

Data-Driven System Level Synthesis

Anton Xue

University of Pennsylvania, Philadelphia, PA

ANTONXUE@SEAS.UPENN.EDU

Nikolai Matni

University of Pennsylvania, Philadelphia, PA

NMATNI@SEAS.UPENN.EDU

Abstract

We establish data-driven versions of the System Level Synthesis (SLS) parameterization of linear systems. In particular, optimization over achievable closed-loop finite-horizon system responses of linear-time-invariant dynamics can be posed using only data from past trajectories without explicitly identifying a system model. We first show an exact equivalence between the traditional and data-driven SLS parameterizations under the idealized assumption of noise-free trajectories. This is then extended to the case with process noise, where techniques from robust SLS are used to characterize and bound the effects of noise on closed-loop performance. Furthermore, we use tools from matrix concentration to show that simple trajectory averaging suffices to mitigate the impact of noise. We end with numerical experiments demonstrating the soundness of our methods.

Keywords: Data-driven control, model-free control, system level synthesis

Modern systems are increasingly dynamic and heterogeneous. These factors complicate modelling and by extension render impractical the techniques from traditional robust and optimal control, which assume the availability of accurate models. Fortunately, contemporary systems are inherently *data-rich*, meaning that data-driven control is possible, practical, and perhaps even necessary.

In the case where the underlying system is linear, approaches in data-driven control include identify-then-control [Dean et al. \(2019\)](#); [Mania et al. \(2019\)](#), policy gradient [Fazel et al. \(2018\)](#); [Malik et al. \(2019\)](#); [Furieri et al. \(2020\)](#), adaptive methods based on robust control [Dean et al. \(2018\)](#), and online-learning [Simchowitz et al. \(2020\)](#); [Hazan et al. \(2020\)](#). For a more extensive overview of recent developments please see [Matni et al. \(2019\)](#) and [Recht \(2019\)](#) for a control-theoretic and machine learning angle, respectively.

This work is motivated by results presented in [De Persis and Tesi \(2019\)](#); [Coulson et al. \(2019a, 2020, 2019b\)](#); [van Waarde et al. \(2020\)](#); [Rotulo et al. \(2019\)](#). Broadly, these paper leverage the behavioral framework of [Willems and Polderman \(1997\)](#) which, under the assumption of *persistence of excitation*, allows for the observed input/output data to exactly parameterize the achievable system trajectories. For example, in [Coulson et al. \(2019a\)](#) it is shown that trajectory tracking in output-feedback based model-predictive-control (MPC) [Garcia et al. \(1989\)](#); [Borrelli et al. \(2017\)](#) can be posed as an optimization problem over a library of past system trajectories; follow-up work establishes connections to distributionally robust programming [Coulson et al. \(2020\)](#) and allow for real-time implementations [Coulson et al. \(2019b\)](#). Similarly, in [Rotulo et al. \(2019\)](#); [De Persis and Tesi \(2019\)](#) it is shown that data-driven synthesis of linear quadratic regulators can be achieved through semi-definite programs without identifying an explicit system model. To the best of our knowledge, the effect of noise is not characterized in the aforementioned works.

Contributions In this paper we establish *data-driven* versions of the System Level Synthesis (SLS) [Anderson et al. \(2019\)](#) parameterization of achievable closed-loop system responses for LTI sys-

tems over finite horizon. SLS is central to breakthroughs in distributed optimal control [Wang et al. \(2019b\)](#), robust and distributed MPC [Amo Alonso and Matni \(2019\)](#); [Wang et al. \(2019a\)](#), and learning-enabled control [Dean et al. \(2019, 2018\)](#): our goal is to take an initial step in extending its advantages to the purely data-driven, model-free setting. We first consider the idealized setting of noise-free trajectories and show an exact equivalence between the traditional and our data-driven formulations of SLS. We then consider the case of systems driven by process noise, and use tools from robust SLS [Matni et al. \(2017\)](#) to characterize the behavior and bound the performance degradation of using noisy trajectory data for closed-loop control. Furthermore, we use matrix concentration [Tropp \(2012\)](#) to show that a simple trajectory averaging technique suffices to mitigate the impact of noise.

Paper structure We define the problem statement in Section 1. A data-driven SLS parameterization is then formulated in Section 2 for the case of LTI systems with no driving noise; this is extended to the noisy setting in Section 3. In Section 4 we derive sub-optimality bounds with norm-bounded assumptions on the noise and characterize the sample complexity. Section 5 contains numerical examples and we end in Section 6 with conclusions and discussion of future work.

Notation We use $x_{[i,j]}$ as shorthand for the signal $[x^\top(i) \ x^\top(i+1) \ \dots \ x^\top(j)]^\top$. Define

$$\mathcal{H}_L(\sigma_{[0,T-1]}) = \begin{bmatrix} \sigma(0) & \sigma(1) & \dots & \sigma(T-L) \\ \sigma(1) & \sigma(2) & \dots & \sigma(T-L+1) \\ \vdots & \vdots & \ddots & \vdots \\ \sigma(L-1) & \sigma(L) & \dots & \sigma(T-1) \end{bmatrix}, \quad \mathbf{R} = \begin{bmatrix} R^{0,0} & & & \\ & R^{1,1} & R^{1,0} & \\ & \vdots & \ddots & \ddots \\ R^{T-1,T-1} & \dots & R^{T-1,1} & R^{T-1,0} \end{bmatrix}$$

We say that the Hankel matrix $\mathcal{H}_L(\sigma_{[0,T-1]})$ is of order L , and when context is clear we overload notation and simply write $\mathcal{H}_L(\sigma)$. A linear, causal operator \mathbf{R} defined over a horizon of T has matrix representation, as shown above: here $R^{i,j} \in \mathbb{R}^{p \times q}$ is a matrix of compatible dimension. We denote the set of such linear causal operators by $\mathcal{L}_{TV}^{T,p \times q}$ and drop the superscript $T, p \times q$ when it is clear: then, an operator $\mathbf{R} \in \mathcal{L}_{TV}^{T,p \times q}$ acts on a signal $\sigma_{[0,T-1]}$ through multiplication, i.e., $y_{[0,T-1]} = \mathbf{R}\sigma_{[0,T-1]}$. We slightly overload notation, and use MATLAB-like syntax to extract block matrices, rows, and columns from linear operators: $\mathbf{R}(i, j)$, $\mathbf{R}(i, :)$, and $\mathbf{R}(:, j)$ denote the (i, j) -th block matrix, i -th block row, and j -th block column of \mathbf{R} , respectively, all indexing from 0.

1. Problem Statement

We consider finite-time optimal state-feedback control of the discrete-time LTI system

$$x(t+1) = Ax(t) + Bu(t) + w(t), \quad \text{for } t = 0, 1, \dots, L-1, \quad (1.1)$$

where $L > 0$ is the control horizon, $x \in \mathbb{R}^n$ is the system state, $u \in \mathbb{R}^m$ is the control input, and $w \in \mathbb{R}^n$ is the disturbance. We assume the pair (A, B) is controllable and by convention write the initial conditions of system (1.1) into the disturbance as $w(-1) = x(0)$.

