

Cautious optimization via data informativity

Jaap Eising^a Jorge Cortés^a

^a*Department of Mechanical and Aerospace Engineering, University of California, San Diego, {jeising,cortes}@ucsd.edu*

Abstract

This paper deals with the problem of accurately determining guaranteed suboptimal values of an unknown cost function on the basis of noisy measurements. We consider a set-valued variant to regression where, instead of finding a best estimate of the cost function, we reason over all functions compatible with the measurements and apply robust methods explicitly in terms of the data. Our treatment provides data-based conditions under which closed-forms expressions of upper bounds of the unknown function can be obtained, and regularity properties like convexity and Lipschitzness can be established. These results allow us to provide tests for point- and set-wise verification of suboptimality, and tackle the cautious optimization of the unknown function in both one-shot and online scenarios. We showcase the versatility of the proposed methods in two control-relevant problems: data-driven contraction analysis of unknown nonlinear systems and suboptimal regulation with unknown dynamics and cost. Simulations illustrate our results.

Key words: Data-based optimization, set-membership identification, data-driven control, contraction analysis

1 Introduction

In many real-world applications involving optimization, the cost or reward function is not fully known and has to be ascertained from measurements. Such situations can arise due to a variety of factors, including system complexity, large-scale structure, or lack of access to the relevant information in unknown or adversarial scenarios. These considerations motivate the need for quantifiable performance guarantees that can be used for control and optimization in safety-critical applications. In this paper, we are therefore interested in cautious optimization of unknown functions, i.e., accurately determining guaranteed suboptimal values on the basis of noisy measurements. We consider both one-shot scenarios, where we are required to make decisions on the basis of a given set of data, and online scenarios, where the suboptimization of the unknown function can be repeatedly refined based on the collection of additional measurements. Our approach, based on data informativity, allows us to derive results which are robust against worst-case situations.

Literature review: Data-based optimization is a widely applicable and thoroughly investigated area. Giving a full overview of all methods and results is not possible, and thus we focus here on those most relevant to our treatment. The problems considered here are robust optimization programs, see e.g. [2, 3] and references therein. Solution approaches to data-based optimization problems usually start by using the data to obtain an estimate of the unknown function which in some sense ‘best’ explains the measurements. This requires the definition of the class of admissible functions. Popular nonparametric techniques are based

on Gaussian processes, as in [4], the approximation properties of neural networks [5, 6], and of particular relevance to our analysis here is the work on set-membership identification [7–9]. These works consider similar set-valued uncertainties, on the basis of Lipschitz interpolation.

In contrast, we make use of basis functions, where the unknown function is parameterized by a linear combination of the basis. Choices that balance expressiveness and the number of parameters are linear, polynomials, Gaussian, and sigmoidal functions. Depending on the problem, one may also choose a basis by taking into account physical considerations or employing spectral methods [10–12]. Once a basis is chosen, the data is used to perform regression [13, 14] on the parameters. What remains to be defined is then the metric chosen to decide the best estimate. Common methods are least squares (minimal Frobenius norm), ridge regression (minimal L_2 norm), sparse or Lasso regression (minimal L_1 norm). The nonlinear system-identification literature has also brought forward methods which determine models that are sparse [15], of low rank (via dynamic mode decomposition) [16–18], or both [19]. Regardless of how one obtains an estimate, the second step in data-based optimization is to employ either a certainty-equivalent or robust method to find optimal values.

The alternative approach we use here is in line with the literature of data-driven control methods on the basis of Willems’ fundamental lemma [20]. More specifically, we consider the concept of data informativity [21]. Simply put, these methods take the viewpoint that guaranteed conclusions can be drawn from data only if *all* systems compatible with the measurements have this property. Of specific interest to this paper are works that deal with nonlinear systems and a number of recent works have worked with systems given in terms of bilinear [22] or polynomial basis functions [23, 24], albeit in discrete time. Simultaneously,

* A preliminary version of this work was submitted as [1] to the IEEE Conference on Decision and Control. Corresponding author J. Eising.

there haven been efforts in generalizing Willems' lemma to continuous time [25, 26] and linking the results in discrete and continuous time by investigating sampling [27].

Our applications here concern data-based contraction analysis of nonlinear systems [28, 29] and the regulation of an unknown system to the suboptimal point of an unknown cost function. Data-based regulation with a known cost function has been investigated in a number of forms, including linear quadratic regulation, see [30] and references therein, and more general cost functions via optimization-based controllers [31, 32]. In [33], such controllers have been employed to optimize unknown cost functions. Finally, the online optimization scenario considered here is inspired by methods such as extremum-seeking control [34–36]. However, our implementation and performance guarantees are of a markedly different nature.

Statement of contributions: We tackle the data-based optimization of an unknown vector-valued function with the goal of obtaining quantifiable performance guarantees. Starting from the assumption that the unknown function is parameterized in terms of a finite number of basis functions, we consider *set-valued regression*, i.e., we reason with the set of *all* parameters compatible with the measurements for which the noise satisfies a bound given in terms of a quadratic matrix inequality. We take a worst-case approach regarding noise realizations, leading us to provide solutions to the following *cautious optimization* problems:

- (i) We provide conditions in terms of the measurements, basis functions, and noise model which guarantee that either the 2–norm or a linear functional of the unknown function is upper bounded by a given $\delta \in \mathbb{R}$ for a given $z \in \mathbb{R}^n$. Moreover, these results make it possible for us to find the smallest worst-case bounds;
- (ii) We also provide data-based conditions for convexity of either the true function or its worst-case norm bounds. These allow us to introduce a gradient-based method for finding the minimizer of the upper bound;
- (iii) We also identify data-based conditions for Lipschitzness which, using interpolation, allows us to derive upper bounds guaranteed to hold over compact sets.

We illustrate the versatility of the results in two application scenarios: contraction analysis of nonlinear systems (for which we provide data-based tests to determine one-sided Lipschitz constants and use them to guarantee contractivity) and suboptimal regulation of unknown plants with unknown cost functions (for which we employ convexity and Lipschitzness to find a fixed input such that the value of the cost function at the corresponding fixed point can be explicitly bounded above). Our final contribution is to online data-based optimization. We describe a framework for incorporating new measurements of the unknown function which includes online local descent, a method that iteratively collects measurements locally and then minimizes an upper bound of the form described above. We prove that under mild assumptions on the collection of data, the set of parameters consistent with all collected measurements shrinks. This allows us to conclude that the upper bounds found with online local descent converge to the true optimizer of the unknown function.

Preliminary results of this paper were presented in the conference article [1], whose focus was restricted to scalar functions and provided upper bounds and analyzed convexity

on the basis of data. The conference paper also presented a simplified version of online local descent for the scalar case. All of these are special cases of the present work. Note that the generalization here from scalar to vector-valued functions is instrumental in expanding the range of possibilities for application in analysis and control of the proposed data-based optimization framework.

Notation: Throughout the paper, we use the following notation. We denote by \mathbb{N} and \mathbb{R} the sets of nonnegative integer and real numbers, respectively. We let $\mathbb{R}^{n \times m}$ denote the space of $n \times m$ real matrices and \mathbb{S}^n the space of symmetric $n \times n$ matrices. For a vector $x \in \mathbb{R}^n$ we denote the standard p –norm by $\|x\|_p$. Given $R \in \mathbb{R}^{n \times n}$, we define the weighted p –norm by $\|x\|_{p,R} := \|Rx\|_p$. The i^{th} standard basis vector is denoted e_i . The induced matrix norm w.r.t. to the p –norm is denoted $\|A\|_p$ for $A \in \mathbb{R}^{n \times m}$. For vectors $v \in \mathbb{R}^n$, we write $v \geq 0$ (resp. $v > 0$) for elementwise non-negativity (resp. positivity). The sets of such vectors are denoted $\mathbb{R}_{\geq 0}^n := \{v \in \mathbb{R}^n | v \geq 0\}$ and $\mathbb{R}_{> 0}^n := \{v \in \mathbb{R}^n | v > 0\}$. For $P \in \mathbb{S}^n$, $P \geq 0$ (resp. $P > 0$) denotes that P is symmetric positive semi-definite (resp. definite). For a square matrix M , we denote $\lambda_{\min}(M)$ and $\lambda_{\max}(M)$ for the smallest and largest eigenvalue of M . We denote the smallest singular value of $M \in \mathbb{R}^{n \times m}$ by $\sigma_-(M)$ and its Moore-Penrose pseudoinverse by M^\dagger . The convex hull and interior of $\mathcal{S} \subseteq \mathbb{R}^n$ are denoted by $\text{conv}(\mathcal{S})$ and $\text{int}(\mathcal{S})$, resp. Given a symmetric matrix $M \in \mathbb{S}^{u+v}$, when u and v are clear from context, we partition it as:

$$M = \begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{bmatrix},$$

where $M_{11} \in \mathbb{S}^u$ and $M_{22} \in \mathbb{S}^v$. We denote the Schur complement of M with respect to M_{22} by $M|M_{22} := M_{11} - M_{12}M_{22}^\dagger M_{21}$. Lastly, a partitioned matrix gives rise to a quadratic matrix inequality (QMI), whose set of solutions is denoted by

$$\mathcal{Z}(M) := \left\{ Z \in \mathbb{R}^{v \times u} \mid \begin{bmatrix} I_u \\ Z \end{bmatrix}^\top M \begin{bmatrix} I_u \\ Z \end{bmatrix} \geq 0 \right\}.$$

2 Problem formulation

Consider a collection of known *basis functions* (or *features*), denoted $\phi_i : \mathbb{R}^n \rightarrow \mathbb{R}$ for $i = 1, \dots, k$. For ease of notation, we collect the basis functions into a vector-function as

$$b : \mathbb{R}^n \rightarrow \mathbb{R}^k, \quad b(z) := \begin{bmatrix} \phi_1(z)^\top & \dots & \phi_k(z)^\top \end{bmatrix}^\top. \quad (1)$$

On the basis of these basis functions, we define linear combinations as parameterized by $\theta \in \mathbb{R}^{k \times m}$:

$$\phi^\theta : \mathbb{R}^n \rightarrow \mathbb{R}^m, \quad \phi^\theta(z) = \theta^\top b(z).$$

Our starting point is an unknown function $\hat{\phi} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ that can be expressed in this form: i.e., there exists $\hat{\theta} \in \mathbb{R}^{k \times m}$ such that $\hat{\phi}(z) = \hat{\theta}^\top b(z)$.

Our goal is to investigate properties and optimize the unknown function $\hat{\phi}$ on the basis of measurements at points $\{z_i\}_{i=1}^T$. We assume the measurements are *noisy*, that is, we collect $y_i = \hat{\phi}(z_i) + w_i$, where w_i denotes an unknown

disturbance vector for each i . To express this in compact form, define the matrices $Y, W \in \mathbb{R}^{m \times T}$ and $\Phi \in \mathbb{R}^{k \times T}$ by

$$Y := \begin{bmatrix} y_1 & \cdots & y_T \end{bmatrix}, \quad W := \begin{bmatrix} w_1 & \cdots & w_T \end{bmatrix},$$

$$\Phi := \begin{bmatrix} b(z_1) & \cdots & b(z_T) \end{bmatrix} = \begin{bmatrix} \phi_1(z_1) & \cdots & \phi_1(z_T) \\ \vdots & & \vdots \\ \phi_k(z_1) & \cdots & \phi_k(z_T) \end{bmatrix}. \quad (2)$$

Then, for the true value of $\hat{\theta}$, we have $Y = \hat{\theta}^\top \Phi + W$. In this equation, the matrices Y and Φ are known, and $\hat{\theta}$ and W are unknown.

A common line of reasoning (see e.g., [13]) to approximate the unknown function $\hat{\phi}$ is the following. Assuming that small noise samples are more probable than large noise samples, we attempt to find a solution θ to $Y = \theta^\top \Phi + W$ for which the Frobenius norm of W , denoted $\|W\|_F$, is as small as possible. This can be formulated as

$$\arg \min_{\theta} \|Y - \theta^\top \Phi\|_F = (Y\Phi^\dagger)^\top.$$

Here instead we consider bounded noise and reason with the set of functions consistent with the measurements. Formally, we assume that the noise conforms to a model defined by a QMI given in terms of a partitioned matrix. The following assumption describes the noise model.

Assumption 1 (Noise model) *We assume that the noise samples satisfy $W^\top \in \mathcal{Z}(\Pi)$, where $\Pi \in \mathbb{S}^{m+T}$ is such that $\Pi_{22} < 0$ and $\Pi[\Pi_{22}] \geq 0$.*

According to [37, Thm. 3.2], under Assumption 1, the set $\mathcal{Z}(\Pi)$ is nonempty, convex, and bounded. A common example of such noise models is the case where $WW^\top = \sum_{i=1}^T w_i w_i^\top \leq Q$ for some $Q \geq 0$. In addition, in the setting of Gaussian noise, these noise models can be employed as confidence intervals corresponding to a given probability.

Under Assumption 1, one can describe the set of parameters consistent with the measurements as

$$\Theta = \{\theta \in \mathbb{R}^{k \times m} \mid Y = \theta^\top \Phi + W, W^\top \in \mathcal{Z}(\Pi)\}. \quad (3)$$

Note that we can write

$$Y = \theta^\top \Phi + W \iff \begin{bmatrix} I_m & W \end{bmatrix} = \begin{bmatrix} I_m & \theta^\top \end{bmatrix} \begin{bmatrix} I_m & Y \\ 0 & -\Phi \end{bmatrix}.$$

Thus, if we define

$$N := \begin{bmatrix} I_m & Y \\ 0 & -\Phi \end{bmatrix} \Pi \begin{bmatrix} I_m & Y \\ 0 & -\Phi \end{bmatrix}^\top, \quad (4)$$

it follows immediately that $\Theta = \mathcal{Z}(N)$. We refer to the procedure of obtaining the set of parameters Θ from the measurements as *set-valued regression*.

