

Data-Enabled Gradient Flow as Feedback Controller: Regulation of Linear Dynamical Systems to Minimizers of Unknown Functions

Liliaokeawawa Cothren

LILIAOKEAWAWA.COTHREN@COLORADO.EDU

Gianluca Bianchin

GIANLUCA.BIANCHIN@COLORADO.EDU

Emiliano Dall’Anese

EMILIANO.DALLANESE@COLORADO.EDU

Department of Electrical, Computer, and Energy Engineering, University of Colorado Boulder

Editors: R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, M. Kochenderfer

Abstract

This paper considers the problem of regulating a linear dynamical system to the solution of a convex optimization problem with an unknown or partially-known cost. We design a data-driven feedback controller – based on gradient flow dynamics – that (i) is augmented with learning methods to estimate the cost function based on infrequent (and possibly noisy) functional evaluations; and, concurrently, (ii) is designed to drive the inputs and outputs of the dynamical system to the optimizer of the problem. We derive sufficient conditions on the learning error and the controller gain to ensure that the error between the optimizer of the problem and the state of the closed-loop system is ultimately bounded; the error bound accounts for the functional estimation errors and the temporal variability of the unknown disturbance affecting the linear dynamical system. Our results directly lead to exponential input-to-state stability of the closed-loop system. The proposed method and the theoretical bounds are validated numerically.

Keywords: Learning-based control; learning-based optimization; gradient flow; output regulation.

1. Introduction

In this paper, we consider the problem of designing feedback controllers to steer the output of a linear time-invariant (LTI) dynamical system towards the solution of a convex optimization problem with unknown costs. The design of controllers inspired by optimization algorithms has received attention recently; see, e.g., [Jokic et al. \(2009\)](#); [Brunner et al. \(2012\)](#); [Lawrence et al. \(2018\)](#); [Hauswirth et al. \(2021a\)](#); [Colombino et al. \(2020\)](#); [Zheng et al. \(2020\)](#); [Bianchin et al. \(2020\)](#) and the recent survey by [Hauswirth et al. \(2021b\)](#). These methods have been utilized to solve control problems in, e.g., power systems in [Hirata et al. \(2014\)](#); [Menta et al. \(2018\)](#), transportation systems in [Bianchin et al. \(2021a\)](#), robotics in [Zheng et al. \(2020\)](#), and epidemics in [Bianchin et al. \(2021b\)](#).

A common denominator in the works mentioned above is that the cost of the optimization problem associated with the dynamical system is *known*, and first- and second-order information is easily accessible at any time. One open research question is whether controllers can be synthesized when the cost of the optimization problem is *unknown* or *partially known*. Towards this direction, in this paper we consider unconstrained convex optimization problems with unknown costs associated with the LTI dynamical systems. We investigate the design of *data-driven* feedback controllers based on online gradient flow dynamics that: (i) leverage learning methods to estimate the cost function based on infrequent (and possibly noisy) functional evaluations; and (ii) are designed to concurrently drive the inputs and outputs of the dynamical system to the optimizer of the problem within a bounded

error. Our learning procedure hinges on a basis expansion for the cost and leverages methods such as least-squares, ridge regression, and sparse linear regression; see [Hastie et al. \(2009\)](#); [Tibshirani \(1996\)](#). In addition, our results are also directly applicable to cases where residual neural networks are utilized to estimate the cost; see [Tabuada and Gharesifard \(2020\)](#).

We consider a cost that includes the sum of a loss associated with the inputs and a loss function associated with the outputs. Both costs may be unknown or parametrized by unknown parameters; we assume that functional evaluations are provided at irregular intervals, due to underlying communication or processing bottlenecks (see, for example, delays in power system metering systems in [Luan et al. \(2013\)](#) and in transportation systems in [Gündling et al. \(2020\)](#)). As an example of a function with unknown parameters associated with the outputs, take $\psi(y) = \|y - r\|^2$, where $y : \mathbb{R}_{>0} \rightarrow \mathbb{R}^p$ is the system output and $r \in \mathbb{R}^p$ is unknown to the controller. Additional examples include cases where $\psi(\cdot)$ represents a barrier function associated with unknown sets; see, e.g., [Robey et al.](#); [Taylor et al. \(2020\)](#). Regarding the function associated with the inputs, another scenario where it is unknown is when it captures objectives of users interacting with the system; see [Simonetto et al. \(2021\)](#); [Notarnicola et al. \(2021\)](#); [Fabiani et al. \(2021\)](#); [Ospina et al. \(2020\)](#); [Luo et al. \(2020\)](#). In this case, the loss models objectives such as dissatisfaction, discomfort, etc. For example, in a platooning problem, the control input is represented by the speed of the vehicles, and the loss function captures the sense of safety for drivers [van Nunen et al. \(2017\)](#). In lieu of synthetic models (that may not represent accurately the user’s objectives [Munir et al. \(2013\)](#); [Bourgin et al. \(2019\)](#)), one learns the loss based on evaluations infrequently provided by the user.