When the system model (A, B) is known, the above problem can be efficiently solved for many cases of interest by making suitable assumptions on the noise signal $w_{[-1,L-2]}$ and control objective. This paper focuses instead on solving an optimal control problem when the model of system (1.1) is unknown, but a collection of state and input trajectories (over a longer horizon $T > L$ to be specified in Section 2) $\{x_{[0,T-1]}^{(i)}, u_{[0,T-1]}^{(i)}\}_{i=1}^N$ are available. Moreover, our goal is to solve this task *without explicitly estimating* the system model.

To make the discussion concrete, we focus on finite-horizon Linear Quadratic Gaussian (LQG) control, wherein the disturbances are assumed to be independently and identically distributed as

$w(t) \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2 I)$, the control policy at time t is given by a linear-time-varying function of past states, i.e., $u(t) = K_t(x_{[0,t]})$, and the cost function to be minimized is given by:

$$\mathbb{E} \left[\sum_{t=0}^{L-2} x^\top(t) Q x(t) + u^\top(t) R u(t) + x^\top(L-1) Q_F x(L-1) \right]. \quad (1.2)$$

We note that much of our analysis extends to other cost functions in a natural way.

2. Data Driven System Level Synthesis

We begin by considering the simplified setting in which there is no driving noise in system (1.1), i.e., $w(-1) = x(0)$ and $w(t) = 0$ for all $t \geq 0$. Our approach is to connect tools from behavioral control theory, namely Willems' Fundamental Lemma [Willems and Polderman \(1997\)](#); [Markovsky and Rapisarda \(2008\)](#), with the *System Level Synthesis* (SLS) [Anderson et al. \(2019\)](#) parameterization of closed-loop controllers.

2.1. Willems' Fundamental Lemma

Tools from behavioral system theory [Willems and Polderman \(1997\)](#); [Willems et al. \(2005\)](#); [De Persis and Tesi \(2019\)](#) provide a natural way of characterizing a dynamical system in terms of its input/output signals. Central to Willems' fundamental lemma is *persistence of excitation*, which is stated as a rank condition on a Hankel matrix constructed from the control input signal u .

Definition 1 *Let $\sigma : \mathbb{Z} \rightarrow \mathbb{R}^p$ be a signal. We say that its finite-horizon restriction $\sigma_{[0,T-1]}$ is persistently exciting (PE) of order L if the Hankel matrix $\mathcal{H}_L(\sigma_{[0,T-1]})$ has full rank.*

The rank condition implies that $T \geq (p+1)L - 1$ is a lower-bound on the horizon T . In what follows, we assume that the order L and data horizon T are chosen such that this bound is satisfied.

Lemma 2 (Willems et al. (2005); De Persis and Tesi (2019)) *Consider the system (1.1) with (A, B) controllable, and assume that there is no driving noise. Let $\{x_{[0,T-1]}, u_{[0,T-1]}\}$ be the state and input signals generated by the system. Then if $u_{[0,T-1]}$ is PE of order $n+L$, the signals $x_{[0,L-1]}^*$ and $u_{[0,L-1]}^*$ are valid trajectories L -length of system (1.1) if and only if*

$$\begin{bmatrix} x_{[0,L-1]}^* \\ u_{[0,L-1]}^* \end{bmatrix} = \begin{bmatrix} \mathcal{H}_L(x_{[0,T-1]}) \\ \mathcal{H}_L(u_{[0,T-1]}) \end{bmatrix} g, \quad \text{for some } g \in \mathbb{R}^{T-L+1}$$

Lemma 2 states that if the underlying system is controllable and $\text{rank } \mathcal{H}_{n+L}(u) = m(n+L)$, then: (1) all initial conditions and inputs are parameterizable from observed signal data; and (2) all valid trajectories, that is to say state/input trajectory pairs $\{x_{[0,L-1]}, u_{[0,L-1]}\}$ that are consistent with the dynamics (1.1) lie in the linear span of a suitable Hankel matrix constructed from the system trajectories. Our goal is to exploit this relationship to characterize valid *closed loop system responses* of the unknown system (1.1) by establishing a connection to the SLS parameterization.

2.2. System Level Synthesis

Consider an L -length trajectory from system (1.1) expressed as block matrix operations

$$x_{[0,L-1]} = \mathcal{Z}\mathcal{A}x_{[0,L-1]} + \mathcal{Z}\mathcal{B}u_{[0,L-1]} + w_{[-1,L-2]} \quad (2.1)$$

where $\mathcal{A} = I_L \otimes A$, and $\mathcal{B} = I_L \otimes B$, and where \mathcal{Z} is the block-downshift operator, i.e., a matrix with identity matrices along the first block subdiagonal and zeros elsewhere. If it is also the case

that the system (2.1) satisfies the linear feedback control law $u_{[0,L-1]} = \mathcal{K}x_{[0,L-1]}$ for a causal linear-time-varying state-feedback control policy $\mathcal{K} \in \mathcal{L}_{TV}^{L,m \times n}$, then rewriting (2.1) we arrive at

$$\begin{aligned} x_{[0,L-1]} &= (I - \mathcal{Z}(\mathcal{A} + \mathcal{B}\mathcal{K}))^{-1}w_{[-1,L-2]} = \Phi_x w_{[-1,L-2]} \\ u_{[0,L-1]} &= \mathcal{K}(I - \mathcal{Z}(\mathcal{A} + \mathcal{B}\mathcal{K}))^{-1}w_{[-1,L-2]} = \Phi_u w_{[-1,L-2]} \end{aligned} \quad (2.2)$$

which captures how the process noise w maps to the state x and control u . We refer to the causal linear operators $\Phi_x \in \mathcal{L}_{TV}^{L,n \times n}$ and $\Phi_u \in \mathcal{L}_{TV}^{L,m \times n}$ as the *system responses*, which characterize the closed-loop system behavior from noise to state and control input, respectively.

Theorem 3 (Theorem 2.1, Anderson et al. (2019)) *For a system (1.1) with state-feedback control law $\mathcal{K} \in \mathcal{L}_{TV}^{L,m \times n}$, i.e., $u_{[0,L-1]} = \mathcal{K}x_{[0,L-1]}$, the following are true*

1. *The affine subspace defined by*

$$[(I - \mathcal{Z}\mathcal{A}) \quad -\mathcal{Z}\mathcal{B}] \begin{bmatrix} \Phi_x \\ \Phi_u \end{bmatrix} = I, \quad \Phi_x \in \mathcal{L}_{TV}^{L,n \times n}, \quad \Phi_u \in \mathcal{L}_{TV}^{L,m \times n} \quad (2.3)$$

parameterizes all possible system responses from $w_{[-1,L-2]} \rightarrow (x_{[0,L-1]}, u_{[0,L-1]})$.

