to be tuned which is a proportional constant of the CE controller. This strategy keeps the number of hyperparameters to be tuned minimal which is advantageous since the larger the number of hyperparameters to be tuned, the larger the variance. Let us discuss the tuning of η with both MLRT and BT approaches.

Evaluating the integrand $\mathbb{E}_{\mathcal{D}}[\mathcal{Q}_{F_{\mathcal{D}},\hat{G}}]$ of the risk function (see (5)) and the integrand $\tilde{Q}_{F_{\mathcal{D}},\hat{G}}$ of the Bayes approach (see (9)) at $F_{\mathcal{D}} = F_{sh}(\eta) = \eta F_{\hat{G}}$, we successively get

$$\begin{aligned} \mathbb{E}_{\mathcal{D}}[\mathcal{Q}_{F_{sh}(\eta),\hat{G}}] &= 1 - 2\eta\Re(\hat{G}\mathbb{E}_{\mathcal{D}}[F_{\hat{G}}]) + \eta^2(|\hat{G}|^2 + \lambda)\mathbb{E}_{\mathcal{D}}[|F_{\hat{G}}|^2] \\ \tilde{Q}_{F_{sh}(\eta),\hat{G}} &= 1 - 2\eta\Re(\hat{G}\mathbb{E}_{\mathcal{D}}[F_{\hat{G}}]) + \eta^2(|\hat{G}|^2 + \Delta_G + \lambda)\mathbb{E}_{\mathcal{D}}[|F_{\hat{G}}|^2] \end{aligned}$$

where $\Delta_G(e^{j\omega}) := \boldsymbol{\Gamma}(e^{j\omega})\tilde{\mathbf{P}}\boldsymbol{\Gamma}^*(e^{j\omega})$ with $\tilde{\mathbf{P}} = \sigma_e^2(\boldsymbol{\Phi}^T\boldsymbol{\Phi})^{-1}$ and $\boldsymbol{\Gamma}(z) = (1, z^{-1}, \dots, z^{-n})$. They are both convex quadratic function with respect to η for which the minimum is easy to compute. However, we face two problems: we have to be able to compute the integral over all the frequencies $[-\pi, \pi]$ which can be computationally expensive depending on the system complexity and both integrands depend on the moments $\mathbb{E}[F_{\hat{G}}]$ and $\mathbb{E}[|F_{\hat{G}}|^2]$ for which there is no closed-form expression. We can solve both problems by approximating the integral over a regular frequency gridding and both moments $\mathbb{E}[F_{\hat{G}}]$ and $\mathbb{E}[|F_{\hat{G}}|^2]$

⁴ In a Bayesian setting \hat{G} and $\tilde{\mathbf{P}}$ are the posterior mean and variance of the frequency function.

can be approximated by drawing several sample systems $G_{sample}(z) = \mathbf{\Gamma}(z)\boldsymbol{\theta}_{sample}$ with $\boldsymbol{\theta}_{sample}$ a sample drawn from $\mathcal{N}(\hat{\boldsymbol{\theta}}, \bar{\mathbf{P}})$.

6. APPLICATION-ORIENTED REGULARIZATION

Using kernel-based regularization is a common practice in parameter estimation especially when some prior knowledge on the system is available such as stability. In [Formentin and Chiuso 2021], the kernel is designed with the control objective in mind for the model reference control problem. A penalty term which corresponds to the control degradation performance cost is added to the identification problem in addition to a model kernel-regularization term. However, the hyperparameters of the kernel are kept fixed in their approach. Our approach differs from [Formentin and Chiuso 2021] on two accounts: (i) the structure of the controller is fixed and equal to the CE controller $F_{ke}(\boldsymbol{\eta})$ computed based on the regularized estimate $\hat{\boldsymbol{\theta}}(\boldsymbol{\eta})$ with, e.g., the SS kernel and (ii) we tune $\boldsymbol{\eta}$ by minimizing the risk or Bayes control cost with $\mathcal{F}_{\mathcal{D}} = F_{ke}(\boldsymbol{\eta})$, i.e., we take into account the direct effect of $\boldsymbol{\eta}$ on both control costs. Let us give more details on the implementation of the MLRT and BT of $\boldsymbol{\eta}$. Injecting $F_{ke}(\boldsymbol{\eta})$ (see (11)) into the integrand $\mathbb{E}_{\mathcal{D}}[\mathcal{Q}_{F_{ke}(\boldsymbol{\eta}), \hat{G}}]$ for the MLRT approach gives