Remark 2.1 (Sufficiently exciting measurements)

Note that the set Θ is compact if and only if $N_{22} < 0$. Since $N_{22} = \Phi \Pi_{22} \Phi^\top$ and $\Pi_{22} < 0$, this holds if and only if Φ has full row rank. In turn, this requires that the basis functions are not identical and that the set of points $\{z_i\}_{i=1}^T$ where measurements are taken is ‘rich’ enough or sufficiently ‘exciting’, cf. [20]. •

Remark 2.2 (Least-squares estimate) One can check that, under Assumption 1, $\theta^{\text{lse}} := -N_{22}^\dagger N_{21} \in \Theta$. This means that the function $(\theta^{\text{lse}})^\top b(z)$ is consistent with the measurements. In fact,

$$N|N_{22} = \begin{bmatrix} I_m \\ -N_{22}^\dagger N_{21} \end{bmatrix}^\top N \begin{bmatrix} I_m \\ -N_{22}^\dagger N_{21} \end{bmatrix} \geq \begin{bmatrix} I_m \\ \theta \end{bmatrix}^\top N \begin{bmatrix} I_m \\ \theta \end{bmatrix},$$

for any $\theta^\top \in \mathcal{Z}(N)$. Therefore, θ^{lse} is the value that maximizes the value of the quadratic inequality. This leads us to refer to the function $\phi^{\text{lse}}(z; \Theta) := (\theta^{\text{lse}})^\top b(z) = -N_{12} N_{22}^\dagger b(z)$ as the *least-squares estimate* of $\hat{\phi}(z)$. •

In regards to the optimization of $\hat{\phi}$, note that changes in the parameter θ might lead to changes in the quantitative behavior and in particular the location of optimal values of the functions ϕ^θ . This means that a small error in estimating the true parameter corresponding to the unknown function $\hat{\phi}$ might lead to significant error. This motivates us to focus on suboptimization problems instead. To be precise, we investigate whether, for a given $\delta \in \mathbb{R}$ and $c \in \mathbb{R}^m$, we can ensure that

$$\|\hat{\phi}(z)\|_2 \leq \delta \quad \text{or} \quad c^\top \hat{\phi}(z) \leq \delta. \quad (5)$$

Thus the following problem statement.

Problem 1 (Cautious optimization) Consider an unknown function $\hat{\phi} : \mathbb{R}^n \rightarrow \mathbb{R}^m$, a noise model Π such that Assumption 1 holds, data (Y, Φ) of the true function, and the set Θ as in (3). Then

- (i) (Verification of suboptimality) given $z \in \mathbb{R}^n$, determine the minimal value of $\delta \in \mathbb{R}$ for which we can guarantee that either inequality in (5) holds;
- (ii) (One-shot cautious suboptimization) given a set $\mathcal{S} \subseteq \mathbb{R}^n$, find $z \in \mathcal{S}$ for which Problem 1(i) yields the minimal value of δ ;
- (iii) (Set-wise verification of suboptimality) given a set $\mathcal{S} \subseteq \mathbb{R}^n$, find the minimal $\delta \in \mathbb{R}$ for which either inequality in (5) holds for all $z \in \mathcal{S}$.

We note that being able to resolve any of these problems can be viewed as a property of the measurements, which is the viewpoint taken in the data informativity framework, e.g., [21, 37]: given $\delta \in \mathbb{R}$ and $\mathcal{S} \subseteq \mathbb{R}^n$, one could say that the data (Y, Φ) is *informative for δ -suboptimization on \mathcal{S}* if there exists $z \in \mathcal{S}$ such that $\|\hat{\phi}(z)\|_2 \leq \delta$.

We rely on the data informativity interpretation throughout our technical discussion in Section 3 to solve Problem 1. Then, in Section 4, we illustrate how to apply our results to solve various problems in data-driven analysis and control. First, we consider data-based contraction analysis for a class of nonlinear systems parameterized in terms of given basis functions. This essentially requires us to characterize one-sided Lipschitz constants for the unknown dynamics in terms of measured data. In terms of data-driven control, we consider the problem of data-based suboptimal regulation of unknown linear systems. Finally, in Section 5 we analyze various formulations of online suboptimization, where we iteratively optimize and collect new measurements to find the optimizer.

3 Cautious optimization

This section describes our solutions to Problem 1. Recall that, based on the data, we cannot distinguish $\hat{\phi}$ from any of the functions ϕ^θ as long as $\theta \in \Theta$. This means that we can only guarantee that $\|\hat{\phi}(z)\|_2 \leq \delta$ holds if $\sup_{\theta \in \Theta} \|\phi^\theta(z)\|_2 \leq \delta$. Similarly, given a vector $c \in \mathbb{R}^m$, we can only conclude that $c^\top \hat{\phi}(z) \leq \delta$ if $\sup_{\theta \in \Theta} c^\top \phi^\theta(z) \leq \delta$. This motivates the following definitions:

$$g(z; \Theta) := \sup_{\theta \in \Theta} \|\phi^\theta(z)\|_2, \quad g_c(z; \Theta) := \sup_{\theta \in \Theta} c^\top \phi^\theta(z). \quad (6)$$

These functions correspond to the elementwise worst-case realization of the unknown parameter $\hat{\theta}$. Therefore, we refer to $g(z; \Theta)$ as the *worst-case norm bound* and $g_c(z; \Theta)$ as the *worst-case linear bound*.

To draw conclusions regarding the true function $\hat{\phi}$, we investigate the functions $g(z; \Theta)$ and $g_c(z; \Theta)$. In fact, problems (i)-(iii) in Problem 1 can be recast as follows. Resolving Problem 1(i) (in Section 3.1) boils down to finding function values of $g(z; \Theta)$ or $g_c(z; \Theta)$. Resolving Problem 1(ii) (in Section 3.3) is equivalent to finding

$$\min_{z \in \mathcal{S}} g(z; \Theta) \quad \text{or} \quad \min_{z \in \mathcal{S}} g_c(z; \Theta). \quad (7)$$

Finally, resolving Problem 1(iii) (in Section 3.4) amounts to finding the minimal δ for which

$$\max_{z \in \mathcal{S}} g(z; \Theta) \leq \delta \quad \text{or} \quad \max_{z \in \mathcal{S}} g_c(z; \Theta) \leq \delta. \quad (8)$$

3.1 Pointwise verification of suboptimality

In order to resolve Problem 1(i), we are interested in obtaining values of the functions in (6). As a first observation, note that if the set Θ is compact (cf. Remark 2.1), then the supremum is attained, and we can replace it with a maximum. In this case, the worst-case norm bound and worst-case linear bound are finite-valued functions.

We rely on the following characterization to obtain function values of $g(z; \Theta)$.

Lemma 3.1 (Data-based conditions for bounds on the norm) *Given data (Y, Φ) , let $\Theta = \mathcal{Z}(N)$, with N as in (4). Suppose that N has at least one positive eigenvalue. Let $\delta \geq 0$ and $z \in \mathbb{R}^n$. Then,*

$$\|\phi^\theta(z)\|_2 \leq \delta \quad \text{for all} \quad \theta \in \Theta \quad (9)$$

if and only if there exists $\alpha \geq 0$ such that

$$\begin{bmatrix} \delta^2 I_m & 0 & 0 \\ 0 & 0 & b(z) \\ 0 & b(z)^\top & 1 \end{bmatrix} - \alpha \begin{bmatrix} N & 0 \\ 0 & 0 \end{bmatrix} \geq 0. \quad (10)$$

PROOF. Note that $\|\phi^\theta(z)\|_2^2 = b(z)^\top \theta \theta^\top b(z)$. This allows us to write (9) equivalently as

$$b(z)^\top \theta \theta^\top b(z) \leq \delta^2 \iff \theta^\top b(z) b(z)^\top \theta \leq \delta^2 I_m,$$

which can also be expressed as

$$\begin{bmatrix} I_m \\ \theta \end{bmatrix}^\top \begin{bmatrix} \delta^2 I_m & 0 \\ 0 & -b(z) b(z)^\top \end{bmatrix} \begin{bmatrix} I_m \\ \theta \end{bmatrix} \geq 0. \quad (11)$$

Thus (9) holds if and only if (11) holds for all $\theta \in \Theta$. Given that $\Theta = \mathcal{Z}(N)$, where N has at least one positive eigenvalue, the statement follows from applying [37, Thm 4.7] to (11) and using a Schur complement. ■

The necessary and sufficient conditions of Lemma 3.1 allow us to provide an alternative expression for $g(z; \Theta)$ that allows us to compute its value efficiently.

Corollary 3.2 (Values of the worst-case norm bound) *Given data (Y, Φ) , let $\Theta = \mathcal{Z}(N)$, with N as in (4). Suppose that N has at least one positive eigenvalue. For $z \in \mathbb{R}^n$, we have*

$$g(z; \Theta) = \min\{\delta \geq 0 \mid \exists \alpha \geq 0 \text{ s.t. (10) holds}\}. \quad (12)$$

Note that technically, the statement (10) is not an LMI. However, since (10) is linear in δ^2 , function values of $g(z; \Theta)$ can be found efficiently via Corollary 3.2. Similar situations pass without mention in the remainder of this paper.

Remark 3.3 (Weakening assumptions on data lead to conservative upper bound) If N does not have positive eigenvalues, i.e., $N \leq 0$, then the ‘if’-side of Lemma 3.1 remains true. Therefore, we can conclude that (12) in Corollary 3.2 holds with ‘ \leq ’ replacing ‘ $=$ ’. This means that we obtain an upper bound for $g(z; \Theta)$, leading to a conservative upper bound of $\hat{\phi}(z)$. •

Moreover, we can use Corollary 3.2 to find bounds independent of the choice of z .

Remark 3.4 (Relaxations of the conditions for bounds on the norm) If M is such that $b(z) b(z)^\top \leq M$, then

$$\begin{bmatrix} \delta^2 I_m & 0 \\ 0 & -b(z) b(z)^\top \end{bmatrix} \geq \begin{bmatrix} \delta^2 I_m & 0 \\ 0 & -M \end{bmatrix}.$$

This implies that

$$\begin{bmatrix} \delta^2 I_m & 0 \\ 0 & -M \end{bmatrix} - \alpha N \geq 0 \quad (13)$$

implies (10). In turn, this means that for any M such that $b(z) b(z)^\top \leq M$ for all z , we have

$$g(z; \Theta) \leq \min\{\delta \geq 0 \mid \exists \alpha \geq 0 \text{ s.t. (13) holds}\},$$

which yields a bound independent from z . •

We rely on the following characterization to obtain function values of $g_c(z; \Theta)$.

Lemma 3.5 (Data-based conditions for scalar upper bounds) *Given data (Y, Φ) , let $\Theta = \mathcal{Z}(N)$, with N as in (4). Let $\delta \in \mathbb{R}$, $z \in \mathbb{R}^n$, and $c \in \mathbb{R}^m$. Suppose $c^\top (N|N_{22}) c > 0$ and let*

$$N_c := \begin{bmatrix} c^\top N_{11} c & c^\top N_{12} \\ N_{21} c & N_{22} \end{bmatrix}. \quad (14)$$

Then it holds that

$$c^\top \phi^\theta(z) \leq \delta \quad \text{for all} \quad \theta \in \Theta \quad (15)$$

if and only if there exists $\alpha \geq 0$ such that

$$\begin{bmatrix} 2\delta & -b(z)^\top \\ -b(z) & 0 \end{bmatrix} - \alpha N_c \geq 0. \quad (16)$$

PROOF. We can rewrite $c^\top \phi^\theta(z)$ as $c^\top \theta^\top b(z) + b(z)^\top \theta c \leq 2\delta$. Thus, (15) holds if and only if

$$\gamma^\top b(z) + b(z)^\top \gamma \leq 2\delta \quad \text{for all } \gamma \in \mathcal{Z}(N)c,$$

where $\mathcal{Z}(N)c := \{Zc \mid Z \in \mathcal{Z}(N)\}$. Since $c \neq 0$, by [37, Thm. 3.4], we have $\mathcal{Z}(N)c = \mathcal{Z}(N_c)$. Therefore, (15) is equivalent to:

$$\begin{bmatrix} 1 \\ \gamma \end{bmatrix}^\top \begin{bmatrix} 2\delta & -b(z)^\top \\ -b(z) & 0 \end{bmatrix} \begin{bmatrix} 1 \\ \gamma \end{bmatrix} \geq 0, \quad \text{for all } \gamma \in \mathcal{Z}(N_c). \quad (17)$$

Given that $N_{22} \leq 0$, we have that $c^\top (N|N_{22})c > 0$ if and only if N_c has a positive eigenvalue. Then, the statement follows from applying [37, Thm 4.7] to (17). \blacksquare

We leverage the characterization in Lemma 3.5 to provide a closed-form expression of $g_c(z; \Theta)$ in terms of the data.

Theorem 3.6 (Explicit bounds in the scalar case) *Given data (Y, Φ) , let $\Theta = \mathcal{Z}(N)$, with N as in (4). Let $z \in \mathbb{R}^n$ and $c \in \mathbb{R}^m$. If $b(z) \in \text{im } \Phi$ or $c = 0$, then*

$$g_c(z; \Theta) = -c^\top N_{12} N_{22}^\dagger b(z) + \sqrt{c^\top (N|N_{22})cb(z)^\top (-N_{22}^\dagger)b(z)},$$

and, if $c \neq 0$ and $b(z) \notin \text{im } \Phi$, then $g_c(z; \Theta) = \infty$.

PROOF. Note that, if either $c = 0$ or $b(z) = 0$, then we have $g_c(z; \Theta) = 0$, thus the result follows. Assume then that $c \neq 0$, $b(z) \neq 0$, and $b(z) \in \text{im } \Phi$. If $c^\top (N|N_{22})c = 0$, then $N_c \leq 0$. As such, $\gamma \in \mathcal{Z}(N_c)$ if and only if

$$N_c \begin{bmatrix} 1 \\ \gamma \end{bmatrix} = 0 \iff \gamma \in -N_{22}^\dagger N_{21}c + \ker N_{22}.$$

Using (17), we see that $c^\top \phi^\theta \leq \delta$ for all $\theta \in \Theta$ if and only if $\gamma^\top b(z) \leq \delta$ for all $\gamma \in \mathcal{Z}(N_c)$. Combining the previous, we see that this holds if and only if $\delta \geq -c^\top N_{12} N_{22}^\dagger b(z)$ (where we have used that $b(z) \in \text{im } \Phi$). Minimizing δ reveals that the statement holds for this case.

Consider the case when $c^\top (N|N_{22})c > 0$. By Lemma 3.5, (15) holds if and only if (16) holds for some $\alpha \geq 0$. Since $b(z) \neq 0$, it must be that $\alpha > 0$. Using the Schur complement on (16), we obtain

$$\begin{aligned} 0 &\leq 2\delta - \alpha c^\top N_{11}c + \alpha^{-1} (b(z)^\top + \alpha c^\top N_{12}) N_{22}^\dagger (b(z) + \alpha N_{21}c) \\ &= 2\delta + c^\top N_{12} N_{22}^\dagger b(z) + b(z)^\top N_{22}^\dagger N_{21}c \\ &\quad - \alpha c^\top (N|N_{22})c + \alpha^{-1} b(z)^\top N_{22}^\dagger b(z). \end{aligned}$$

Clearly, there exists α for which this holds if and only if it holds for the value of α that maximizes the expression on the right. Note that $\ker N_{22} = \ker \Phi^\top$. Thus, since $b(z) \in \text{im } \Phi$ and $b(z) \neq 0$, we see that $b(z)^\top N_{22}^\dagger b(z) < 0$. The expression is then maximized over $\alpha \geq 0$ with

$$\alpha = \sqrt{\frac{-b(z)^\top N_{22}^\dagger b(z)}{c^\top (N|N_{22})c}}.$$

This yields that (16) holds if and only if

$$\begin{aligned} &2\delta + c^\top N_{12} N_{22}^\dagger b(z) + b(z)^\top N_{22}^\dagger N_{21}c \\ &\quad - 2\sqrt{c^\top (N|N_{22})cb(z)^\top (-N_{22}^\dagger)b(z)} \geq 0. \end{aligned}$$

To obtain $g_c(z; \Theta)$, we minimize this over δ , which yields the expression in the statement.