2. *For any causal linear operators Φ_x, Φ_u satisfying (2.3), the controller $\mathcal{K} = \Phi_u \Phi_x^{-1} \in \mathcal{L}_{TV}^{L,m \times n}$ achieves the desired closed-loop responses (2.2).*

Theorem 3 allows for the problem of controller synthesis to be equivalently posed as a search over the affine space of system responses characterized by constraint (2.3) by setting $x_{[0,L-1]} = \Phi_x w_{[-1,L-2]}$ and $u_{[0,L-1]} = \Phi_u w_{[-1,L-2]}$. In particular, the LQG problem in Section 1 can be recast as a search over system responses (see Section 2.2 of Anderson et al. (2019)) as:¹

$$\underset{\Phi_x, \Phi_u}{\text{minimize}} \quad \left\| \begin{bmatrix} \mathcal{Q}^{1/2} & \\ & \mathcal{R}^{1/2} \end{bmatrix} \begin{bmatrix} \Phi_x \\ \Phi_u \end{bmatrix} \right\|_F \quad \text{subject to (2.3)}, \quad (2.4)$$

where $\mathcal{Q} = I_L \otimes Q$, $\mathcal{R} = I_L \otimes R$, and $\|\cdot\|_F$ is the Frobenius norm. In the noise free setting, i.e., when the initial condition $w(-1) = x(0)$ is known and $w(t) = 0$ for $t \geq 0$, the objective function of this problem instead simplifies to

$$\underset{\Phi_x(:,0), \Phi_u(:,0)}{\text{minimize}} \quad \left\| \begin{bmatrix} \mathcal{Q}^{1/2} & \\ & \mathcal{R}^{1/2} \end{bmatrix} \begin{bmatrix} \Phi_x(:,0) \\ \Phi_u(:,0) \end{bmatrix} x(0) \right\|_F, \quad (2.5)$$

and similarly, because only the initial condition is nonzero, the affine constraint (2.3) reduces to

$$[(I - \mathcal{Z}\mathcal{A}) \quad -\mathcal{Z}\mathcal{B}] \begin{bmatrix} \Phi_x(:,0) \\ \Phi_u(:,0) \end{bmatrix} = I(:,0). \quad (2.6)$$

2.3. A Data-Driven Formulation

We now show how the simplified achievability constraints (2.6) can be replaced by a data-driven representation through the use of Lemma 2. Our key insight is to recognize that the i -th column of Φ_x and Φ_u are the impulse response of the state and control input, respectively, to the i -th disturbance channel — which are themselves valid system trajectories that can be characterized using Willems' fundamental lemma.

1. We drop both the squaring of the objective function, and the scaling factor σ^2 , for brevity of notation going forward, as neither affect the optimal solution, or the order-wise scaling of the derived bounds.

Theorem 4 Consider the system (1.1) with (A, B) controllable, and assume that there is no driving noise. Suppose that a state/input signal pair $\{x_{[0,T-1]}, u_{[0,T-1]}\}$ is collected, and assume that $u_{[0,T-1]}$ is PE of order at least $n + L$. We then have that the set of feasible solutions to constraint (2.6) defined over a time horizon $t = 0, 1, \dots, L - 1$ can be equivalently characterized as:

$$\begin{bmatrix} \mathcal{H}_L(x) \\ \mathcal{H}_L(u) \end{bmatrix} G, \quad \text{for all } G \in \Gamma(x) := \{G : \mathcal{H}_1(x)G = I\}. \quad (2.7)$$

Proof (Sketch)² Our goal is to prove the following relationship

$$\text{LHS} := \left\{ \begin{bmatrix} \Phi_x(:, 0) \\ \Phi_u(:, 0) \end{bmatrix} \text{ satisfying (2.6)} \right\} = \left\{ \begin{bmatrix} \mathcal{H}_L(x) \\ \mathcal{H}_L(u) \end{bmatrix} G : G \in \Gamma(x) \right\} =: \text{RHS}.$$

(\subseteq) Let $\{\Phi_x(:, 0), \Phi_u(:, 0)\} \in \text{LHS}$ and take e_1, \dots, e_n to be the standard basis vectors. By Lemma 2 there exists g_1, \dots, g_n such that each $\Phi_x(:, 0)e_i = \mathcal{H}_L(x)g_i$ and $\Phi_u(:, 0)e_i = \mathcal{H}_L(u)g_i$. Define G to be the horizontal concatenation of g_1, \dots, g_n .

(\supseteq) Substitute $\{\mathcal{H}_L(x)G, \mathcal{H}_L(u)G\} \in \text{RHS}$ into constraint (2.6) and note that $G \in \Gamma(x)$. \blacksquare

Thus, if a state/input pair $\{x_{[0,T-1]}, u_{[0,T-1]}\}$ is generated by a PE input signal of order at least $n + L$, Theorem 4 gives conditions under which $\{\mathcal{H}_L(x), \mathcal{H}_L(u)\}$ can be used to parameterize achievable system responses for system (1.1) under no driving noise. In particular, one can then reformulate the deterministic optimal control problem formulated in equations (2.5) and (2.6) as

$$\underset{G \in \Gamma(x)}{\text{minimize}} \quad \left\| \begin{bmatrix} \mathcal{Q}^{1/2} & \\ & \mathcal{R}^{1/2} \end{bmatrix} \begin{bmatrix} \mathcal{H}_L(x) \\ \mathcal{H}_L(u) \end{bmatrix} Gx(0) \right\|_F. \quad (2.8)$$

3. Robust-Data Driven System Level Synthesis

We now turn our attention to the original stochastic LQG optimal control problem (2.4), where for notational convenience we set $w(-1) = x(0) = 0$ and driving noise $w(t) \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2 I)$ for $t = 0, \dots, T - 2$. To differentiate between the state of the noise-free and noisy system, we will denote the state signal by $\tilde{x}_{[0,T-1]}$ when driving noise is present. This additional *unmeasurable input* means that valid system trajectories can no longer be solely characterized in terms of the Hankel matrices $\mathcal{H}_L(\tilde{x})$ and $\mathcal{H}_L(u)$ as the effect of the process noise, captured by a corresponding Hankel matrix $\mathcal{H}_L(w)$, must also be accounted for. To address this challenge, we relate the state-trajectories of system (1.1) under driving noise to those of system (1.1) under no driving noise, and use this relationship to construct *approximate system responses* that lie a bounded distance from the affine subspace defined in (2.3). We then leverage a robust SLS parameterization to bound the effects of this approximation error on the closed loop behavior.

3.1. Robust System Level Synthesis

We begin with a robust variant of Theorem 3 that characterizes the behavior achieved by a controller constructed from system responses lying near the affine subspace characterized by constraint (2.3).