$$\mathbb{E}_{\mathcal{D}}[\mathcal{Q}_{F_{ke}(\boldsymbol{\eta}), \hat{G}}] = 1 - 2\Re(\hat{G}\mathbb{E}_{\mathcal{D}}[F_{ke}(\boldsymbol{\eta})]) + (|\hat{G}|^2 + \lambda)\mathbb{E}_{\mathcal{D}}[|F_{ke}(\boldsymbol{\eta})|^2]$$

while the integrand $\tilde{\mathcal{Q}}_{\mathcal{F}_{\mathcal{D}}, \hat{G}}$ for the BT approach is

$$\tilde{\mathcal{Q}}_{F_{ke}(\boldsymbol{\eta}), \hat{G}} = 1 - 2\Re(\hat{G}\mathbb{E}_{\mathcal{D}}[F_{ke}(\boldsymbol{\eta})]) + (|\hat{G}|^2 + \Delta_G + \lambda)\mathbb{E}_{\mathcal{D}}[|F_{ke}(\boldsymbol{\eta})|^2]$$

However, we get the same issues as for the shrinkage. Hence, we should approximate the integral with a frequency gridding and both moments with a sampling approach. But, another issue is met: the optimization problems are both non-convex and we do not have a closed-form expression of the integrand with respect to $\boldsymbol{\eta}$. One can consider a gradient free optimization algorithms such as the simplex method in order to circumvent this issue. This renders this control design computationally demanding.

7. INPUT PENALIZED CE

The motivation for the input penalized CE-controller $F_{ip}(\boldsymbol{\eta})$ (12) stems from Bayes control, cf. (14). In fact, if the function η is determined using Bayes tuning, we will recover (14) corresponding to a constant weighting $p(\boldsymbol{\theta})$.

However, one may also tune the function η using risk-minimization. This will give another bias-variance trade-off. The approximation $\tilde{\mathcal{Q}}_{F_{\hat{G}}(\boldsymbol{\eta}), G}(e^{j\omega})$ (see the paragraph on risk minimization in Section 3.3) to $\mathbb{E}[\mathcal{Q}_{F_{\hat{G}}(\boldsymbol{\eta}), G}]$, see (6), in this case becomes a rational function in $\eta(\omega)$, having a derivative with a third order numerator. One can thus, frequency-by-frequency, obtain $\eta(\omega)$ by computing the solutions to this third order polynomial and taking the solution resulting in the smallest cost. We omit the details due to space limitations.

Remark 1. In this paper, we focused on the design of the frequency response of the proposed controllers. One can compute a filter which will approximate the obtained frequency responses using, e.g., the Yule-Walker approach so that the controllers can be implemented. However, it is important to check that the generated filters will stabilize the loop before. This can be done using robust control tools as described in, e.g., [Formentin and Chiuso 2021].

Table 1. Average control cost with different controllers.

	MLRT	BT
$F_{\hat{G}}$ (CE)	11.638	13.160
$F_{B, \mathcal{D}}$	7.845	5.007
Optimal	9.346	8.945

8. NUMERICAL EXAMPLE

For $G(z)$ in (1), we consider a truncation of the impulse response of the following system

$$\frac{0.5000 - 0.4150z^{-1} + 0.0816z^{-2}}{1 + 0.0200z^{-1} - 0.7395z^{-2}}$$

to obtain a FIR system whose first 35 impulse response coefficients (i.e., $n = 34$) are $\boldsymbol{\theta} = [\theta_0, \dots, \theta_{34}]^T$ and we set $G(z) = \sum_{i=0}^{34} \theta_i z^{-i}$. The noise variance is to $\sigma_e^2 = 0.1$. The reference output signal $r(t)$ considered for the feedforward control is a zero-mean white Gaussian noise filtered by $0.5/(1 - 0.5z^{-1})$.