Finally, consider the case when $c \neq 0$ and $b(z) \notin \text{im } \Phi$. It is straightforward to check that if $\theta \in \Theta$, then $\theta + vw^\top \in \Theta$, for all $w \in \mathbb{R}^m$ and $v \in \ker \Phi^\top \subseteq \mathbb{R}^k$. Since $b(z) \notin \text{im } \Phi$ and $\text{im } \Phi$ is orthogonal to $\ker \Phi^\top$, we can take v such that $v^\top b(z) > 0$. Now, for arbitrary $\beta \in \mathbb{R}$, take $w = \beta c$. Then,

$$\sup_{\theta \in \Theta} \phi^\theta(z) \geq c^\top (\theta + vw^\top)^\top b(z) = c^\top \theta^\top b(z) + \beta c^\top cv^\top b(z).$$

Since $c^\top c > 0$ and $v^\top b(z) > 0$, and this is valid for any $\beta \in \mathbb{R}$, the last part of the statement follows. \blacksquare

Remark 3.7 (The special case of $N_{22} < 0$) Note that when Θ is compact (or equivalently, Φ has full row rank, cf. Remark 2.1), it is immediate that $b(z) \in \text{im } \Phi$ for any $z \in \mathbb{R}^n$, and hence $g_c(z; \Theta)$ is always finite valued. \bullet

The following illustrates the introduced concepts in the familiar setting of linear regression.

Example 1 (Linear regression) Suppose that $\hat{\phi}: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and $k = 3$. We consider basis functions $\phi_1(z) = 1$ and $\phi_2(z) = \begin{pmatrix} 1 & 0 \end{pmatrix} z$, $\phi_3(z) = \begin{pmatrix} 0 & 1 \end{pmatrix} z$. In other words, we take $\theta \in \mathbb{R}^{3 \times 2}$, and let $\phi^\theta(z) = \theta^\top \begin{bmatrix} 1 & z^\top \end{bmatrix}^\top$. We collect three measurements,

$$\begin{pmatrix} 1 \\ 1 \end{pmatrix} = \hat{\phi}(0) + w_1, \quad \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \hat{\phi}(e_1) + w_2, \quad \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \hat{\phi}(e_2) + w_3,$$

where we assume that $WW^\top \leq I_2$. This leads to $\Theta = \mathcal{Z}(N)$ with N such that

$$N_{11} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad N_{12} = \begin{bmatrix} 2 & 0 & 1 \\ 2 & 1 & 0 \end{bmatrix}, \quad N_{22} = \begin{bmatrix} -3 & -1 & -1 \\ -1 & -1 & 0 \\ -1 & 0 & -1 \end{bmatrix}.$$

Now, the least-squares estimate is equal to

$$\theta^{\text{lsc}} = \begin{bmatrix} 1 & 1 \\ -1 & 0 \\ 0 & -1 \end{bmatrix}, \quad \phi^{\text{lsc}}(z; \Theta) = \begin{pmatrix} 1 - z_1 \\ 1 - z_2 \end{pmatrix}.$$

In fact, using Theorem 3.6 we can conclude that

$$g_c(z; \Theta) = c^\top \begin{pmatrix} 1 - z_1 \\ 1 - z_2 \end{pmatrix} + \sqrt{c^\top c} \sqrt{(1 - z_1 - z_2)^2 + z_1^2 + z_2^2}.$$

As such, on the basis of the data we can guarantee that, for instance, $c^\top \hat{\phi}(0) \leq c^\top \begin{pmatrix} 1 \\ 1 \end{pmatrix} + \|c\|_2$. \bullet

3.2 Uncertainty on function values

The results in Section 3.1 allow us to find bounds for the unknown function $\hat{\phi}$, but do not consider how much these bounds deviate from its true value. For this, note that the result of Theorem 3.6 can be interpreted as quantifying




Figure 1. A plot depicting the situation of Example 1. For $c = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$, we plot a number function $c^\top \phi^\theta$, with θ consistent with the data. Any of these functions could be the unknown function. The least squares estimate is shown in yellow. If we are interested in bounds on $\hat{\phi}$, we can employ $g_c(z; \Theta)$ (shown in red). In particular note that this is not a linear map. Below the plot, we show the slices $z_1 = z_2$ and $z_1 = -z_2$ respectively.

the distance between the true function value and the least-squares estimate ϕ^{lse} (defined in Remark 2.2) in terms of the basis functions and the data. This observation leads us to define the *uncertainty* corresponding to (6) at z by

$$U(z; \Theta) := \sup_{\theta \in \Theta} \|\phi^\theta(z) - \phi^{\text{lse}}(z; \Theta)\|_2, \quad (18a)$$

$$U_c(z; \Theta) := \sup_{\theta \in \Theta} c^\top (\phi^\theta(z) - \phi^{\text{lse}}(z; \Theta)). \quad (18b)$$

Thus, if the uncertainty $U(z; \Theta)$ at z is small, then $\phi^{\text{lse}}(z; \Theta)$ is quantifiably a good estimate of $\hat{\phi}(z)$.

The uncertainty also allows us describe a useful variant to Problem 1(ii). In order to balance the demands of a low upper bound on the value of the true function with an associated low uncertainty, it is reasonable to consider the following generalization of the cautious suboptimization problem in (7): for $\lambda \geq 0$, consider

$$\min_{z \in \mathcal{S}} g(z; \Theta) + \lambda U(z; \Theta), \quad (19a)$$

$$\min_{z \in \mathcal{S}} g_c(z; \Theta) + \lambda U_c(z; \Theta). \quad (19b)$$

The following result provides a closed-form expression for the objective function in (19a).

Lemma 3.8 (Explicit expression of uncertainty in norms) *Given data (Y, Φ) , let $\Theta = \mathcal{Z}(N)$, with N as in (4). For $z \in \mathbb{R}^n$, if $b(z) \in \text{im } \Phi$, then*

$$U(z; \Theta) = \sqrt{\lambda_{\max}(N|N_{22})b(z)^\top (-N_{22}^\dagger)b(z)}. \quad (20)$$

If $b(z) \notin \text{im } \Phi$, then $U(z; \Theta) = \infty$.

PROOF. First suppose that N has at least one positive

eigenvalue. Similar to the proof of Lemma 3.1, we can conclude that $\|\phi^\theta(z) - \phi^{\text{lse}}(z; \Theta)\|_2 \leq \varepsilon$ for $\varepsilon \geq 0$ and for all $\theta \in \Theta$ if and only if there exists $\alpha \geq 0$ such that

$$\begin{bmatrix} \varepsilon^2 I & 0 \\ 0 & -b(z)b(z)^\top \end{bmatrix} - \alpha \begin{bmatrix} N|N_{22} & 0 \\ 0 & N_{22} \end{bmatrix} \geq 0. \quad (21)$$

Thus, in line with Corollary 3.2, we obtain

$$U(z; \Theta) = \min\{\varepsilon \geq 0 \mid \exists \alpha \geq 0 \text{ s.t. (21)}\}. \quad (22)$$

If $b(z) \in \text{im } \Phi$ then we can minimize ε by taking $\alpha = b(z)^\top (-N_{22}^\dagger)b(z)$, leading to (20). If N does not have at least one positive eigenvalue, then $N|N_{22} = 0$. As in Remark 3.3, we can conclude (22) with an inequality instead. Moreover, it immediately follows that $U(z; \Theta) \leq 0$, proving that (20) holds. The second part of the statement can be concluded in a similar fashion as the last part of the proof of Theorem 3.6. ■

In addition to expressing the uncertainty, this result also gives rise to a useful explicit bound on the function $g(z; \Theta)$.

Corollary 3.9 (Explicit upper bound of the worst-case norm bound) *Given data (Y, Φ) , let $\Theta = \mathcal{Z}(N)$, with N as in (4). For $z \in \mathbb{R}^n$, if $b(z) \in \text{im } \Phi$, then*

$$g(z; \Theta) \leq \|\phi^{\text{lse}}(z; \Theta)\|_2 + \sqrt{\lambda_{\max}(N|N_{22})b(z)^\top (-N_{22}^\dagger)b(z)}. \quad (23)$$

The following result provides a closed-form expression for the objective function in (19b).

Lemma 3.10 (Explicit expression of scalar uncertainty) *Given data (Y, Φ) , let $\Theta = \mathcal{Z}(N)$, with N as in (4). Let $z \in \mathbb{R}^n$ and $c \in \mathbb{R}^m$. If $b(z) \in \text{im } \Phi$ or $c = 0$, then*

$$U_c(z; \Theta) = \sqrt{(c^\top (N|N_{22})c)(b(z)^\top (-N_{22}^\dagger)b(z))}.$$

If $b(z) \notin \text{im } \Phi$ and $c \neq 0$, then $U_c(z; \Theta) = \infty$. Regardless, if $\lambda \geq 0$ and we define

$$N_\lambda := \begin{bmatrix} N_{11} & N_{12} \\ N_{21} & N_{22} \end{bmatrix} + \begin{bmatrix} \lambda(2 + \lambda)(N|N_{22}) & 0 \\ 0 & 0 \end{bmatrix},$$

then $g_c(z; \Theta) + \lambda U_c(z; \Theta) = g_c(z; \mathcal{Z}(N_\lambda))$.

PROOF. The first part follows immediately from Theorem 3.6. Now note that

$$N_\lambda = \begin{bmatrix} (1 + \lambda)^2(N|N_{22}) + N_{12}N_{22}^{-1}N_{21} & N_{12} \\ N_{21} & N_{22} \end{bmatrix}.$$

Therefore $N_\lambda|N_{22} = (1 + \lambda)^2(N|N_{22})$. Thus, the result follows from the application of Theorem 3.6. ■

Lemma 3.10 means that, even though problem (19b) is more general than the corresponding cautious suboptimization problem in (7), both problems can be resolved in the same fashion.

3.3 Convexity and suboptimization

In order to resolve the one-shot cautious suboptimization problem efficiently, cf. Problem 1(ii), we investigate conditions under which $g(\cdot; \Theta)$ and $g_c(z; \Theta)$ are convex. Towards this, we first investigate conditions under which we can guarantee that the true function $\hat{\phi}$ is convex. Since, non-negative combinations of convex functions are convex, we provide conditions that ensure the set of parameters consistent with the measurements are nonnegative.

Lemma 3.11 (Nonnegativity of parameters) *Given data (Y, Φ) , let $\Theta = \mathcal{Z}(N)$, with N is as in (4). Let $c \neq 0$. Then, $\theta c \in \mathbb{R}_{\geq 0}^k$ for all $\theta \in \Theta$ if and only if Φ has full row rank and one of the following conditions hold*

- (i) $c^\top(N|N_{22})c = 0$ and $-N_{22}^{-1}N_{21}c \in \mathbb{R}_{\geq 0}^k$, or
- (ii) $c^\top(N|N_{22})c > 0$, $-N_{22}^{-1}N_{21}c \in \mathbb{R}_{> 0}^k$ and for all $i = 1, \dots, k$,

$$c^\top(N|N_{22})c + \frac{(e_i^\top N_{22}^{-1} N_{21} c)^2}{e_i^\top N_{22}^{-1} e_i} \leq 0.$$

PROOF. Recall the definition of N_c from (14). Since $c \neq 0$ we have that $\mathcal{Z}(N)c = \mathcal{Z}(N_c)$. Thus, $\theta c \in \mathbb{R}_{\geq 0}^k$ for all $\theta \in \Theta$ if and only if $\mathcal{Z}(N_c) \subseteq \mathbb{R}_{\geq 0}^k$. Suppose that $\mathcal{Z}(N_c) \subseteq \mathbb{R}_{\geq 0}^k$. This implies that the set $\mathcal{Z}(N_c)$ contains no subspace, which implies that $N_{22} < 0$. In turn, this holds if and only if Φ has full row rank.

To complete the proof, we prove that if Φ has full row rank, then $\mathcal{Z}(N_c) \subseteq \mathbb{R}_{\geq 0}^k$ if and only if either (i) or (ii) holds. If $c^\top(N|N_{22})c = 0$, then $N_c \leq 0$. As such, $\mathcal{Z}(N_c) = \{-N_{22}^{-1}N_{21}\}$, that is, *only* the least-squares estimate is compatible with the data. The first part of the proof then follows immediately.

Next, consider $c^\top(N|N_{22})c > 0$, that is, N_c has one positive eigenvalue. Note that the Slater condition holds, and we can apply the S-Lemma [38, Thm 2.2] to see that $\mathcal{Z}(N_c) \subseteq \mathbb{R}_{\geq 0}^k$ is equivalent to the existence of $\alpha_1, \dots, \alpha_k \geq 0$ such that

$$\begin{bmatrix} 0 & e_i^\top \\ e_i & 0 \end{bmatrix} - \alpha_i N_c \geq 0,$$

for each standard basis vector e_i . This requires $\alpha_i \neq 0$ for each $i \in \{1, \dots, k\}$. Thus, this can equivalently written as

$$N_c - \frac{1}{\alpha_i} \begin{bmatrix} 0 & e_i^\top \\ e_i & 0 \end{bmatrix} \leq 0.$$

Since $N_{22} < 0$, we can define $\beta_i = \frac{1}{\alpha_i}$ and conclude that the latter holds if and only if

$$\begin{aligned} & c^\top N_{11} c - (c^\top N_{12} - \beta_i e_i^\top) N_{22}^{-1} (N_{21} c - \beta_i e_i) \\ &= c^\top (N|N_{22}) c + \beta_i (c^\top N_{12} N_{22}^{-1} e_i + e_i^\top N_{22}^{-1} N_{21} c) \\ & \quad - \beta_i^2 e_i^\top N_{22}^{-1} e_i \leq 0. \end{aligned}$$

In turn, there exist such β_i if and only if the minimal value over all $\beta_i > 0$ satisfies this. Since this is a scalar expression, we can explicitly minimize it to prove the statement. \blacksquare

Now we can apply this result to identify conditions that ensure the convexity of $\hat{\phi}$ and its upper bound.