Theorem 5 (Theorem 2.2, Anderson et al. (2019)) Let Δ be a strictly causal linear operator (i.e., its matrix representation is strictly block-lower-triangular), and suppose that $\{\hat{\Phi}_x, \hat{\Phi}_u\}$ satisfy

$$[(I - \mathcal{Z}\mathcal{A}) \quad -\mathcal{Z}\mathcal{B}] \begin{bmatrix} \hat{\Phi}_x \\ \hat{\Phi}_u \end{bmatrix} = I + \Delta, \quad \hat{\Phi}_x \in \mathcal{L}_{TV}^{L,n \times n}, \quad \hat{\Phi}_u \in \mathcal{L}_{TV}^{L,m \times n} \quad (3.1)$$

2. A full version is available at <https://arxiv.org/abs/2011.10674>

Then the controller $\hat{\mathcal{K}} = \hat{\Phi}_u \hat{\Phi}_x^{-1}$ achieves the system responses

$$\begin{bmatrix} x[0, L-1] \\ u[0, L-1] \end{bmatrix} = \begin{bmatrix} \hat{\Phi}_x \\ \hat{\Phi}_u \end{bmatrix} (I + \Delta)^{-1} w_{[-1, L-2]} \quad (3.2)$$

Equation (3.2) shows that the effect of the error term Δ in the approximate achievability constraint (3.1) is to map the original disturbance to $w_{[-1, L-2]} \rightarrow \tilde{w}_{[-1, L-2]} := (I + \Delta)^{-1} w_{[-1, L-2]}$. This makes clear that we must design the full system responses $\{\hat{\Phi}_x, \hat{\Phi}_u\}$, and not just their first block-columns as in the idealized setting considered in the previous section.

3.2. A Robust Data-Driven Formulation

We now construct causal linear operators $\{\hat{\Phi}_x, \hat{\Phi}_u\}$ from noisy data $\{\tilde{x}_{[0, T-1]}, u_{[0, T-1]}\}$ that satisfy equation (3.1). Our approach is to construct each block-column of the approximate system responses separately and then suitably concatenate them to yield a feasible solution. We emphasize that only $\{\tilde{x}_{[0, T-1]}, u_{[0, T-1]}\}$ is available, but it is instructive to also consider $w_{[-1, T-2]}$ in the analysis. To begin, since each column of $\mathcal{H}_L(\tilde{x}), \mathcal{H}_L(u), \mathcal{H}_L(w)$ satisfies (2.1), it follows that

$$[(I - \mathcal{Z}\mathcal{A}) \quad -\mathcal{Z}\mathcal{B}] \begin{bmatrix} \mathcal{H}_L(\tilde{x}) \\ \mathcal{H}_L(u) \end{bmatrix} = \begin{bmatrix} \mathcal{H}_1(\tilde{x}) \\ 0 \end{bmatrix} + \mathcal{Z}\mathcal{H}_L(w) \quad (3.3)$$

Then fix $\hat{G} \in \Gamma(\tilde{x})$ and let $\hat{\Phi}_x(:, 0) = \mathcal{H}_L(\tilde{x})\hat{G}$, $\hat{\Phi}_u(:, 0) = \mathcal{H}_L(u)\hat{G}$ as in the proof of Theorem 4,

$$[(I - \mathcal{Z}\mathcal{A}) \quad -\mathcal{Z}\mathcal{B}] \begin{bmatrix} \hat{\Phi}_x(:, 0) \\ \hat{\Phi}_u(:, 0) \end{bmatrix} = I(:, 0) + \underbrace{\mathcal{Z}\mathcal{H}_L(w)\hat{G}}_{\Delta(:, 0)} \quad (3.4)$$

If block down-shifting is accounted for, (3.4) demonstrates the construction of a single block-column of $\{\hat{\Phi}_x, \hat{\Phi}_u\}$ and Δ . The construction of the other columns is similar; in general consider $\hat{G}_0, \dots, \hat{G}_{L-1} \in \Gamma(\tilde{x})$, where each \hat{G}_{k-1} is used to construct the k th column of $\hat{\Phi}_x, \hat{\Phi}_u$. Note that \mathcal{Z} commutes with block-diagonal matrices with identical block-diagonal entries (adjusting for dimensions): this can be seen by observing that left-multiplication by \mathcal{Z} (down-shifting) is equivalent to right-multiplication by \mathcal{Z} (left-shifting). Thus, we can construct down-shifted block-columns of the form (3.4) as follows

$$[(I - \mathcal{Z}\mathcal{A}) \quad -\mathcal{Z}\mathcal{B}] \begin{bmatrix} \mathcal{Z}^{k-1}\mathcal{H}_L(x) \\ \mathcal{Z}^{k-1}\mathcal{H}_L(u) \end{bmatrix} \hat{G}_{k-1} = \mathcal{Z}^{k-1}I(:, 0) + \mathcal{Z}^k\mathcal{H}_L(w)\hat{G}_{k-1}, \quad (3.5)$$

from which full approximate system responses can be constructed, as formalized in the following.

Theorem 6 *For system (1.1) with (A, B) controllable, and control input $u_{[0, T-1]}$ and disturbance process $w_{[-1, T-2]}$ PE of order $n + L$, the approximate system response matrices $\{\hat{\Phi}_x, \hat{\Phi}_u\}$ and perturbation term Δ are defined as*

$$\hat{\Phi}_x = \mathcal{Z}_L(I_L \otimes \mathcal{H}_L(\tilde{x}))\hat{\mathcal{G}}, \quad \hat{\Phi}_u = \mathcal{Z}_L(I_L \otimes \mathcal{H}_L(u))\hat{\mathcal{G}}, \quad \Delta = \mathcal{Z}_L(I_L \otimes \mathcal{Z}\mathcal{H}_L(w))\hat{\mathcal{G}} \quad (3.6)$$

and satisfy the approximate achievability constraint (3.1), where $\mathcal{Z}_L := [I \quad \mathcal{Z} \quad \dots \quad \mathcal{Z}^{L-1}]$ and $\hat{\mathcal{G}} \in \mathcal{L}_{TV}$ with block-diagonal elements $\hat{G}(i, i) \in \Gamma(\tilde{x})$ for $i = 0, 1, \dots, L-1$, and off-diagonal blocks $\mathcal{H}_1(\tilde{x})\hat{G}(i, j) = 0$ for $i \neq j$.