For the identification, we generate $N_{id} = 500$ data with an input $u(t)$ which is a white Gaussian noise of variance 0.05. The average signal to noise ratio is -8 dB which corresponds to a highly noisy case. Hence, we could expect relatively poor control performance with the CE approach.

The design of the proposed methods will be done in the frequency domain as explained in the previous sections. However, the frequency responses are defined for all the frequencies in the continuous interval $]-\pi, \pi]$ and thus cannot be evaluated numerically. We will consider a regular gridding of the half interval $[0, \pi]$ with 500 gridding points since the control cost on the other interval is the same. All the control designs and their performance evaluation will be based on this grid. Finally, we will consider 100 different noise realizations in order to approximate $\mathcal{V}(F, G)$ in (5) with the average of the obtained costs.

In Table 1, we give the average control cost $\mathcal{V}(F, G)$ observed with the CE controller $F_{\hat{G}}$, the Bayes *unstructured* controller $F_{B, \mathcal{D}}$ (see (14)) and the three proposed parametrized controllers F_{sh} , F_{ke} and F_{ip} whose hyperparameters are tuned with both MLRT and BT. Firstly, all the proposed methods improve upon the CE controller to a varying degree. Secondly, we observe that which tuning method that performs the best depends on the controller. Risk-based tuning performs best for the shrinkage and input penalized controllers, while Bayes tuning is better for the kernel-based controller. Note that the cost for F_{ip} with the BT approach is the same as for $F_{B, \mathcal{D}}$, which was mentioned in Section 7.

In this example, application-oriented hyperparameter tuning strategies are the best ones, with a better performance obtained with the risk minimization tuning. As a comparison, we also simulated the CE controller regularized with the SS kernel for which both hyperparameters are tuned using the Empirical Bayes methods and we obtained an average cost of 11.659 which is twice worse than the application-oriented design with MLRT. This shows that taking the intended application of the identified model (here control) into account in the hyperparameter tuning may be a good practice as it improves the control performance.

9. CONCLUSIONS

In this contribution we have elucidated on the bias-variance trade-off that exists for all data-driven control design techniques. Drawing from parameter estimation, we have outlined two basic techniques that exists: Risk minimization and Bayes tuning. We have also, in what can be seen as a case study, studied how these techniques manifest themselves when applied to three different control structures. In regards to this case study, it is interesting to notice that Bayes controller (14) strongly connects to standard practice of detuning the control criterion in case of uncertainty. However, rather than the engineering solution of increasing the input penalty λ , a more elaborate scheme is employed where the penalty is made frequency dependent, with higher penalty at frequencies where the uncertainty is large. This suggests that research in this area may significantly impact engineering solutions of data-driven control.

We noted early on for our setting that, thanks to that $\hat{\theta}$ is a sufficient statistic, it was sufficient to study model-based data-driven control methods. Let us reassure ourselves that this is indeed the case. It is clear that all controllers (10)-(12), and (14), are employing the ML-estimate \hat{G} . Now, also η is tuned using data. However, as can be easily verified, the only way the data enters in the used tuning criteria is through \hat{G} , and thus all approaches are model-based as claimed. We remark that in the case of an ARX-model, the parameter estimate $\hat{\theta}$ and $\Phi^T \Phi$ form a sufficient statistic. The situation is thus slightly more complicated in this case but notice that the dimension of this statistic is invariant with respect to the sample size, and thus any data-driven controller can be based on this statistic without sacrificing statistical properties. Other model structures such as output-error, ARMAX, or Box-Jenkins do not have this property.

While we have employed the simple setting of single-input single-output open loop quadratic control, we believe that it is rather straightforward to generalize our findings to other, essentially, quadratic control problems, such as LQG and MPC.