Corollary 3.12 (Convexity of the true function)

Given data (Y, Φ) , let $\Theta = \mathcal{Z}(N)$, with N is as in (4). Assume the basis functions ϕ_i are convex. Let $c \neq 0$ with $\theta c \in \mathbb{R}_{\geq 0}^k$ for all $\theta \in \Theta$. Then,

- $c^\top \phi^\theta$ is convex for all $\theta \in \Theta$ and $g_c(\cdot; \Theta)$ is a finite-valued convex function;
- If, in addition, the functions ϕ_i are strictly convex and $0 \notin \mathcal{Z}(N_c)$, then $c^\top \phi^\theta$ is strictly convex for all $\theta \in \Theta$ and $g_c(\cdot; \Theta)$ is strictly convex.

Moreover, $\|\phi^\theta(\cdot)\|_2$ is convex if $e_i^\top \phi^\theta$ is convex and nonnegative for all $i = 1, \dots, m$.

PROOF. Lemma 3.11 shows that, if $\theta c \in \mathbb{R}_{\geq 0}^k$ for all $\theta \in \Theta$, then Φ has full row rank or, equivalently, $N_{22} < 0$. This implies, cf. Remark 2.1, that $\mathcal{Z}(N)$ is compact and therefore $g_c(\cdot; \Theta)$ is finite-valued. The first result now follows from the facts that nonnegative combinations of convex functions are convex and that the maximum of any number of convex functions is also convex. The last part follows from standard composition rules for convexity, see e.g., [39, Example 3.14]. \blacksquare

Note that $0 \in \mathcal{Z}(N_c)$ if and only if $c^\top N_{11} c \geq 0$. Equivalently, this means that $Y^\top c \in \mathcal{Z}(\Pi_c)$, where Π_c is defined analogously to (14). This condition is easy to check in terms of data. Lemma 3.11 and Corollary 3.12 taken together mean that, if the basis functions are convex, we can test for convexity on the basis of data.

Looking back at Example 1, we can make two interesting observations. First, the conditions of Corollary 3.12 do *not* hold in the example, since for instance $\theta^{\text{lse}} c \in \mathbb{R}_{\geq 0}^3$ only if $c = 0$. Yet $c^\top \phi^\theta$ is convex for all $\theta \in \Theta$. This is because the basis functions ϕ_i are linear, and therefore the coefficients θ_i are not required to be nonnegative. Second, $g_c(\cdot; \Theta)$ in Example 1 is strictly convex for $c \neq 0$ even though the basis functions are not. Recall that we are interested in the optimization problem (7), and thus not necessarily in properties of the true function $\hat{\phi}$, but of the worst-case linear bound $g_c(\cdot; \Theta)$. These observations motivate our ensuing discussion to provide conditions for convexity of the upper bound instead. Under the assumptions of Theorem 3.6:

$$g_c(z; \Theta) = c^\top \phi^{\text{lse}}(z; \Theta) + U_c(z; \Theta).$$

Thus, if (i) $c^\top \phi^{\text{lse}} = -c^\top N_{12} N_{22}^{-1} b(\cdot)$ is convex and (ii) U_c is convex, then so is g_c . Moreover, if in addition either is strictly convex, then so is g_c . Condition (i) can be checked directly if all basis functions ϕ_i are twice continuously differentiable by computing the Hessian of $c^\top \phi^{\text{lse}}$. Here, we present a simple criterion, also derived from composition rules for convexity (again, see e.g., [39, Example 3.14]), to test for condition (ii).

Corollary 3.13 (Convexity of the uncertainties)

Given data (Y, Φ) , let $\Theta = \mathcal{Z}(N)$ and assume $N_{22} < 0$. Then both $U(\cdot; \Theta)$ and $U_c(\cdot; \Theta)$ are convex if each basis function ϕ_i is convex and $-N_{22}^{-1} b(z) \geq 0$ for all $z \in \mathbb{R}^n$.

In addition to guaranteeing that $g_c(\cdot; \Theta)$ is convex, this result can be used to find a convex upper bound when optimizing the norm, by bounding $g(\cdot; \Theta)$ using Corollary 3.9.

Equipped with the results of this section, one can solve the cautious suboptimization problems (7) efficiently. Under

the assumptions of Theorem 3.6, we can write a closed-form expression for the function $g_c(\cdot; \Theta)$. This, in turn, allows us to test for (strict) convexity using e.g., Corollaries 3.12 or 3.13. If so, we can apply (projected) gradient descent for the optimization. For this, let $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be differentiable, and denote its Jacobian by

$$J(f) := \begin{bmatrix} \frac{df}{dz_1} & \dots & \frac{df}{dz_n} \\ \vdots & & \vdots \\ \frac{df_m}{dz_1} & \dots & \frac{df_m}{dz_n} \end{bmatrix}.$$

Clearly, $J(\phi^\theta) = \theta^\top J(b)$. The following result, consequence of Theorem 3.6, provides an expression for the Jacobian of the worst-case linear bound function.

Corollary 3.14 (Jacobians of the worst-case linear bound) *Given data (Y, Φ) , let $\Theta = \mathcal{Z}(N)$ and assume $N_{22} < 0$. Let the basis functions $\phi_i(z)$ be differentiable and such that $b(z) \neq 0$ for all z . Then*

$$J(g_c(z; \Theta)) = \left(c^\top N_{12} + \frac{\sqrt{c^\top (N|N_{22})c}}{\sqrt{b(z)^\top (-N_{22}^{-1})b(z)}} b(z)^\top \right) (-N_{22}^{-1}) J(b(z)).$$

In a similar fashion, we can also obtain gradients of the bound of $g(z; \Theta)$ given in (23). Note that, if $N_{22} < 0$ this requires that both $b(z) \neq 0$ and $\phi^{\text{lse}}(z; \Theta) \neq 0$ for all z .

3.4 Set-wise verification of suboptimality

Here we solve Problem 1(iii) and provide upper (and lower) bounds as in (5) which are guaranteed to hold for all $z \in \mathcal{S} \subseteq \mathbb{R}^n$. In terms of g and g_c , this holds only if (8) holds.

Note that $g_{-c}(z; \Theta) = -g_c(z; \Theta)$, and therefore $g_c(z; \Theta)$ is concave if and only if $g_{-c}(\cdot; \Theta)$ is convex. Moreover,

$$\max_{z \in \mathcal{S}} g_c(z; \Theta) = -\min_{z \in \mathcal{S}} g_{-c}(z; \Theta).$$

Therefore, we can test for concavity of g_c and apply the minimization results to g_{-c} of the Section 3.3. Beyond the case of concavity, we can also employ convexity in order to efficiently provide upper bounds over the convex hull of a finite set, as the following result shows.

Proposition 3.15 (Convexity and maximal values) *Given data (Y, Φ) , let $\Theta = \mathcal{Z}(N)$. Let $\mathcal{F} \subseteq \mathbb{R}^n$ be a finite set and let $\mathcal{S} = \text{conv } \mathcal{F}$. Suppose that $g(z; \Theta)$ is convex. Then (9) for all $z \in \mathcal{S}$ if and only if*

$$g(z; \Theta) \leq \delta \quad \forall z \in \mathcal{F}.$$

Similarly, if $g_c(z; \Theta)$ is convex, then (15) for all $z \in \mathcal{S}$ if and only if

$$g_c(z; \Theta) \leq \delta \quad \forall z \in \mathcal{F}.$$

The proof of this result follows immediately from the fact that any convex function attains its maximum over \mathcal{S} on \mathcal{F} . Thus, convexity (resp., concavity) allows us to find minimal upper bounds (resp., maximal lower bounds) over a set, leading to a solution of Problem 1.(iii). In order to provide bounds over sets in scenarios where convexity or concavity cannot be guaranteed, we investigate Lipschitz continuity. This leads to conditions that can be checked at a finite

number of points. Therefore, we are first interested in the question whether for all $z, z^* \in \mathcal{S} \subseteq \mathbb{R}^n$, it holds that

$$\|\hat{\phi}(z) - \hat{\phi}(z^*)\|_2 \leq L \|z - z^*\|_2.$$

More generally, and for later reference, the following result considers the more general case of weighted norms.

Lemma 3.16 (Lipschitz constants of functions consistent with the measurements) *Given data (Y, Φ) , let $\Theta = \mathcal{Z}(N)$, with N as in (4), and assume N has at least one positive eigenvalue. Let $P \in \mathbb{S}^m$ and $Q \in \mathbb{S}^n$ with $P > 0$ and $Q > 0$. For $L \geq 0$ and $z, z^* \in \mathcal{S} \subseteq \mathbb{R}^n$, we have*

$$\|\phi^\theta(z) - \phi^\theta(z^*)\|_{2, P^{1/2}} \leq L \|z - z^*\|_{2, Q^{1/2}} \quad \text{for all } \theta \in \Theta$$

if and only if there exists $\alpha \geq 0$ such that

$$\begin{bmatrix} L^2 P^{-1} & 0 & 0 \\ 0 & 0 & \frac{b(z) - b(z^*)}{\|z - z^*\|_{2, Q^{1/2}}} \\ \hline 0 & \frac{b(z)^\top - b(z^*)^\top}{\|z - z^*\|_{2, Q^{1/2}}} & 1 \end{bmatrix} - \alpha \begin{bmatrix} N & 0 \\ 0 & 0 \end{bmatrix} \geq 0. \quad (24)$$

If, in addition the basis functions ϕ_i are differentiable, then

$$\|J(\phi^\theta)\|_2 \leq L \quad \text{for all } \theta \in \Theta$$

if and only if there exists $\alpha \geq 0$ such that

$$\begin{bmatrix} L^2 I_m & 0 & 0 \\ 0 & 0 & J(b) \\ \hline 0 & J(b)^\top & I_n \end{bmatrix} - \alpha \begin{bmatrix} N & 0 \\ 0 & 0 \end{bmatrix} \geq 0. \quad (25)$$

PROOF. Note that $\phi^\theta(z) - \phi^\theta(z^*) = \theta^\top (b(z) - b(z^*))$. Similarly, $J(\phi^\theta) = \theta^\top J(b)$. The result can then be proven by following the steps of the proof of Lemma 3.1. ■

Remark 3.17 (Single check for establishing Lipschitz constant) To check for Lipschitzness, we can employ the method of Remark 3.4 to avoid checking (24) for all $z, z^* \in \mathbb{R}^n$. Suppose the basis satisfies a Lipschitz condition of the form

$$\|b(z) - b(z^*)\|_2 \leq L_b \|z - z^*\|_{2, Q^{1/2}},$$

for all $z, z^* \in \mathbb{R}^n$. Then,

$$\frac{(b(z) - b(z^*))(b(z) - b(z^*))^\top}{\|z - z^*\|_{2, Q^{1/2}}^2} \leq L_b^2 I,$$

and thus, as in Remark 3.4 we have

$$\begin{bmatrix} L^2 P^{-1} & 0 \\ 0 & -L_b^2 I_k \end{bmatrix} - \alpha N \geq 0 \quad (26)$$

implies (24). This gives a single condition to check to establish a Lipschitz constant valid for all $z, z^* \in \mathbb{R}^n$. •

The characterization of Lipschitz constants in terms of data in Lemma 3.16 allows us to bound the unknown function $\hat{\phi}$ as follows. To do this, we use the following notion: The set $\mathcal{G} \subset \mathbb{R}^n$ is an ε -covering of $\mathcal{S} \subseteq \mathbb{R}^n$ if for every $z \in \mathcal{S}$ there exists $z^* \in \mathcal{G}$ such that $\|z - z^*\|_2 \leq \varepsilon$. If \mathcal{S} is a bounded set, then for any $\varepsilon > 0$, there exists an ε -covering with finitely many elements.

Theorem 3.18 (Coverings and bounds) *Given data (Y, Φ) , let $\Theta = \mathcal{Z}(N)$. For $\varepsilon > 0$, let \mathcal{G} be a finite ε -covering of $\mathcal{S} \subseteq \mathbb{R}^n$. Suppose that $g(z; \Theta)$ is Lipschitz with constant L on \mathcal{S} . Then, for all $z \in \mathcal{S}$,*

$$\min_{z^* \in \mathcal{G}} g(z^*; \Theta) - \varepsilon L \leq g(z; \Theta) \leq \max_{z^* \in \mathcal{G}} g(z^*; \Theta) + \varepsilon L.$$

If $g_c(z; \Theta)$ is Lipschitz with constant L , then for all $z \in \mathcal{S}$,

$$\min_{z^* \in \mathcal{G}} g_c(z^*; \Theta) - \varepsilon L \leq g_c(z; \Theta) \leq \max_{z^* \in \mathcal{G}} g_c(z^*; \Theta) + \varepsilon L.$$

We omit the proof, which follows directly from combining the definitions of Lipschitz continuity and ε -coverings.

This result means that we can find guaranteed upper and lower bounds of either $\|\hat{\phi}(z)\|$ or $c^\top \hat{\phi}(z)$ over the bounded set \mathcal{S} in terms of a finite number of evaluations of $g(z; \Theta)$ or $g_c(z; \Theta)$, resp. In turn, recall that Corollary 3.2 allows us to efficiently find function values of $g(z; \Theta)$ on the basis of measurements and, similarly, Theorem 3.6 allows us to directly calculate values of $g_c(z; \Theta)$.

4 Applications to system analysis and control

In this section, we exploit the proposed solutions to Problem 1 in two applications: contraction analysis of unknown nonlinear systems and regulation of an unknown linear system to a suboptimal point of an unknown cost function.

4.1 Data-based contraction analysis for nonlinear systems

Consider the autonomous discrete- or continuous-time system given by either

$$z_{k+1} = \hat{\phi}(z_k), \text{ or } \dot{z}(t) = \hat{\phi}(z(t)), \quad (27)$$

where $\hat{\phi} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is an unknown function. We consider noisy measurements (2) as described in Section 2. Therefore $\hat{\phi} = \phi^\theta$ for some $\theta \in \Theta = \mathcal{Z}(N)$. In the continuous-time case, this means that we collect noisy measurements of the derivative of the state at a finite set of states (see [27] for a discussion on this assumption).

The discrete-time system in (27) is *strongly contracting* (cf. [29, Sec. 3.4]) with respect to the weighted norm $\|\cdot\|_{2, P^{1/2}}$ if $\hat{\phi}$ admits a Lipschitz constant $L < 1$, that is,

$$\|\hat{\phi}(z) - \hat{\phi}(z^*)\|_{2, P^{1/2}} \leq L \|z - z^*\|_{2, P^{1/2}}.$$

One can now directly employ Lemma 3.16 in order to obtain data-based conditions under which this property holds.