Proof (Sketch) Repeated application of (3.5) to construct each block column of $\hat{\Phi}_x, \hat{\Phi}_u, \Delta$. ■

Theorem 6 thus allows us to apply the robust SLS parameterization of Theorem 5 to characterize the closed-loop behavior (3.2) achieved by a controller constructed from the data-driven approximate system responses $\{\hat{\Phi}_x, \hat{\Phi}_u\}$ in terms of the perturbation term Δ , as described in equation (3.6). In particular, if we assume that $\|\mathcal{H}(w)\|_2 \leq \varepsilon$ but is otherwise acting adversarially, we can pose the following robust LQG problem:

$$\begin{aligned} & \text{minimize}_{\hat{\mathcal{G}} \in \mathcal{L}_{TV}} \max_{\|\mathcal{H}_L(w)\|_2 \leq \varepsilon} \left\| \begin{bmatrix} \mathcal{Q}^{1/2} & \\ & \mathcal{R}^{1/2} \end{bmatrix} \begin{bmatrix} \mathcal{Z}_L(I_L \otimes \mathcal{H}_L(\tilde{x})) \\ \mathcal{Z}_L(I_L \otimes \mathcal{H}_L(u)) \end{bmatrix} \hat{\mathcal{G}} (I + \mathcal{Z}_L(I_L \otimes \mathcal{Z}\mathcal{H}_L(w))\hat{\mathcal{G}})^{-1} \right\|_F \\ & \text{subject to} \quad \hat{G}(i, i) \in \Gamma(\tilde{x}) \text{ for all } i, \quad \mathcal{H}_1(x)\hat{G}(i, j) = 0 \text{ for all } i \neq j. \end{aligned} \quad (3.7)$$

Note that although we assume that the disturbance process is drawn as $w(t) \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2 I)$, we conservatively treat the effects of the unknown Hankel matrix $\mathcal{H}_L(w)$ on the estimated system responses as adversarial in our analysis. This approach also allows our method to generalize naturally to other optimal control settings, such as those with \mathcal{H}_∞ and \mathcal{L}_1 cost functions.

The objective function of optimization problem (3.7) is non-convex, but its structure allows for a transparent and data-independent quasi-convex upper-bound to be derived. First, we observe that we can upper bound $\|\Delta\|_2$ as given in equation (3.6), by

$$\|\Delta\|_2 = \|\mathcal{Z}_L(I_L \otimes \mathcal{Z}\mathcal{H}_L(w))\hat{\mathcal{G}}\|_2 \leq \|\mathcal{Z}_L\|_2 \|\mathcal{H}_L(w)\|_2 \|\hat{\mathcal{G}}\|_2 \leq \sqrt{L}\varepsilon \|\hat{\mathcal{G}}\|_2, \quad (3.8)$$

from which it follows immediately that if $\sqrt{L}\varepsilon \|\hat{\mathcal{G}}\|_2 < 1$, then $\|(I + \Delta)^{-1}\|_2 \leq \frac{1}{1 - \sqrt{L}\varepsilon \|\hat{\mathcal{G}}\|_2}$. This observation allows us to follow a similar argument as in Dean et al. (2019) to derive the following quasi-convex upper bound to problem (3.7):

$$\begin{aligned} & \text{minimize}_{\gamma \in [0, 1], \hat{\mathcal{G}}} \quad \frac{1}{1 - \gamma} \left\| \begin{bmatrix} \mathcal{Q}^{1/2} & \\ & \mathcal{R}^{1/2} \end{bmatrix} \begin{bmatrix} \mathcal{Z}_L(I_L \otimes \mathcal{H}_L(\tilde{x})) \\ \mathcal{Z}_L(I_L \otimes \mathcal{H}_L(u)) \end{bmatrix} \hat{\mathcal{G}} \right\|_F \\ & \text{subject to} \quad \|\hat{\mathcal{G}}\|_2 \leq \frac{\gamma}{\sqrt{L}\varepsilon}, \quad \hat{G}(i, i) \in \Gamma(\tilde{x}) \text{ for all } i, \quad \mathcal{H}_1(\tilde{x})\hat{G}(i, j) = 0 \text{ for all } i \neq j \end{aligned} \quad (3.9)$$

which is quasi-convex in $(\gamma, \hat{\mathcal{G}})$, allowing for an efficient solution via bisection.

4. Sub-optimality Analysis

In this section we prove the following sub-optimality result, which relates the performance \hat{J} achieved by the controller synthesized via the robust problem (3.9) to the optimal J^* achieved by the optimal LQG controller. To state the our main result, let G_k^* be the optimal solution to the $L - 1 - k$ (indexing starts at zero) horizon LQG problem, as in Theorem 4. Further, let:

$$\mathcal{O}_L(A) := \begin{bmatrix} I \\ A \\ \vdots \\ A^{L-1} \end{bmatrix}, \quad \mathcal{T}_L(X) := \begin{bmatrix} 0 & & & & \\ X & 0 & & & \\ AX & X & \ddots & & \\ \vdots & \vdots & \ddots & \ddots & \\ A^{L-2} & A^{L-3} & \cdots & X & 0 \end{bmatrix}.$$

Theorem 7 *Let $(\hat{\mathcal{G}}, \hat{\gamma})$ be the optimal solution to (3.9), let \hat{J} be the LQG cost that the controller $\hat{\mathcal{K}} = \hat{\Phi}_u \hat{\Phi}_x^{-1}$ constructed from the system responses (3.6) achieves on system (1.1). Assume that $T \geq 2L + 1$, that $\|\mathcal{H}_L(w)\|_2 \leq \varepsilon$, and that ε satisfies the bounds (4.1). Let $\mathcal{G}^* = \text{blkdiag}(G_0^*, \dots, G_{L-1}^*)$, with $G_0^*, \dots, G_{L-1}^* \in \Gamma(x)$ the parameters to the optimal LQG system responses as defined above, and let $\{\Phi_x^*, \Phi_u^*\} = \{\mathcal{Z}_L(I_L \otimes \mathcal{H}_L(x))\mathcal{G}^*, \mathcal{Z}_L(I_L \otimes \mathcal{H}_L(u))\mathcal{G}^*\}$. Letting*

J^* be the optimal LQG cost achieved by the resulting optimal controller $\mathcal{K}_* = \Phi_u^*(\Phi_x^*)^{-1}$, then

$$\frac{\hat{J} - J^*}{J^*} \leq 6\|\mathcal{G}^*\|_2 \varepsilon \left(2\sqrt{L} + \|\mathcal{T}_{T-L+1}(I)\|_2 + \frac{L(1 + \|\mathcal{O}_L(A)\|_2)\|Q^{1/2}\|_F\|\mathcal{T}_{T-L+1}(I)\|_2}{J^*} \right)$$

Our strategy is to construct a feasible solution to problem (3.9) using the optimal \mathcal{G}^* defined in the theorem statement, such that $\{\Phi_x^*, \Phi_u^*\} = \{\mathcal{Z}_L(I_L \otimes \mathcal{H}_L(x))\mathcal{G}^*, \mathcal{Z}_L(I_L \otimes \mathcal{H}_L(u))\mathcal{G}^*\}$ for data $\{x_{[0,T-1]}, u_{[0,T-1]}\}$ generated by system (1.1) with no driving noise.³ First, we introduce the following lemma relating Hankel matrices of system (1.1) state trajectories with and without noise.