REFERENCES

- Bazanella, A.S., Campestrini, L., and Eckhard, D. (2023). The data-driven approach to classical control theory. *Annual Reviews in Control*, 56, 100906.
- 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.
- Carè, A., Carli, R., Libera, A.D., Romeres, D., and Pillonetto, G. (2023). Kernel methods and gaussian processes for system identification and control: A road map on regularized kernel-based learning for control. *IEEE Control Systems Magazine*, 43(5), 69–110.
- Chiuso, A. (2023). ERNSI workshop 2023: Optimal data driven predictive control for linear stochastic systems. URL https://www.kth.se/polopoly_fs/1.1284895.1696833266!/ERNSI_2023_Alessandro_Chiuso.pdf.
- Coulson, J., Lygeros, J., and Dörfler, F. (2019). Regularized and distributionally robust data-enabled predictive control. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, 2696–2701. IEEE.
- Dörfler, F. (2023a). Data-driven control: Part one of two: A special issue sampling from a vast and dynamic landscape. *IEEE Control Systems Magazine*, 43(5), 24–27. doi:10.1109/MCS.2023.3291624.
- Dörfler, F. (2023b). Data-driven control: Part two of two: Hot take: Why not go with models? *IEEE Control Systems Magazine*, 43(6), 27–31. doi:10.1109/MCS.2023.3310302.
- Ferizbegovic, M., Mattsson, P., Schön, T., and Hjalmarsson, H. (2021). Bayes control of Hammerstein systems. In *19th IFAC Symposium on System Identification*. Padua, Italy.
- Formentin, S. and Chiuso, A. (2021). Control-oriented regularization for linear system identification. *Automatica*, 127, 109539.
- Formentin, S., Van Heusden, K., and Karimi, A. (2014). A comparison of model-based and data-driven controller tuning. *International Journal of Adaptive Control and Signal Processing*, 28(10), 882–897.
- Gevers, M. (2005). Identification for control: From the early achievements to the revival of experiment design. *European Journal of Control*, 11(4-5), 335–352. Semiplenary lecture at IEEE Conference on Decision and Control - European Control Conference.
- Hjalmarsson, H. (2005). From experiment design to closed loop control. *Automatica*, 41(3), 393–438.
- Hjalmarsson, H., Gevers, M., Gunnarsson, S., and Lequin, O. (1998). Iterative Feedback Tuning: theory and applications. *IEEE Control Systems Magazine*, 18(4), 26–41.
- Lehmann, E.L. and Casella, G. (1998). *Theory of Point Estimation*. John Wiley & Sons, New York, second edition.
- Ljung, L. (1999a). *System Identification: Theory for the User*. Prentice-Hall, Englewood Cliffs, NJ, 2nd edition.
- Ljung, L. (1999b). *System identification, Theory for the user*. System sciences series. Prentice Hall, Upper Saddle River, NJ, USA, second edition.
- Pillonetto, G., Chiuso, A., and Nicolao, G.D. (2011). Prediction error identification of linear systems: A nonparametric gaussian regression approach. *Automatica*, 47(2), 291–305.
- Pillonetto, G., Chen, T., Chiuso, A., De Nicolao, G., and Ljung, L. (2022). *Regularized System Identification: Learning Dynamic Models from Data*. Springer Nature.
- Scampicchio, A., Chiuso, A., Formentin, S., and Pillonetto, G. (2019). Bayesian kernel-based linear control design. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, 822–827.
- Stein, C. (1981). Estimation of the mean of a multivariate normal distribution. *Annals of Statistics*, 9(6), 1135–1151.
- van Waarde, H.J., Camlibel, M.K., and Mesbahi, M. (2022). From noisy data to feedback controllers: Non-conservative design via a matrix s-lemma. *IEEE Transactions on Automatic Control*, 67(1), 162–175.
- Willems, J.C., Rapisarda, P., Markovskiy, I., and De Moor, B.L. (2005). A note on persistency of excitation. *Systems & Control Letters*, 54(4), 325–329.