To deal with the situation of continuous-time systems, we first require a number of prerequisites. Given $P > 0$, a function $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is *one-sided Lipschitz with respect to* $\|\cdot\|_{2, P^{1/2}}$ if there exists $\gamma \in \mathbb{R}$ such that, for all $z, z^* \in \mathbb{R}^n$,

$$(z - z^*)^\top P(f(z) - f(z^*)) \leq \gamma \|z - z^*\|_{2, P^{1/2}}^2. \quad (28)$$

Such γ is called a *one-sided Lipschitz constant* and the smallest such γ is denoted $\text{osLip}(f)$. The autonomous system $\dot{z}(t) = f(z(t))$ is *strictly contracting* with rate $|b|$ if $\text{osLip}(f) \leq b < 0$. Given any two trajectories $x(t), \bar{x}(t)$ of a strictly contracting system, one has $\|x(t) - \bar{x}(t)\|_{2, P^{1/2}} \leq e^{b(t-s)} \|x(s) - \bar{x}(s)\|_{2, P^{1/2}}$ for any $t \geq s \geq 0$.

The following result, derived from Theorem 3.6, establishes a test for strict contractivity of the continuous-time system in (27) on the basis of a set of measurements of $\hat{\phi}$.

Theorem 4.1 (Data-based test for strict contractivity) *Given data (Y, Φ) , let $\Theta = \mathcal{Z}(N)$, with N as in (4). Let $P > 0$ and $z, z^* \in \mathbb{R}^n$. Assume Φ has full row rank. Then,*

$$(z - z^*)^\top P \theta^\top (b(z) - b(z^*)) \leq \gamma \|z - z^*\|_{2, P^{1/2}}^2$$

for all $\theta \in \Theta$ if and only if

$$\begin{aligned} & (z - z^*)^\top P (-N_{12} N_{22}^{-1}) (b(z) - b(z^*)) \\ & + \sqrt{(z - z^*)^\top P (N | N_{22}) P (z - z^*)} \\ & \cdot \sqrt{(b(z) - b(z^*))^\top (-N_{22}^{-1}) (b(z) - b(z^*))} \\ & \leq \gamma \|z - z^*\|_{2, P^{1/2}}^2. \end{aligned}$$

As a consequence of this result, we can provide a test to establish strict contractivity in terms only of the least-squares estimate ϕ^{lse} , the data matrix N , and a Lipschitz constant for the basis functions.

Corollary 4.2 (Strict contractivity in terms of the least-squares estimate) *Given data (Y, Φ) , let $\Theta = \mathcal{Z}(N)$, with N as in (4). Assume Φ has full row rank. Let L be such that $\|b(z) - b(z^*)\|_{2, (-N_{22})^{-1/2}} \leq L \|z - z^*\|_2$, for all $z, z^* \in \mathbb{R}^n$. Then, for $P > 0$,*

$$(z - z^*)^\top P \theta^\top (b(z) - b(z^*)) \leq \gamma \|z - z^*\|_{2, P^{1/2}}^2$$

for all $\theta \in \Theta$ and $z, z^* \in \mathbb{R}^n$ if

$$\text{osLip}(\phi^{\text{lse}}) < \gamma - L \sqrt{\lambda_{\max}(N | N_{22})} \frac{\lambda_{\max}(P)}{\lambda_{\min}(P)}. \quad (29)$$

Remark 4.3 (Contraction w.r.t. different norms)

The results of this section can be readily generalized to contractivity (and thus one-sided Lipschitzness) with respect to $\|\cdot\|_{p, R}$ for $p \in [1, \infty)$. To do this, one replaces the one-sided Lipschitz condition (28) with the respective condition from [40, Table I] and adjust the statements accordingly. Note that the case of $\|\cdot\|_{\infty, R}$ is not amenable to this treatment because the corresponding one-sided Lipschitz condition $\text{osLip}(f)$ is not linear in f . •

4.2 Suboptimal regulation of unknown systems

Consider the problem of regulating an unknown linear system to a suboptimal point of the norm of an unknown cost function $\hat{\phi} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ on the basis of measurements. We consider noisy measurements (2) as described in Section 2. In short, $\hat{\phi} = \phi^\theta$ for some $\theta \in \Theta = \mathcal{Z}(N)$. Moreover, let

$$x_{k+1} = \hat{A}x_k + \hat{B}u_k,$$

with state $x_k \in \mathbb{R}^n$ and input $u_k \in \mathbb{R}^r$. Here, $\hat{A} \in \mathbb{R}^{n \times n}$ and $\hat{B} \in \mathbb{R}^{n \times r}$. If \hat{A} is stable, we know that any fixed point corresponds to a constant input u^* . Indeed, if we apply the constant input $u_k = u^*$, no matter the initial condition x_0 , the state asymptotically converges to the fixed point

$$\lim_{k \rightarrow \infty} x_k = x^* := (I - \hat{A})^{-1} \hat{B}u^*.$$

This means that we can find a static input to regulate the system towards any state in $\text{im}(I - \hat{A})^{-1} \hat{B} \subseteq \mathbb{R}^n$.

We are interested in regulating the state to a point x^* for which we can guarantee that $\|\hat{\phi}(x^*)\|_2 \leq \delta$, with a value of δ as small as possible. If \hat{A} and \hat{B} are known, this corresponds to Problem 1(ii) with $\mathcal{S} = \text{im}(I - \hat{A})^{-1}\hat{B}$.

Now assume that we do not have access to the matrices \hat{A} and \hat{B} . Instead, we assume that we have not only measurements of $\hat{\phi}$, but additionally a set of noisy measurements corresponding to the unknown system. To be precise, we collect measurements of the system

$$x_{k+1} = \hat{A}x_k + \hat{B}u_k + w_k. \quad (30)$$

Using the basis functions

$$\psi_i(x, u) = \begin{cases} x_i & \text{for } i = 1, \dots, n \\ u_i & \text{for } i = n + 1, \dots, n + r, \end{cases}$$

and collecting measurements of the state and input signals reduces the setup of Section 2 to that of e.g., [37]. To avoid some unnecessary repetition, we simply state here that on the basis of measurements, we can conclude that $(A, B)^\top \in \Sigma = \mathcal{Z}(M)$, where $M \in \mathbb{S}^{n+(n+r)}$ is analogous to N in (4).

In order to be able to regulate the unknown system in the same manner as before, recall that we require that \hat{A} is stable. Given that, on the basis of measurements we can not distinguish \hat{A} from other matrices A such that $(A, B)^\top \in \Sigma$, we therefore require all such A to be stable. In the following, we describe a test for the existence of a shared quadratic Lyapunov function, which is a stronger stability condition. For a detailed discussion on the conservativeness of this assumption, see e.g., [41]. Now we can easily adapt e.g., [37, Thm. 5.1] to obtain the following.

Lemma 4.4 (Informativity for stability) *Let $\Sigma = \mathcal{Z}(M)$, where the corresponding noise model satisfies Assumption 1. Then, there exists $P \in \mathbb{S}^n$ with $P > 0$ and $APA^\top < P$ for all $(A, B)^\top \in \Sigma$ if and only if there exist $\bar{P} \in \mathbb{S}^n$ with $\bar{P} > 0$ and $\beta > 0$ such that*

$$\begin{bmatrix} \bar{P} - \beta I_n & 0 & 0 \\ 0 & -\bar{P} & 0 \\ 0 & 0 & 0 \end{bmatrix} - M \geq 0. \quad (31)$$

Remark 4.5 (Stability and contractivity) Applying a fixed input u^* to the unknown system yields an autonomous discrete-time system. We know that for any initial condition x_0 , this system converges to a fixed point. Moreover, if $P > 0$ and $APA^\top < P$ for all $(A, B)^\top \in \Sigma$, then $A^\top P^{-1}A < P^{-1}$. Therefore, one can see that the system is strongly contracting with respect to $\|\cdot\|_{2, P^{-1/2}}$. •

When each system matrix A in the set Σ of systems consistent with the data is stable, we can characterize the fixed points resulting from applying the same input to each.

Lemma 4.6 (Characterizing fixed points) *Let $\Sigma = \mathcal{Z}(M)$, where the corresponding noise model satisfies Assumption 1 and M has at least one positive eigenvalue. Assume A is stable for all $(A, B)^\top \in \Sigma$. Given $x \in \mathbb{R}^n$ and $u \in \mathbb{R}^r$, $\|x - (I - A)^{-1}Bu\|_2 \leq \varepsilon$ for all $(A, B)^\top \in \Sigma$ if and*

only if there exists $\gamma \geq 0$ such that

$$\begin{bmatrix} I_n & -I_n & 0 & \vdots & x \\ -I_n & I_n & 0 & \vdots & -x \\ 0 & 0 & 0 & \vdots & -u \\ \hline x^\top & -x^\top & -u^\top & \vdots & \varepsilon^2 \end{bmatrix} - \gamma \begin{bmatrix} M & 0 \\ 0 & 0 \end{bmatrix} \geq 0 \quad (32)$$

PROOF. Let $(A, B)^\top \in \Sigma$. First, note that we have

$$\begin{aligned} \|x - (I_n - A)^{-1}Bu\|_2 &\leq \varepsilon \\ \Leftrightarrow (x - (I_n - A)^{-1}Bu)(x - (I_n - A)^{-1}Bu)^\top &\leq \varepsilon^2 I. \end{aligned}$$

Since A is stable, we know that $I - A$ is nonsingular. Thus, the previous holds if and only if

$$((I - A)x - Bu)((I - A)x - Bu)^\top \leq \varepsilon^2 (I - A)(I - A)^\top,$$

or equivalently,

$$\begin{bmatrix} I_n \\ A^\top \\ B^\top \end{bmatrix}^\top \left(\begin{bmatrix} \varepsilon^2 I_n & -\varepsilon^2 I_n & 0 \\ -\varepsilon^2 I_n & \varepsilon^2 I_n & 0 \\ 0 & 0 & 0 \end{bmatrix} - \begin{pmatrix} x \\ -x \\ -u \end{pmatrix} \begin{pmatrix} x \\ -x \\ -u \end{pmatrix}^\top \right) \begin{bmatrix} I_n \\ A^\top \\ B^\top \end{bmatrix} \geq 0.$$

As such, the set of (A, B) that satisfy the condition can be written as the solution set of a QMI. By assumption $\Sigma = \mathcal{Z}(M)$. We can then apply the matrix S-Lemma [37, Thm 4.7] and use a Schur complement to prove the statement. ■

Motivated by Lemma 4.6, we define

$$\varepsilon^-(x) := \min\{\varepsilon \mid \exists \gamma \geq 0, u \in \mathbb{R}^r \text{ s.t. (32) holds}\}. \quad (33)$$

Given $x \in \mathbb{R}^n$, the function $\varepsilon^-(x)$ thus gives the minimal radius of a ball around x , for which there exists a u^* such that the corresponding set of fixed points of each system in Σ is contained in this ball. The next result describes useful properties of this function.

Lemma 4.7 (Properties of the minimal radius function) *Let $\Sigma = \mathcal{Z}(M)$, where the corresponding noise model satisfies Assumption 1 and M has at least one positive eigenvalue. Assume A is stable for all $(A, B)^\top \in \Sigma$. Then $\varepsilon^-(x) \leq \|x\|_2$ for any $x \in \mathbb{R}^n$ and ε^- is convex.*

PROOF. The statement $\varepsilon^-(x) \leq \|x\|_2$ follows by noting that, for any x and M , (32) holds with $u = 0$, $\gamma = 0$ and $\varepsilon = \|x\|_2$. Next, suppose that for all $(A, B)^\top \in \Sigma$,

$$\|x - (I - A)^{-1}Bu\|_2 \leq \varepsilon, \quad \|\bar{x} - (I - A)^{-1}B\bar{u}\|_2 \leq \bar{\varepsilon}.$$

Let $\lambda \in [0, 1]$ and define $x_\lambda := \lambda x + (1 - \lambda)\bar{x}$ and similarly for u_λ and ε_λ . Then,

$$\begin{aligned} \|x_\lambda - (I - A)^{-1}Bu_\lambda\|_2 &\leq \lambda \|x - (I - A)^{-1}Bu\|_2 \\ &\quad + (1 - \lambda) \|\bar{x} - (I - A)^{-1}B\bar{u}\|_2 \leq \varepsilon_\lambda. \end{aligned}$$

Therefore, $\varepsilon^-(x_\lambda) \leq \varepsilon_\lambda$, proving that ε^- is convex in x . ■

The following result uses this function in order to provide conditions for suboptimal regulation.

Theorem 4.8 (Suboptimal regulation) Let $\Sigma = \mathcal{Z}(M)$ and $\Theta = \mathcal{Z}(N)$, where the corresponding noise models satisfy Assumption 1. Assume A is stable for all $(A, B)^\top \in \Sigma$ and ϕ^θ is Lipschitz with constant L for all $\theta \in \Theta$. Then

$$\min_u \|\hat{\phi}((I - \hat{A})^{-1} \hat{B}u)\|_2 \leq \min_x (g(x; \Theta) + L\varepsilon^-(x)).$$

PROOF. Using the triangle inequality and Lipschitzness,

$$\|\hat{\phi}((I - \hat{A})^{-1} \hat{B}u)\|_2 \leq \|\hat{\phi}(x)\|_2 + L\|x - (I - \hat{A})^{-1} \hat{B}u\|_2,$$

for any $x \in \mathbb{R}^n$. Note that this separates the unknown parameter $\hat{\theta}$ from the unknown pair (\hat{A}, \hat{B}) . From here, we deduce

$$\begin{aligned} & \min_u \|\hat{\phi}((I - \hat{A})^{-1} \hat{B}u)\|_2 \\ & \leq \min_{x, u} \left(\max_{\theta \in \Theta} \|\phi^\theta(x)\| + L \left(\max_{(A, B) \in \Sigma} \|x - (I - A)^{-1} Bu\| \right) \right). \end{aligned}$$

Since only the second term is dependent on u , we can use the definition of $\varepsilon^-(x)$ and Lemma 4.6 to conclude the statement. ■

Now, since we know that ε^- is convex, we can resolve this problem efficiently if so is $g(\cdot, \Theta)$ (see Corollary 3.13). We conclude this section by noting that Theorem 4.8 requires only stability of each system matrix A . Under the stronger assumption of quadratic stability, which can be determined using Lemma 4.4, we can provide a transient guarantee on the values of the unknown function.