Lemma 8 *Let $x_{[0,T-1]}$ and $\tilde{x}_{[0,T-1]}$ be the state signals for system (1.1), driven by noise-free $u_{[0,T-1]}$ and noisy $\{u_{[0,T-1]}, w_{[-1,T-2]}\}$, respectively. Then the following holds*

$$\mathcal{H}_L(\tilde{x}) = \mathcal{H}_L(x) + \mathcal{T}_L(I)\mathcal{H}_L(w) + \mathcal{O}_L(A)\mathcal{W}_{[0,T-L]},$$

where $\mathcal{W}(t) = \sum_{k=0}^{t-1} A^{t-1-k}w(k)$ are columns of $\mathcal{W}_{[0,T-L]} = [\mathcal{W}(0) \ \cdots \ \mathcal{W}(T-L)]$.⁴

Proof (Sketch) Let $x_{[t,t-L+1]}$ and $\tilde{x}_{[t,t+L-1]}$ be any pair of columns of $\mathcal{H}_L(x)$ and $\mathcal{H}_L(\tilde{x})$. Then,

$$\begin{aligned} x_{[t,t+L-1]} &= \mathcal{O}_L(A)x(t) + \mathcal{T}_L(B)u_{[t,t+L-1]} \\ \tilde{x}_{[t,t+L-1]} &= \mathcal{O}_L(A)\tilde{x}(t) + \mathcal{T}_L(B)u_{[t,t+L-1]} + \mathcal{T}_L(I)w_{[t,t+L-1]} \\ \tilde{x}_{[t,t+L-1]} &= \mathcal{O}_L(A)(x(t) + \mathcal{W}(t)) + \mathcal{T}_L(B)u_{[t,t+L-1]} + \mathcal{T}_L(I)w_{[t,t+L-1]} \end{aligned}$$

■

We now use Lemma 8 to construct a feasible solution to the robust optimization problem (3.9) using the optimal solution \mathcal{G}^* , which is subsequently used to prove the main result of this section.

Lemma 9 *Let $x_{[0,T-1]}$ and $\tilde{x}_{[0,T-1]}$ be the state signals for system (1.1), driven by noiseless $u_{[0,T-1]}$ and noisy $\{u_{[0,T-1]}, w_{[-1,T-2]}\}$, respectively, and suppose that $\{u_{[0,T-1]}, w_{[-1,T-2]}\}$ are PE of order $n+L$. Let $\mathcal{G}^* = \text{blkdiag}(G_0^*, \dots, G_{L-1}^*)$, with $G_0^*, \dots, G_{L-1}^* \in \Gamma(x)$, be the parameter to the optimal LQG system responses in Theorem 4. Then, if*

$$\varepsilon \leq \min \left\{ \frac{1}{3\sqrt{L}\|\mathcal{G}^*\|_2}, \frac{1}{2\|\mathcal{G}^*\|_2 \cdot \|\mathcal{T}_{T-L+1}(I)\|_2} \right\}, \quad (4.1)$$

the pair $\{\mathcal{G}^0 = \mathcal{G}^*(I + \mathcal{D})^{-1}, \ \gamma^0 = 2\varepsilon\|\mathcal{G}^*\|_2\sqrt{L}\}$ is a feasible solution to (3.9) where

$$\mathcal{D} = \text{blkdiag}(D_0, \dots, D_{L-1}), \quad D_k = \mathcal{W}_{[0,T-L]}G_k^*, \text{ for } 0 \leq k \leq L-1.$$

Proof (Sketch) These are sufficient for $(I + \mathcal{D})^{-1}$ to exist, and for each $\|D_k\|_2 \leq 1/2$. ■

Proof [Theorem 7] (Sketch) By (3.9), the approximate system response $\{\hat{\Phi}_x, \hat{\Phi}_u\} = \{\mathcal{Z}_L(I_L \otimes \mathcal{H}_L(\tilde{x}))\hat{\mathcal{G}}, \mathcal{Z}_L(I_L \otimes \mathcal{H}_L(u))\hat{\mathcal{G}}\}$ achieves cost \hat{J} on the true dynamics, which is bounded as

$$\hat{J} \leq \frac{1}{1 - \hat{\gamma}} \left\| \begin{bmatrix} \mathcal{Q}^{1/2} & \hat{\Phi}_x \\ \mathcal{R}^{1/2} & \hat{\Phi}_u \end{bmatrix} \right\|_F \leq \frac{1}{1 - \gamma^0} \left\| \begin{bmatrix} \mathcal{Q}^{1/2} & \mathcal{Z}_L(I_L \otimes \mathcal{H}_L(\tilde{x})) \\ \mathcal{R}^{1/2} & \mathcal{Z}_L(I_L \otimes \mathcal{H}_L(u)) \end{bmatrix} \mathcal{G}_0 \right\|_F \quad (4.2)$$

3. Such a \mathcal{G}^* exists by Theorem 4, and we select the minimum norm \mathcal{G}^* satisfying the desired relationship. Future work will seek explicit relationships between the norms of the system responses, data matrices, and \mathcal{G}^* .

4. We let $\mathcal{W}(0) = 0$ by convention.

where the first inequality follows from (3.7), and the second from the optimality of $\{\hat{\Phi}_x, \hat{\Phi}_u \hat{\gamma}\}$ and Lemma 9. Set $\alpha(T) = \|\mathcal{T}_{T-L+1}(I)\|_2 \|\mathcal{G}^*\|_2$, then apply Lemma 8 and the relations

$$\mathcal{G}_0 = \mathcal{G}^*(I + \mathcal{D})^{-1}, \quad \begin{bmatrix} \Phi_x^* \\ \Phi_u^* \end{bmatrix} = \begin{bmatrix} \mathcal{Z}_L(I_L \otimes \mathcal{H}_L(x)) \\ \mathcal{Z}_L(I_L \otimes \mathcal{H}_L(u)) \end{bmatrix} \mathcal{G}^*, \quad \|(I + \mathcal{D})^{-1}\|_2 \leq \frac{1}{1 - \alpha(T)\varepsilon}$$

and $\|\mathcal{T}_L\|_2 \leq \|\mathcal{T}_{T-L+1}\|_2$ to yield

$$\frac{J - J^*}{J^*} \leq \frac{\gamma^0 + \alpha(T)\varepsilon}{(1 - \gamma^0)(1 - \alpha(T)\varepsilon)} + \frac{L(1 + \|\mathcal{O}_L(A)\|_2) \|Q^{1/2}\|_F \|\mathcal{T}_{T-L+1}(I)\|_2 \|\mathcal{G}^*\|_2 \varepsilon}{(1 - \gamma^0)(1 - \alpha(T)\varepsilon) J^*}$$

Recall that $\gamma^0 = 2\varepsilon \|\mathcal{G}^*\|_2 \sqrt{L}$, and by the assumptions we have that $\alpha(T)\varepsilon \leq 1/2$ and $\gamma^0 \leq 2/3$:

$$\frac{\gamma^0 + \alpha(T)\varepsilon}{(1 - \gamma^0)(1 - \alpha(T)\varepsilon)} \leq 6\varepsilon(2\|\mathcal{G}^*\|_2 \sqrt{L} + \alpha(T)).$$

Similarly, the second term is bounded as $6L(1 + \|\mathcal{O}_L(A)\|_2) \|Q^{1/2}\|_F \|\mathcal{T}_{T-L+1}(I)\|_2 \|\mathcal{G}^*\|_2 \varepsilon / J^*$. ■

4.1. Sample Complexity

Simple *trajectory averaging* suffices to reduce noise. Let $\{\tilde{x}_{[0,T-1]}^{(i)}, u_{[0,T-1]}^{(i)}, w_{[-1,T-2]}^{(i)}\}_{i=1}^N$ be a sample of N trajectories from system (1.1). By linearity, the averages $\bar{x}_{[0,T-1]}, \bar{u}_{[0,T-1]}, \bar{w}_{[-1,T-2]}$ also form a valid trajectory. Since $\bar{w}(t) \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \frac{\sigma^2}{N} I_n)$, then by matrix concentration Tropp (2012):

Lemma 10 $\mathbb{P}[\|\mathcal{H}_L(\bar{w})\|_2 \geq t] \leq 2nT \exp(-\frac{t^2 N}{2\sigma^2 nT})$ for all $t \geq 0$.

As in Alpago et al. (2020), the control input is not averaged out by setting $u_{[0,T-1]}^{(i)} \equiv u_{[0,T-1]}^{(1)}$ for all i such that $\bar{u}(t) = u^{(1)}(t)$ for all t . We then obtain an end-to-end sample complexity by combining Lemma 10 and Theorem 7, and inverting the probability bound of Lemma 10:

Corollary 11 *If $N \geq 2\sigma^2 nT \log(2nT/\delta) \max\{9L\|\mathcal{G}^*\|_2^2, 4\|\mathcal{G}^*\|_2^2 \|\mathcal{T}_{T-L+1}(I)\|_2^2\}$, then with probability at least $1 - \delta$, we have that*

$$\frac{\hat{J} - J^*}{J^*} \leq 6\|\mathcal{G}^*\|_2 \sqrt{\frac{2\sigma^2 nT}{N} \log\left(\frac{2nT}{\delta}\right)} \left(2\sqrt{L} + \|\mathcal{T}_{T-L+1}(I)\|_2 \left(1 + \frac{L(1 + \|\mathcal{O}_L(A)\|_2) \|Q^{1/2}\|_F}{J^*}\right)\right)$$

5. Experiments

We present experiments on the system from Dean et al. (2019)

$$A = \begin{bmatrix} 1.01 & 0.01 & 0.00 \\ 0.01 & 1.01 & 0.01 \\ 0.00 & 0.01 & 1.01 \end{bmatrix}, \quad B = I, \quad \sigma^2 = 0.1, \quad Q = 10^{-3}I, \quad R = I,$$

which corresponds to a slightly unstable graph Laplacian system with input much more penalized than output. All experiments were done in Julia v1.3.1 using the JuMP v0.21.2 with MOSEK v9.2.9. We found it effective to use JuMP's `dual_optimizer` due to the semi-definite programs.⁵

Bootstrap Estimation of Noise A vanilla bootstrap Chernick et al. (2011) is used to empirically estimate confidence on $\|\mathcal{H}_L(\bar{w})\|_2$ for $L = 10, T = 45$ in terms of the number of trajectory samples: Fig. 1 shows the estimated 95-th percentile compared to true 95-th percentile over 1000 trials.

5. All code is open source and available at <https://github.com/unstable-zeros/data-driven-sls>.

Controller Performance in MPC Loop: We consider four types of unconstrained MPC controllers based on the following finite-horizon LTV feedback gains: \mathcal{K}^* the optimal LQG controller synthesized with noise-free data; \mathcal{K}_B and \mathcal{K}_T the robust controllers synthesized using the bootstrap value of ε_B and true $\varepsilon = \|\mathcal{H}_L(w)\|_2$ respectively in problem (3.9); and \mathcal{K}_N the naive controller is synthesized by dropping the robustness constraint in problem (3.9). For the selected values of N , random trajectories $\{\tilde{x}^{(i)}_{[0,T-1]}, u^{(i)}_{[0,T-1]}, w^{(i)}_{[-1,T-2]}\}$ of length $T = 45$ are generated and used to form the appropriate Hankel matrices $\mathcal{H}_L(\tilde{x}), \mathcal{H}_L(\bar{u}), \mathcal{H}_L(\bar{w})$ by averaging trajectories. We set $u^{(1)}(t) \stackrel{\text{iid}}{\sim} \mathcal{N}(0, I)$ and replay $u^{(1)}$ in all trials. Each finite-horizon controller is then synthesized with the running cost matrices (Q, R) specified above. To remove the effects of the terminal cost Q_L on stability and optimality, we set $Q_L = P^*$, for P^* the solution to the discrete algebraic Riccati equation for the infinite horizon LQG problem specified in terms of (A, B, Q, R) , thus ensuring that \mathcal{K}^* is both stabilizing and equal to the optimal infinite horizon LQG controller. An MPC loop is then implemented over a horizon of $H = 1000$ time-steps starting from an initial $x(0) = 0$ with $w(t) \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2 I)$ noise. For robust controllers, we constrain the $\hat{\mathcal{G}}$ to be block-diagonal, with block diagonals $\hat{\mathcal{G}}(i, i) \in \Gamma(\tilde{x})$ in a restricted setting of Theorem 6.

We evaluate 50 trials at each value of N , with empirically computed costs shown in Figs. 2. We omit the nominal controllers as they consistently fail to stabilize the system. While the robust controllers tend to have worse cost than the optimal controller (Fig. 2 (Left)), they achieve better disturbance rejection, as seen by the smaller norm of the states (Fig. 2 (Right)), at the expense of larger control effort (not shown).

6. Conclusion & Discussion

We defined and analyzed data-driven SLS parameterizations of controllers for LTI systems. We show that using noise-free trajectories gives an exact equivalence between traditional and data-driven SLS. We then consider the setting with noisy data, and use tools from robust SLS and matrix concentration to characterize and bound the effect of noise on closed-loop performance. Future work will look to extend these results to the infinite horizon, distributed & robust MPC, and output-feedback settings.

Acknowledgments

The authors would like to thank Karl Schmeckpeper and Alexandre Amice for helpful feedback and comments. This work is funded in part by NSF awards CPS-2038873 and CAREER award ECCS-2045834, and a Google Research Scholar award.