Theorem 4.9 (Transient values of the cost function)

Let $\Sigma = \mathcal{Z}(M)$ and $\Theta = \mathcal{Z}(N)$, where the corresponding noise models satisfy Assumption 1. Let $P > 0$ such that $APA^\top < P$ for all $(A, B)^\top \in \Sigma$ and assume ϕ^θ is Lipschitz with constant L for all $\theta \in \Theta$. Let $\delta > 0$ and u^* be a fixed input such that $\|\hat{\phi}((I - \hat{A})^{-1} \hat{B}u^*)\|_2 \leq \delta$. Given an initial condition $\hat{x}_0 \in \mathbb{R}^n$, let \hat{x}_k denote the trajectory of the unknown system (\hat{A}, \hat{B}) corresponding to the input u^* and starting from \hat{x}_0 . Then, there exists $\lambda \in [0, 1)$ such that

$$\|\hat{\phi}(\hat{x}_k)\|_2 \leq \delta + \lambda^k L \frac{\lambda_{\max}(P)}{\lambda_{\min}(P)} \|\hat{x}_0 - x^*\|_2.$$

PROOF. Let x^* be the fixed point corresponding to the input u^* for the system (\hat{A}, \hat{B}) . If $APA^\top < P$ for all $(A, B)^\top \in \Sigma$, then the set $\{A \mid (A, B)^\top \in \Sigma\}$ is bounded. Since the set is closed by definition, we can conclude that $APA^\top < P$ for all $(A, B)^\top \in \Sigma$ if and only if there exists $\lambda \in [0, 1)$ such that $APA^\top < \lambda P$ for all $(A, B)^\top \in \Sigma$. In particular, this implies that the *affine* system resulting from the application of input u^* to (\hat{A}, \hat{B}) is strongly contracting (cf. Remark 4.5). Thus, for any $k \geq 1$,

$$\|\hat{x}_k - x^*\|_{2, P^{-1/2}} \leq \lambda \|\hat{x}_{k-1} - x^*\|_{2, P^{-1/2}}.$$

By applying this repeatedly, we conclude $\|\hat{x}_k - x^*\|_{2, P^{-1/2}} \leq \lambda^k \|\hat{x}_0 - x^*\|_{2, P^{-1/2}}$, and thus

$$\|\hat{x}_k - x^*\|_2 \leq \lambda^k \frac{\lambda_{\max}(P^{-1})}{\lambda_{\min}(P^{-1})} \|\hat{x}_0 - x^*\|_2.$$

Using the triangle inequality and the fact that $\hat{\phi}$ is Lipschitz with constant L ,

$$\|\hat{\phi}(\hat{x}_k)\|_2 \leq \|\hat{\phi}(x^*)\|_2 + L\|\hat{x}_k - x^*\|_2.$$

Combining this previous with the fact that $\lambda_{\max}(P^{-1}) = 1/\lambda_{\min}(P)$ and $\lambda_{\min}(P^{-1}) = 1/\lambda_{\max}(P)$ yields

$$\|\hat{\phi}(\hat{x}_k)\|_2 \leq \delta + \lambda^k L \frac{\lambda_{\max}(P)}{\lambda_{\min}(P)} \|\hat{x}_0 - x^*\|_2.$$

This proves the statement. ■

Remark 4.10 (Suboptimal regulation of unstable systems)

Our discussion above assumes that the unknown system and all of those consistent with the measurements are stable. When this is not the case, we can instead proceed by first testing if the data is sufficiently informative for *quadratic stabilization*, i.e., to guarantee the existence of a static feedback K such that $(A + BK)P(A + BK)^\top < P$ for all $(A, B)^\top \in \Sigma$, cf [37, Thm. 5.1]. Then, we can characterize the behavior of the fixed points of the resulting systems by considering

$$\|x - (I - A + BK)^{-1} Bu\|_2 \leq \varepsilon.$$

Adapting Lemma 4.6 and the definition of ε^- accordingly is then straightforward. This gives rise to a two-step approach: (i) determine a gain K which stabilizes all systems in Σ and (ii) minimize $g(x; \Theta) + L\varepsilon^-(x)$. Note that the choice of K among potentially many stabilizing gains influences the resulting ε^- in a nontrivial way. We leave the analysis of jointly minimizing $g(x; \Theta) + L\varepsilon^-(x)$ over all x and all possible choices K as an open problem. •

5 Online cautious optimization

Our exposition so far has considered one-shot optimization settings where, given a set of measurements, we determine regularity properties and address suboptimization of the unknown function. In this section, we consider *online* scenarios, where one can repeatedly combine optimization of the unknown function with the collection of additional measurements, aiming at reducing the optimality gap. Consistent with our exposition, this means that we are interested in finding guaranteed upper bounds on $c^\top \hat{\phi}$ on the basis of initial measurements, then refining this bound on the basis of newly collected measurements¹. We aim at doing this in a monotonic fashion, that is, in a way that has the obtained upper bounds do not increase as time progresses.

Consider repeated measurements of the form

$$Y_k = \hat{\theta}^\top \Phi_k + W_k, \text{ with } W_k^\top \in \mathcal{Z}(\Pi), \quad (34)$$

where Y_k and W_k are as in (2) and $k \in \mathbb{N}$. Let $\Theta_k := \mathcal{Z}(N_k)$ be the set of parameters which are compatible with the k^{th} set of measurements, where

$$N_k := \begin{bmatrix} I & Y_k \\ 0 & -\Phi_k \end{bmatrix} \Pi \begin{bmatrix} I & Y_k \\ 0 & -\Phi_k \end{bmatrix}^\top.$$

Our online optimization strategy is described as follows.

¹ To keep the presentation contained, we do not consider minimization of $\|\hat{\phi}\|_2$, but the results can be adapted to this case.

Algorithm 1 (Online cautious optimization) Consider an initial candidate optimizer $z_0 \in \mathbb{R}^n$, an initial set of data (Y_0, Φ_0) , and $c \in \mathbb{R}^m$. Define the initial set of compatible parameters as $\bar{\Theta}_0 = \mathcal{Z}(N_0)$, where N_0 is given as in (4). For $k \geq 1$, we alternate two steps:

(i) Employ the set $\bar{\Theta}_{k-1}$ to improve the candidate optimizer by finding z_k such that

$$g_c(z_k; \bar{\Theta}_{k-1}) \leq g_c(z_{k-1}; \bar{\Theta}_{k-1}). \quad (35)$$

(ii) Measure the function as in (34) and update the set of parameters by finding a set $\bar{\Theta}_k$ such that $\hat{\theta} \in \bar{\Theta}_k$ and

$$g_c(z_k; \bar{\Theta}_k) \leq g_c(z_k; \bar{\Theta}_{k-1}). \quad (36)$$

In the following, we enforce (36) by considering parameters consistent with *all* previous measurements,

$$\bar{\Theta}_k := \Theta_0 \cap \dots \cap \Theta_k. \quad (37)$$

Note that Algorithm 1 provides a sequence of upper bounds to true function values, on the basis of measurements. This means that after *any* number of iterations, we obtain a worst-case estimate of the function value $c^\top \hat{\phi}(z_k)$ and, as such, for the minimum of $c^\top \hat{\phi}$. To be precise, for each $k \geq 1$,

$$\min_{z \in \mathbb{R}^n} c^\top \hat{\phi}(z) \leq c^\top \hat{\phi}(z_k) \leq g_c(z_k; \bar{\Theta}_{k-1}), \quad (38)$$

and these bounds are monotonically nonincreasing

$$g_c(z_{k+1}; \bar{\Theta}_k) \leq g_c(z_k; \bar{\Theta}_{k-1}) \quad (39)$$

Therefore this sequence of upper bounds is nonincreasing and bounded below, and we can conclude that the algorithm converges. However, without further assumptions, one cannot guarantee convergence of the upper bounds to the minimal value of $c^\top \hat{\phi}$, or respectively of z_k to the global minimum of $c^\top \hat{\phi}$. This is because, even though when repeatedly collecting measurements, it is reasonable to assume that the set of consistent parameters would decrease in size, this is not necessarily the case in general. In particular, a situation might arise where repeated measurements corresponding to a worst-case, or adversarial, noise signal give rise to convergence to a fixed bound with nonzero uncertainty. This is the point we address next by considering random noise realizations.

5.1 Collection of uniformly distributed measurements

We consider a scenario where the noise samples are not only bounded but distributed uniformly over the set $\mathcal{Z}(\Pi)$. To formalize this, let Π be such that $\mathcal{Z}(\Pi)$ is bounded. Consider a measure μ , and probability distribution p ,

$$p(W) := \begin{cases} \frac{1}{\mu(\mathcal{Z}(\Pi))} & W^\top \in \mathcal{Z}(\Pi), \\ 0 & \text{otherwise.} \end{cases}$$

As a notational shorthand, we write $W^\top \sim \text{Uni}(\mathcal{Z}(\Pi))$. The following result shows that, under such uniformly distributed noise samples, the size of the set of parameters indeed shrinks with increasing k .

Theorem 5.1 (Repeated measurements leading to shrinking set of consistent parameters) Suppose that the measurements (Y_k, Φ_k) are collected such that $W_k^\top \sim \text{Uni}(\mathcal{Z}(\Pi))$ and $\sigma_-(\Phi_k(z)) \geq a$ for some $a \in \mathbb{R}_{>0}$ and all $z \in \mathbb{R}^n$. Then, for any probability $0 < \pi < 1$ and any $\varepsilon > 0$,

there exists k such that if $B_\varepsilon(\hat{\theta})$ denotes the ball of radius $\varepsilon > 0$ centered at $\hat{\theta}$ in $\mathbb{R}^{k \times m}$, then

$$p(\Theta_0 \cap \dots \cap \Theta_{k-1} \subseteq B_\varepsilon(\hat{\theta})) > \pi. \quad (40)$$

PROOF. Let $k \geq 0$ and consider the set $\Theta_k = \mathcal{Z}(N_k)$. In terms of the true parameter $\hat{\theta}$, we obtain $Y_k = \hat{\theta}^\top \Phi_k + W_k$, with $W_k^\top \sim \text{Uni}(\mathcal{Z}(\Pi))$. This implies that

$$Y_k^\top \sim \text{Uni}(\mathcal{Z}(\Pi) + \Phi_k^\top \hat{\theta}).$$

For $\theta \in \mathbb{R}^k$, we know that $\theta \in \Theta_k$ if and only if $Y_k^\top - \Phi_k^\top \theta \in \mathcal{Z}(\Pi)$. Therefore, $Y_k^\top - \Phi_k^\top \theta \sim \text{Uni}(\mathcal{Z}(\Pi) + \Phi_k^\top (\hat{\theta} - \theta))$. We can now calculate the probability of $\theta \in \Theta_k$ to be

$$p(\theta \in \Theta_k) = \frac{\mu((\Phi_k^\top (\hat{\theta} - \theta) + \mathcal{Z}(\Pi)) \cap \mathcal{Z}(\Pi))}{\mu(\mathcal{Z}(\Pi))}.$$

Recall, however, that the true parameter $\hat{\theta}$ is unknown. Let $\theta \notin B_\varepsilon(\hat{\theta})$, then, since by assumption $\sigma_-(\Phi_k) \geq a$, we can derive that

$$p(\theta \in \Theta_k) \leq \max_{\|V\|_2=1} \frac{\mu((a\varepsilon V + \mathcal{Z}(\Pi)) \cap \mathcal{Z}(\Pi))}{\mu(\mathcal{Z}(\Pi))} =: \eta(a\varepsilon).$$

Note that, since $\mathcal{Z}(\Pi)$ is bounded and convex, we have in addition that $\eta(\lambda) < 1$ for $\lambda > 0$. This means that, regardless of the values of the measurements, the probability that $\theta \in \Theta_k$ is strictly less than 1. Now, by repeating such measurements, we see that for $\theta \notin B_\varepsilon(\hat{\theta})$, we have

$$p(\theta \in \Theta_0 \cap \dots \cap \Theta_{k-1}) = p(\theta \in \Theta_0) \cdots p(\theta \in \Theta_{k-1}) \leq \eta(a\varepsilon)^k.$$

From this, if k is large enough, then (40) holds. \blacksquare

Remark 5.2 (Extension to other probability distributions) Note that the proof of Theorem 5.1 essentially depends on the fact that $p(W) \neq 0$ for all $W^\top \in \mathcal{Z}(\Pi)$. Therefore, we can draw the same conclusion for other probability distributions that satisfy this assumption, such as truncated Gaussian distributions. \bullet

Thus, if the noise is uniformly random, Theorem 5.1 guarantees that the set of parameters converges. This allows us to conclude that, under Algorithm 1, the sequence $g_c(z_k; \bar{\Theta}_k)$ converges to the sequence $c^\top \hat{\phi}(z_k)$ pointwise.

5.2 Methods to improve the candidate optimizer

In this section we discuss methods to guarantee that (35) holds. A direct way of enforcing this is by simply optimizing the expression, that is, updating z_k as

$$z_k := \arg \min_{z \in \mathbb{R}^n} g_c(z; \bar{\Theta}_{k-1}). \quad (41)$$

However, this requires us to find a global optimizer of the function $g_c(z; \bar{\Theta}_{k-1})$, which might be difficult to obtain depending on the scenario. As an alternative, we devise a procedure where the property (35) is guaranteed by *locally* updating the candidate optimizer along with collecting new measurements *near* the candidate optimizer. We show that, using only local information in this manner, the algorithm converges to the true optimizer under appropriate technical assumptions.

In order to formalize this notion of ‘local’, we both measure and optimize on a polyhedral set around the candidate optimizer. For this, let $\mathcal{F} = \{f_i\}_{i=1}^T \subseteq \mathbb{R}^n$ be a finite set such

that $0 \in \text{int}(\text{conv } \mathcal{F})$. Let $\mathcal{S}(z) := z + \text{conv } \mathcal{F}$. Then, (35) holds if we update the optimizer by

$$z_k := \arg \min_{z \in \mathcal{S}(z_{k-1})} g_c(z; \Theta_0 \cap \dots \cap \Theta_{k-1}). \quad (42)$$

We measure the function at all points in $z_k + \mathcal{F}$, that is, we take $\Phi_k = \Phi^{\mathcal{F}}(z_k)$, where

$$\Phi^{\mathcal{F}}(z) := \begin{bmatrix} \phi_1(z + f_1) & \dots & \phi_1(z + f_T) \\ \vdots & & \vdots \\ \phi_k(z + f_1) & \dots & \phi_k(z + f_T) \end{bmatrix}.$$

The online optimization procedure then incrementally incorporates these measurements to refine the computation of the candidate optimizer. The following result investigates the properties of the resulting *online local descent*.

Theorem 5.3 (Online local descent) *Let \mathcal{F} be a finite set such that $0 \in \text{int}(\text{conv } \mathcal{F})$. Let $\mathcal{S}(z) := z + \text{conv } \mathcal{F}$. Consider Algorithm 1, employing local update rules (42) and (37), where $\theta_k = \mathcal{Z}(N_k)$ and $\Phi_k = \Phi^{\mathcal{F}}(z_k)$. Suppose that the initial point $z_0 \in \mathbb{R}^n$ and the data $(Y_0, \Phi^{\mathcal{F}}(z_0))$ are such that $c^\top \phi^\theta$ is strictly convex for all $\theta \in \Theta_0$. Then the following hold:*

- (i) For any $k \geq 1$, the problem (42) is strictly convex;
- (ii) If $z_k \neq z_{k+1}$, then

$$g_c(z_{k+1}; \bar{\Theta}_k) < g_c(z_k; \bar{\Theta}_{k-1}),$$

that is, (39) holds with a strict inequality.