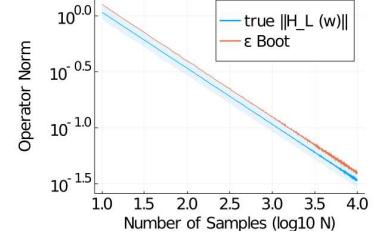


Fig. 1: Bootstrap Estimation of ε : In red is the 95-th percentile bound on the bootstrapped estimate of $\|\mathcal{H}_L(\bar{w})\|_2$ over 1000 trials. In blue are the median, 5-th, and 95-th percentiles of $\|\mathcal{H}_L(\bar{w})\|_2$ across an additional 1000 trials.

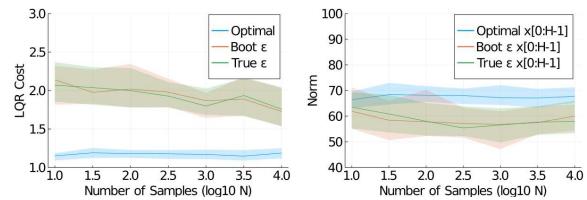


Fig. 2: (Left) Median and quartiles of MPC controller performances, with infeasibility assumed to be $+\infty$ in cost. (Right) Median and quartiles of running state trajectory norm $\|x_{[0:H-1]}\|_2$, with infeasibility assumed to be $+\infty$ in norm.

We show that using noise-free trajectories gives an exact equivalence between traditional and data-driven SLS. We then consider the setting with noisy data, and use tools from robust SLS and matrix concentration to characterize and bound the effect of noise on closed-loop performance. Future work will look to extend these results to the infinite horizon, distributed & robust MPC, and output-feedback settings.

References

Daniele Alpago, Florian Dörfler, and John Lygeros. An extended kalman filter for data-enabled predictive control. *IEEE Control Systems Letters*, 2020.

Carmen Amo Alonso and Nikolai Matni. Distributed and localized model predictive control via system level synthesis. *arXiv preprint arXiv:1909.10074*, 2019.

James Anderson, John C Doyle, Steven H Low, and Nikolai Matni. System level synthesis. *Annual Reviews in Control*, 2019.

Francesco Borrelli, Alberto Bemporad, and Manfred Morari. *Predictive control for linear and hybrid systems*. Cambridge University Press, 2017.

Michael R Chernick, Wenceslao González-Manteiga, Rosa M Crujeiras, and Erniel B Barrios. Bootstrap methods, 2011.

Jeremy Coulson, John Lygeros, and Florian Dörfler. Data-enabled predictive control: In the shallows of the deepc. In *2019 18th European Control Conference (ECC)*, pages 307–312. IEEE, 2019a.

Jeremy Coulson, John Lygeros, and Florian Dörfler. Regularized and distributionally robust data-enabled predictive control. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 2696–2701. IEEE, 2019b.

Jeremy Coulson, John Lygeros, and Florian Dörfler. Distributionally robust chance constrained data-enabled predictive control. *arXiv preprint arXiv:2006.01702*, 2020.

Claudio De Persis and Pietro Tesi. On persistency of excitation and formulas for data-driven control. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 873–878. IEEE, 2019.

Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. Regret bounds for robust adaptive control of the linear quadratic regulator. In *Advances in Neural Information Processing Systems*, pages 4188–4197, 2018.

Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. On the sample complexity of the linear quadratic regulator. *Foundations of Computational Mathematics*, pages 1–47, 2019.

Maryam Fazel, Rong Ge, Sham Kakade, and Mehran Mesbahi. Global convergence of policy gradient methods for the linear quadratic regulator. In *International Conference on Machine Learning*, pages 1467–1476. PMLR, 2018.

Luca Furieri, Yang Zheng, and Maryam Kamgarpour. Learning the globally optimal distributed lq regulator. In *Learning for Dynamics and Control*, pages 287–297. PMLR, 2020.

Carlos E Garcia, David M Prett, and Manfred Morari. Model predictive control: theory and practice—a survey. *Automatica*, 25(3):335–348, 1989.

Elad Hazan, Sham Kakade, and Karan Singh. The nonstochastic control problem. In *Algorithmic Learning Theory*, pages 408–421. PMLR, 2020.

Dhruv Malik, Ashwin Pananjady, Kush Bhatia, Koulik Khamaru, Peter Bartlett, and Martin Wainwright. Derivative-free methods for policy optimization: Guarantees for linear quadratic systems. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2916–2925. PMLR, 2019.

Horia Mania, Stephen Tu, and Benjamin Recht. Certainty equivalence is efficient for linear quadratic control. In *Advances in Neural Information Processing Systems*, pages 10154–10164, 2019.

Ivan Markovsky and Paolo Rapisarda. Data-driven simulation and control. *International Journal of Control*, 81(12):1946–1959, 2008.

Nikolai Matni, Yuh-Shyang Wang, and James Anderson. Scalable system level synthesis for virtually localizable systems. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 3473–3480. IEEE, 2017.

Nikolai Matni, Alexandre Proutiere, Anders Rantzer, and Stephen Tu. From self-tuning regulators to reinforcement learning and back again. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 3724–3740. IEEE, 2019.

Benjamin Recht. A tour of reinforcement learning: The view from continuous control. *Annual Review of Control, Robotics, and Autonomous Systems*, 2:253–279, 2019.

Monica Rotulo, Claudio De Persis, and Pietro Tesi. Data-driven linear quadratic regulation via semidefinite programming. *arXiv preprint arXiv:1911.07767*, 2019.

Max Simchowitz, Karan Singh, and Elad Hazan. Improper learning for non-stochastic control. *arXiv preprint arXiv:2001.09254*, 2020.

Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12(4):389–434, 2012.

Henk J van Waarde, Claudio De Persis, M Kanat Camlibel, and Pietro Tesi. Willems’ fundamental lemma for state-space systems and its extension to multiple datasets. *IEEE Control Systems Letters*, 4(3):602–607, 2020.

Han Wang, Shaoru Chen, Victor M Preciado, and Nikolai Matni. Robust model predictive control via system level synthesis. *arXiv*, pages arXiv–1911, 2019a.

Yuh-Shyang Wang, Nikolai Matni, and John C Doyle. A system-level approach to controller synthesis. *IEEE Transactions on Automatic Control*, 64(10):4079–4093, 2019b.

Jan C Willems, Paolo Rapisarda, Ivan Markovsky, and Bart LM De Moor. A note on persistency of excitation. *Systems & Control Letters*, 54(4):325–329, 2005.

J.C. Willems and J.W. Polderman. *Introduction to Mathematical Systems Theory: A Behavioral Approach*. Texts in Applied Mathematics. Springer New York, 1997. ISBN 9780387982663. URL <https://books.google.com/books?id=AbmTA1HbhKSC>.