PROOF. Statement (i) follows from Corollary 3.12. Then, note that if $z_k \neq z_{k+1}$, we have that

$$g_c(z_{k+1}; \Theta_0 \cap \dots \cap \Theta_k) < g_c(z_k; \Theta_0 \cap \dots \cap \Theta_k),$$

thus proving (ii). \blacksquare

Remark 5.4 (Reduction of the complexity) Most optimization schemes require a large number of function evaluations in order to find an optimizer of e.g., (42). In turn, finding values of $g_c(\cdot; \bar{\Theta}_k)$ requires us to resolve another optimization problem. This nested nature of the problem means that reducing the number of evaluations can speed up computation significantly. This can be done by relaxing (42) as follows. Assume a Lipschitz constant of $c^\top \hat{\phi}$ is known or obtained from measurements. Since \mathcal{F} is a finite set, there exists $\nu \in \mathbb{R}_{\geq 0}$ such that each $z \in \mathcal{S}(z_{k-1})$ is at most a distance ν from a point in $z_{k-1} + \mathcal{F}$. As such, we can instead find

$$z_k := \arg \min_{z \in z_{k-1} + \mathcal{F}} g_c(z; \Theta_0 \cap \dots \cap \Theta_{k-1}),$$

and employ the value of ν and the Lipschitz constant to approximate the optimal value of (42). \bullet

Recall that, without making further assumptions, we can only guarantee that the algorithm converges, but not that it converges to the true optimizer. Moreover, even when the assumptions of Theorem 5.1 hold, we do not yet have a way of concluding that the algorithm has converged. The following result employs the uncertainty near the optimizer, to provide a criterion for convergence.

Lemma 5.5 (Stopping criterion) *Let Θ be compact and $\mathcal{S} \subseteq \mathbb{R}^n$ closed. Define*

$$\bar{z} := \arg \min_{z \in \mathcal{S}} g_c(z; \Theta), \quad \hat{z} := \arg \min_{z \in \mathcal{S}} c^\top \hat{\phi}(z).$$

Then $g_c(\bar{z}; \Theta) \geq c^\top \hat{\phi}(\hat{z}) \geq g_c(\hat{z}; \Theta) - 2 \max_{z \in \mathcal{S}} U_c(z; \Theta)$.

PROOF. For the first inequality, note that by definition,

$$g_c(\bar{z}; \Theta) \geq c^\top \hat{\phi}(\bar{z}) \geq c^\top \hat{\phi}(\hat{z}).$$

To show the second inequality, note that $g_c(\bar{z}; \Theta) \leq g_c(\hat{z}; \Theta)$, and

$$2 \max_{z \in \mathcal{S}} U_c(z; \Theta) \geq g_c(\hat{z}; \Theta) - c^\top \hat{\phi}(\hat{z}; \Theta).$$

Combining these, we obtain

$$g_c(\bar{z}; \Theta) - 2 \max_{z \in \mathcal{S}} U_c(z; \Theta) \leq c^\top \hat{\phi}(\hat{z}; \Theta),$$

which proves the result. \blacksquare

Importantly, we note that the upper and lower bounds in Lemma 5.5 can be computed purely in terms of data.

If the maximal uncertainty on the set \mathcal{S} is equal to zero, then the local minima of g_c and $c^\top \hat{\phi}$ coincide. In addition, if $c^\top \hat{\phi}$ is strictly convex, we have that any local minimum in the interior of \mathcal{S} is a global minimum.

Corollary 5.6 (Uncertainty under repetition) *Under the assumptions of Theorem 5.3, suppose in addition that for $k \geq 1$, the measurements $(Y_k, \Phi^{\mathcal{F}}(z_k))$ are collected such that $W_k^\top \sim \text{Uni}(\mathcal{Z}(\Pi))$ and $\sigma_-(\Phi^{\mathcal{F}}(z_k)) \geq a$ for some $a \in \mathbb{R}_{>0}$ and all k . Then, for any $z \in \mathbb{R}^n$, the expected value of the uncertainty monotonically converges to 0, that is,*

$$U_c(z; \Theta_0 \cap \dots \cap \Theta_{k-1}) \geq U_c(z; \Theta_0 \cap \dots \cap \Theta_k),$$

and $\lim_{k \rightarrow \infty} \mathbb{E}(U_c(z; \Theta_0 \cap \dots \cap \Theta_k)) = 0$.

PROOF. The monotonic decrease of the uncertainty readily follows from its definition. If $\Theta_0 \cap \dots \cap \Theta_{k-1} \subseteq B_\epsilon(\hat{\theta})$, then it is immediate from the definitions that

$$U_c(z; \Theta_0 \cap \dots \cap \Theta_{k-1}) \leq U_c(z; B_\epsilon(\hat{\theta})).$$

Given that the map $\theta \mapsto c^\top \phi^\theta(z)$ is linear in θ , the right-hand side can be seen to converge to 0 for $\epsilon \rightarrow 0$. Combining these pieces with Theorem 5.1 proves the statement. \blacksquare

As a consequence of Lemma 5.5 and Corollary 5.6, we conclude that the expected difference between the optimal value $\min_{z \in \mathcal{S}} c^\top \hat{\phi}(z)$ of the unknown function and the optimal value $\min_{z \in \mathcal{S}} g_c(z; \Theta_0 \cap \dots \cap \Theta_k)$ provided by online local descent both converge to zero.

6 Simulation examples

Here we provide two simulation examples to illustrate the proposed framework and our results. We consider an application to data-based contraction analysis based on Section 4.1, and the data-based suboptimal regulation of an unmanned aerial vehicle, using Section 4.2. We refer the

interested reader to our conference paper [1] for an illustration of a simplified version of online cautious suboptimization, cf. Section 5.

6.1 Application to data-based contraction analysis

Consider continuous-time systems $\dot{z} = \phi^\theta(z)$, where $z \in \mathbb{R}^2$, the parameter $\theta \in \mathbb{R}^{6 \times 2}$, and with basis functions

$$\begin{aligned} \phi_1(z) &= z_1, & \phi_3(z) &= \sin(z_1) - 1, & \phi_5(z) &= \cos(z_1), \\ \phi_2(z) &= z_2, & \phi_4(z) &= \sin(z_2) - 1, & \phi_6(z) &= \cos(z_2). \end{aligned}$$

These define the vector-function $b : \mathbb{R}^2 \rightarrow \mathbb{R}^6$, cf. (1). Suppose that the parameter of the unknown function $\hat{\phi}$ is

$$\hat{\theta}^\top = \begin{bmatrix} -6 & 1 & -1 & 1 & -1 & 1 \\ 0 & -6 & 1 & -1 & 1 & -1 \end{bmatrix}.$$

We are interested in determining whether the system $\dot{z} = \hat{\phi}(z)$ is strictly contracting with respect to the unweighted 2–norm on the basis of noisy measurements.

As a first step, note that $\|J(b)\|_2 = \sqrt{2}$. Then writing $\hat{\phi}(z) = -6z + h(z)$, we can conclude that

$$\begin{aligned} \text{osLip}(\hat{\phi}) &\leq -6 + \text{osLip}(h) \leq -6 + \|J(h)\|_2 \\ &\leq -6 + \left\| \begin{bmatrix} 0 & 1 & -1 & 1 & -1 & 1 \\ 0 & 0 & 1 & -1 & 1 & -1 \end{bmatrix} \right\|_2 \cdot \|J(b)\|_2 \quad (43) \\ &\approx -1.8694. \end{aligned}$$

Thus, the system $\dot{z} = \hat{\phi}(z)$ is indeed strictly contracting.

However, as before we assume that we do not have access to the value of $\hat{\theta}$ and can only assess contractivity based on measurements of $\hat{\phi}$. To be precise, we assume that we have access to samples of a single continuous-time trajectory starting from $z(0) = (10, -20)^\top$ at the time instances

$$t \in \{0, 0.01, \dots, 0.25\}.$$

The noise is assumed to satisfy $WW^\top \leq 10 \cdot I_2$, giving rise to a set $\Theta = \mathcal{Z}(N)$. Two examples of parameters compatible with the measurements are given by

$$\begin{bmatrix} -5.2588 & 1.3807 & -1.4810 & 0.8087 & -0.8825 & 1.5444 \\ -0.2166 & -6.1490 & 1.1110 & -0.5087 & 0.4157 & -0.6826 \\ -6.3111 & 0.8091 & -0.6433 & 1.0725 & -0.5355 & 0.9938 \\ 0.7750 & -5.6377 & 0.9785 & -1.1375 & 1.3976 & -0.8905 \end{bmatrix},$$

To illustrate the different systems compatible with the measurements, Figure 2 shows the trajectories emanating from the origin for a number of such systems.

In order to guarantee strict contraction in terms of the measurements, we employ the condition of Corollary 4.2. With these measurements Φ has full row rank and thus $\Theta = \mathcal{Z}(N)$ is compact. For the last assumption, from the fact that $\|J(b)\|_2 = \sqrt{2}$, we obtain that for any $z, z^* \in \mathbb{R}^2$,

$$\|b(z) - b(z^*)\|_{2, (-N_{22})^{-1/2}} \leq \frac{\sqrt{2}}{\sqrt{-\lambda_{\max}(N_{22})}} \|z - z^*\|_2.$$

This means that we can use Corollary 4.2, which shows that the unknown system is strictly contracting if




Figure 2. Trajectories of $\dot{z} = \phi^\theta(z)$ from the origin for 100 randomly generated realization of the parameter $\theta \in \Theta$. The black dots indicate the resulting equilibria.

$$\text{osLip}(\phi^{\text{lse}}) < -\frac{\sqrt{2\lambda_{\max}(N|N_{22})}}{\sqrt{-\lambda_{\max}(N_{22})}} \approx -1.78.$$

Using a line of reasoning as in (43) allows us to conclude that this indeed holds. As such, the system $\dot{z} = \phi^\theta$ is strictly contracting for all $\theta \in \Theta$. We plot a number of trajectories of the systems corresponding to 4 different realizations of the parameter θ in Figure 3. It can be seen that each is strictly contracting.

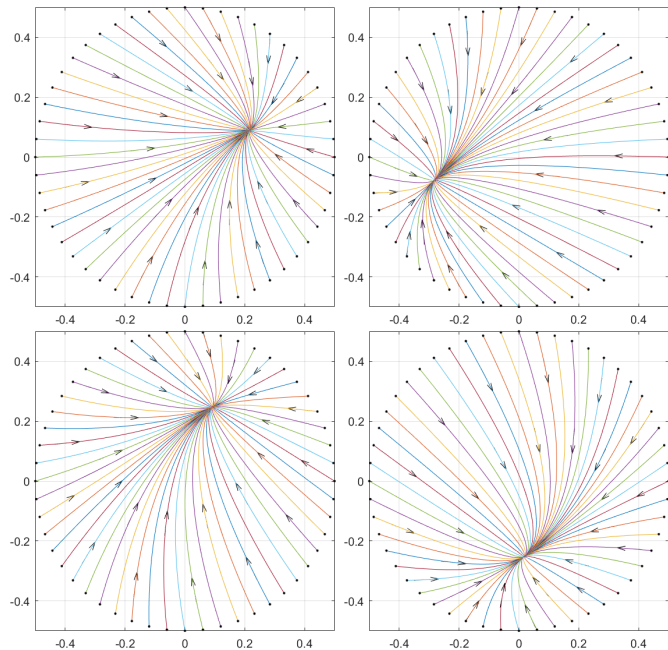


Figure 3. Trajectories of $\dot{z} = \phi^\theta(z)$ starting from a number of different initial conditions on the circle with radius $\frac{1}{2}$ for four different realizations of $\theta \in \Theta$.

6.2 Suboptimal regulation of unmanned aerial vehicle

This example considers the suboptimal regulation of a fixed-wing unmanned aerial vehicle (UAV). We consider a model of the longitudinal motion for which the relevant equations of motion are derived in e.g., [42, Ch. 2]. For

the values of the parameters and the linearization we follow [43], which derives a benchmark model on the basis of 10 different models of real-world UAV's of different specifications. The derived system is a state-space model

$$\dot{x}(t) = A_{\text{cont}}x(t) + B_{\text{cont}}u(t),$$

given in [43, Eq. (8)-(10)], with states $x(t) \in \mathbb{R}^4$ and inputs $u(t) \in \mathbb{R}$. Here, $x_1(t)$ denotes the *deviation from the nominal forward velocity*, $x_2(t)$ the *deviation from the nominal angle of attack*, $x_3(t)$ the *pitch angle*, and $x_4(t)$ the *pitch rate*. The input $u(t)$ represents the *elevator deflection*. We use the parameters derived in [43, Eq. (13)-(14)], and thus

$$A_{\text{cont}} = \begin{bmatrix} -0.240 & 0.345 & -0.411 & 0 \\ -1.905 & -10.695 & 0 & 0.941 \\ 0 & 0 & 0 & 1 \\ 0.457 & -250.513 & 0 & -8.844 \end{bmatrix}, B_{\text{cont}} = \begin{bmatrix} 0 \\ -0.301 \\ 0 \\ -98.658 \end{bmatrix}.$$

We consider a discretization of this model with stepsize $\frac{1}{20}$, thus arriving at a system of the form (30), where $\hat{A} = I_4 + \frac{1}{20}A_{\text{cont}}$ and $\hat{B} = \frac{1}{20}B_{\text{cont}}$. These matrices are unknown to us, and instead we have access to a single second-long trajectory of measurements of the state $x(t)$ with $x(0) = (1 \ 1 \ 1 \ 1)^\top$, and fixed input $u(t) = -4$. Thus, $T = 20$. Denoting the matrix collecting the noise samples by W_s , we assume that the noise on the measurements is such that $W_s W_s^\top \leq 10^{-4}I_4$. These measurements give rise to a set of systems consistent with the data given by $\Sigma = \mathcal{Z}(M)$.

Following Section 4.2, we aim at finding suboptimal values of the norm of an unknown cost function. In this example, we consider a simple cost function $\hat{\phi} : \mathbb{R}^4 \rightarrow \mathbb{R}^4$ of the form $\hat{\phi}(x) = x - \hat{x}$, where $\hat{x} = (1 \ 0 \ 0 \ 0)^\top$. This situation might arise, for instance, when the UAV is tasked to mirror the orientation and speed of another UAV on the basis of noisy measurements. If the cost function were known, the problem of minimizing $\|\hat{\phi}\|_2$ would be to simply regulate x as close to \hat{x} as possible. However, we assume that we only know that $\hat{\phi}$ is affine in x , and thus of the form

$$\phi^\theta(x) := \theta^\top \begin{pmatrix} 1 \\ x \end{pmatrix}, \text{ where } \hat{\theta} = \begin{bmatrix} -\hat{x}^\top \\ I_4 \end{bmatrix}.$$

We assume that the measurements available to us are collected concurrently with the collection of the measurements of the unknown system. In line with Section 2, this means that we measure at $z_i := x(i-1)$ for $i = 1, \dots, T = 20$. We assume a noise model of the form $WW^\top \leq I_4$, and obtain a set of parameters consistent with the measurements given by $\Theta = \mathcal{Z}(N)$.

We are interested in finding upper bounds for the quantity

$$\min_u \|\hat{\phi}((I - \hat{A})^{-1}\hat{B}u)\|_2.$$

Before applying Theorem 4.8, we need to check its assumptions. First, we employ Lemma 4.4 in order to check whether each A consistent with the measurements is stable. Solving the LMI (31) yields

$$\bar{P} \approx \begin{bmatrix} 0.0072 & 0.0015 & 0.0004 & -0.0102 \\ 0.0015 & 0.1665 & 0.1218 & 0.1496 \\ 0.0004 & 0.1218 & 0.1648 & -1.3489 \\ -0.0102 & 0.1496 & -1.3489 & 45.00437 \end{bmatrix}.$$

Next, we employ Lemma 3.16 to obtain a Lipschitz constant. Given that the basis functions are linear, we can easily minimize L over (25) to obtain $L = 2.0171$. Thus, L is a Lipschitz constant for ϕ^θ for all $\theta \in \mathcal{Z}(N)$.

Therefore, the required pieces are in place, and we can employ Theorem 4.8 to show that

$$\min_u \|\hat{\phi}((I - \hat{A})^{-1}\hat{B}u)\|_2 \leq 2.3347. \quad (44)$$

This value is attained for the input $u^* = 0.2321$ and the minimizer in the state is $x = (0.809 \ 0.054 \ -0.355 \ -0.006)^\top$.

Figure 4 shows the results of applying this input to different realizations of the system matrices compatible with the measurements. We make the following observations. First, one can see that each of the systems converges to a fixed point, and each of these fixed points are within a distance $\varepsilon^-(x) = 0.4491$ of x . Second, for certain systems the transients, especially in states x_1 and x_2 , converge quite slowly. This holds because the true system is marginally stable and the same holds for elements of set $\mathcal{Z}(M)$ (to be precise, the upper bound of Theorem 4.9 can only be shown to hold for $0.999 < \lambda < 1$).

Corresponding to these trajectories, Figure 5 shows the resulting values of $\|\phi^\theta(x_k)\|_2$, for both the true value and for different realizations of $\theta \in \mathcal{Z}(N)$. The plot reveals that, with the true parameter $\hat{\theta}$, the values at the fixed points are well below the bound (44). In fact, resolving the problem with known $\hat{\phi}$ yields $u^* = 0.1820$ and an upper bound of 0.7963. Indeed, as the left plot shows, the unknown nature of $\hat{\phi}$ has a large effect and therefore has to be taken into account when providing formal guarantees.

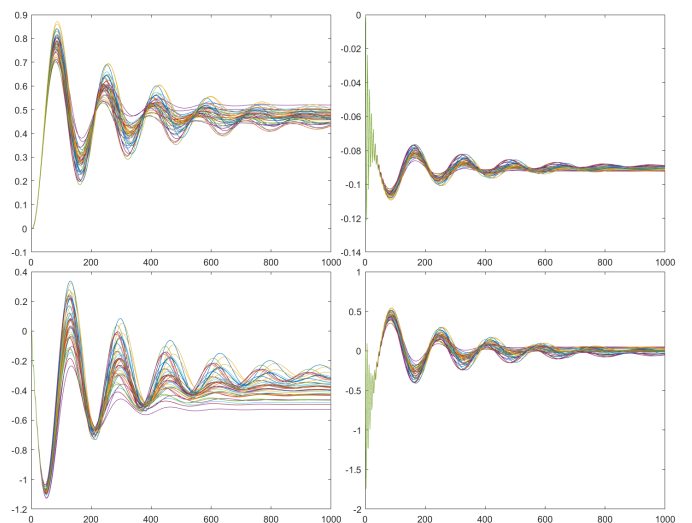


Figure 4. State trajectory corresponding to the found input $u^* = 0.2321$ for 40 different realizations of the matrices (A, B) consistent with the measurements.




Figure 5. Values of $\|\phi^\theta(x_k)\|_2$ for the state trajectories of Figure 4. On the left, we pick for each system a different realization of $\theta \in \mathcal{Z}(N)$. On the right, we use the true value of $\hat{\theta}$.

7 Conclusions

We have developed a set-valued regression approach to data-based optimization of an unknown vector-valued function. Taking a worst-case perspective, we have provided a range of guarantees in terms of measurements on the minimization of least conservative upper bounds of both the norm and any linear combination of the components of the unknown function. Our analysis has yielded closed-form expressions for the proposed upper bounds, conditions to ensure their convexity (as an enabler for the use of gradient-based methods for their optimization), and Lipschitzness characterizations to facilitate simplifications based on interpolation. We have illustrated the applicability of the proposed cautious optimization approach in systems and control scenarios with unknown dynamics: first, by providing conditions on the data which guarantee that the dynamics is strongly contracting and, second, by providing conditions under which we can regulate an unknown system to a point with a guaranteed maximal value of an unknown cost function. Finally, we have considered online scenarios, where one repeatedly combines the optimization of the unknown function with the collection of additional data. Under mild assumptions, we show that repeated measurements lead to a strictly shrinking set of compatible parameters and we build on this result to design an online procedure that provides sequential upper bounds which converge to a true optimizer.

Future work will investigate conditions to ensure other regularity properties of the unknown function and its least conservative upper bounds, the impact of the choice of basis functions on the guarantees and computational efficiency of the proposed cautious optimization methods (particularly, the use of polynomial bases coupled with sum-of-squares optimization techniques), and the application of the results to the certification of Lyapunov stability and controller design for nonlinear systems.

References

- [1] J. Eising and J. Cortés, “Set-valued regression and cautious suboptimization: from noisy data to optimality,” in *IEEE Conf. on Decision and Control*, Singapore, 2023, submitted.
- [2] A. Ben-Tal, L. E. Ghaoui, and A. Nemirovski, *Robust Optimization*, ser. Applied Mathematics Series. Princeton, NJ: Princeton University Press, 2009.
- [3] D. Bertsimas, D. B. Brown, and C. Caramanis, “Theory and applications of robust optimization,” *SIAM Review*, vol. 53, no. 3, pp. 464–501, 2011.
- [4] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. The MIT Press, Nov. 2005.
- [5] G. Cybenko, “Dynamic load balancing for distributed memory multiprocessors,” *Journal of Parallel and Distributed*

Computing, vol. 7, no. 2, pp. 279–301, 1989.

- [6] P. Tabuada and B. Ghahserifard, “Universal approximation power of deep residual neural networks through the lens of control,” *IEEE Transactions on Automatic Control*, vol. 68, no. 5, pp. 2715–2728, 2023.
- [7] M. Milanese and A. Vicino, “Optimal estimation theory for dynamic systems with set membership uncertainty: An overview,” *Automatica*, vol. 27, no. 6, pp. 997–1009, 1991.
- [8] M. Milanese and C. Novara, “Set membership estimation of nonlinear regressions,” in *IFAC World Congress*, vol. 35, no. 1, Barcelona, Spain, 2002, pp. 7–12.
- [9] J. Calliess, S. J. Roberts, C. E. Rasmussen, and J. Maciejowski, “Lazily adapted constant kinky inference for nonparametric regression and model-reference adaptive control,” *Automatica*, vol. 122, p. 109216, 2020.
- [10] M. Budišić, R. Mohr, and I. Mezić, “Applied Koopmanism,” *Chaos*, vol. 22, no. 4, p. 047510, 2012.
- [11] M. Korda and I. Mezić, “Optimal construction of Koopman eigenfunctions for prediction and control,” *IEEE Transactions on Automatic Control*, vol. 65, no. 12, pp. 5114–5129, 2020.
- [12] M. Haseli and J. Cortés, “Generalizing dynamic mode decomposition: balancing accuracy and expressiveness in Koopman approximations,” *Automatica*, vol. 153, p. 111001, 2023.
- [13] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed., ser. Springer Series in Statistics. New York: Springer, 2013.
- [14] R. Tibshirani, “Regression shrinkage and selection via the Lasso,” *Journal of the Royal Statistical Society. Series B*, vol. 58, no. 1, pp. 267–288, 1996.
- [15] S. L. Brunton, J. L. Proctor, and J. N. Kutz, “Discovering governing equations from data by sparse identification of nonlinear dynamical systems,” *Proceedings of the National Academy of Sciences*, vol. 113, no. 15, pp. 3932–3937, 2016.
- [16] P. J. Schmid, “Dynamic mode decomposition of numerical and experimental data,” *Journal of Fluid Mechanics*, vol. 656, pp. 5–28, 2010.
- [17] —, “Dynamic mode decomposition and its variants,” *Annual Review of Fluid Mechanics*, vol. 54, no. 1, pp. 225–254, 2022.
- [18] J. N. Kutz, S. L. Brunton, B. W. Brunton, and J. L. Proctor, *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems*, ser. Other Titles in Applied Mathematics. Philadelphia, PA: SIAM, 2016, vol. 149.
- [19] M. R. Jovanović, P. J. Schmid, and J. W. Nichols, “Sparsity-promoting dynamic mode decomposition,” *Physics of Fluids*, vol. 26, no. 2, p. 024103, 2014.
- [20] J. C. Willems, P. Rapisarda, I. Markovsky, and B. L. M. De Moor, “A note on persistency of excitation,” *Systems & Control Letters*, vol. 54, no. 4, pp. 325–329, 2005.
- [21] H. J. van Waarde, J. Eising, H. L. Trentelman, and M. K. Camlibel, “Data informativity: a new perspective on data-driven analysis and control,” *IEEE Transactions on Automatic Control*, vol. 65, no. 11, pp. 4753–4768, 2020.
- [22] I. Markovsky, “Data-driven simulation of generalized bilinear systems via linear time-invariant embedding,” *IEEE Transactions on Automatic Control*, vol. 68, no. 2, pp. 1101–1106, 2022.
- [23] M. Guo, C. D. Persis, and P. Tesi, “Data-driven stabilization of nonlinear polynomial systems with noisy data,” *IEEE Transactions on Automatic Control*, vol. 67, no. 8, pp. 4210–4217, 2021.
- [24] C. De Persis, M. Rotulo, and P. Tesi, “Learning controllers from data via approximate nonlinearity cancellation,” *IEEE Transactions on Automatic Control*, 2023, to appear.
- [25] V. G. Lopez and M. A. Müller, “On a continuous-time version of Willems’ lemma,” *arXiv preprint arxiv:2203.03702*, 2022.

- [26] P. Rapisarda, M. K. Camlibel, and H. van Waarde, “A persistency of excitation condition for continuous-time systems,” *IEEE Control Systems Letters*, vol. 7, pp. 589–594, 2022.
- [27] J. Eising and J. Cortés, “When sampling works in data-driven control: informativity for stabilization in continuous time,” *IEEE Transactions on Automatic Control*, 2023, submitted. Available at <https://arxiv.org/abs/2301.10873>.
- [28] W. Lohmiller and J.-J. E. Slotine, “On contraction analysis for non-linear systems,” *Automatica*, vol. 34, no. 6, pp. 683–696, 1998.
- [29] F. Bullo, *Contraction Theory for Dynamical Systems*, 1.1 ed. Kindle Direct Publishing, 2023.
- [30] F. Dörfler, P. Tesi, and C. De Persis, “On the certainty-equivalence approach to direct data-driven LQR design,” *IEEE Transactions on Automatic Control*, 2023, to appear.
- [31] G. Bianchin, M. Vaquero, J. Cortés, and E. Dall’Anese, “Online stochastic optimization of unknown linear dynamical systems: data-driven controller synthesis and analysis,” *IEEE Transactions on Automatic Control*, 2021, submitted.
- [32] A. Hauswirth, S. Bolognani, G. Hug, and F. Dörfler, “Optimization algorithms as robust feedback controllers,” *arXiv preprint arXiv:2103.11329*, 2021.
- [33] L. Cothren, G. Bianchin, and E. Dall’Anese, “Data-enabled gradient flow as feedback controller: Regulation of linear dynamical systems to minimizers of unknown functions,” in *Conference on Learning for Dynamics and Control*, ser. Proceedings of Machine Learning Research, vol. 168. PMLR, 2022, pp. 234–247.
- [34] M. Krstić and H.-H. Wang, “Stability of extremum seeking feedback for general nonlinear dynamic systems,” *Automatica*, vol. 36, no. 4, pp. 595–601, 2000.
- [35] K. B. Ariyur and M. Krstić, *Real-Time Optimization by Extremum-Seeking Control*. New York: Wiley, 2003.
- [36] A. Teel and D. Popovic, “Solving smooth and nonsmooth multivariable extremum seeking problems by the methods of nonlinear programming,” in *American Control Conference*, Arlington, VA, 2001, pp. 2394–2399.
- [37] H. J. van Waarde, M. K. Camlibel, J. Eising, and H. L. Trentelman, “Quadratic matrix inequalities with applications to data-based control,” *arXiv preprint arXiv:2203.12959*, 2022.
- [38] I. Pólik and T. Terlaky, “A survey of the S-Lemma,” *SIAM Review*, vol. 49, no. 3, pp. 371–418, 2007.
- [39] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, UK: Cambridge University Press, 2009.
- [40] A. Davydov, S. Jafarpour, and F. Bullo, “Non-Euclidean contraction theory for robust nonlinear stability,” *IEEE Transactions on Automatic Control*, vol. 67, no. 12, pp. 6667–6681, 2022.
- [41] H. J. van Waarde, M. K. Camlibel, and H. L. Trentelman, “Data-driven analysis and design beyond common Lyapunov functions,” in *IEEE Conf. on Decision and Control*, 2022, pp. 2783–2788.
- [42] J. H. Blakelock, *Automatic Control of Aircraft and Missiles*. John Wiley & Sons, 1991.
- [43] E. Bertran, P. Tercero, and A. Sànchez-Cerdà, “UAV generalized longitudinal model for autopilot controller designs,” *Aircraft Engineering and Aerospace Technology*, vol. 94, no. 3, pp. 380–391, 2022.