Related Works. We note that a key differentiating aspect relative to extremum seeking methods (see, e.g., [Krstic and Wang \(2000\)](#); [Ariyur and Krstić \(2003\)](#); [Teel and Popovic \(2001\)](#) and many others), the Q-learning of [Devraj and Meyn \(2017\)](#), and methods based on concurrent learning [Chowdhary and Johnson \(2010\)](#); [Chowdhary et al. \(2013\)](#); [Poveda et al. \(2021\)](#) is that we consider a setting where only sporadic functional evaluations are available (i.e., we do not have continuous access to functional evaluations). Regarding the problem of regulating LTI systems towards solutions of optimization problems, existing approaches leveraged gradient flows in [Menta et al. \(2018\)](#); [Bianchin et al. \(2020\)](#), proximal-methods in [Colombino et al. \(2020\)](#), saddle-flows in [Brunner et al. \(2012\)](#), prediction-correction methods in [Zheng et al. \(2020\)](#), and the hybrid accelerated methods proposed in [Bianchin et al. \(2020\)](#). Plants with (smooth) nonlinear dynamics were considered in [Brunner et al. \(2012\)](#); [Hauswirth et al. \(2021a\)](#), and switched LTI systems in [Bianchin et al. \(2021a\)](#). A joint stabilization and regulation problem was considered in [Lawrence et al. \(2021, 2018\)](#). See also the recent survey by [Hauswirth et al. \(2021b\)](#). In all these works, the cost function is assumed to be known; here, we tackle the problem of jointly learning the cost and performing the regulation task.

Our setup is aligned with [Simonetto et al. \(2021\)](#); [Ospina et al. \(2020\)](#), where Gaussian Processes are utilized to learn cost functions based on infrequent functional evaluations, and [Notarnicola et al. \(2021\)](#), where the cost is estimated via recursive least squares method. However, these works focus on discrete-time algorithms and, more importantly, have no dynamical system implemented in closed-loop with the algorithms. Few recent works considered controllers that are learned using neural networks; see, e.g., [Karg and Lucia \(2020\)](#); [Yin et al. \(2021\)](#); [Marchi et al. \(2022\)](#), and the work on reinforcement learning in [Jin and Lavaei \(2020\)](#). With respect to this literature, we utilize learning methods to estimate the cost, and we use a gradient-flow controller based on the estimated cost. Finally, similarly to [Sontag \(2022\)](#), we study the input-to-state stability (ISS) property of perturbed gradient flows. Differently from [Sontag \(2022\)](#), in this work we consider the interconnection between a perturbed gradient-flow and an LTI system.

Contributions. Our contribution is threefold. (C1) We design a data-driven feedback controller to steer the inputs and outputs of an LTI system towards the optimizer of a convex optimization problem; the controller does not require knowledge of the unknown and time-varying exogenous inputs affecting the system. The controller leverages methods that learn the cost functions of the optimization problem from historical information (as a starting estimate) and through infrequent functional evaluations during the operation of the controller. Our setting accounts for cases where the cost function admits a representation through a finite set of basis functions, and the more general case where we approximate the function using a truncated basis expansion. (C2) We leverage singular-perturbation arguments (as in (Khalil, 2002, Ch. 11) and in, e.g., Hauswirth et al. (2021a); Bianchin et al. (2020)) and the theory of perturbed systems (Khalil, 2002, Ch. 9) to derive sufficient conditions on the learning error and the controller gain to ensure that the error between the optimizer of the problem and the state of the closed-loop system is ultimately bounded. (C3) We verify the stability claims and the analytical bounds through a representative set of simulations.

Organization. Section 2 outlines the problem formulation and the main assumptions; Section 3 presents the main data-driven control framework and the stability results. Representative numerical simulations are presented in Section 4, and Section 5 concludes the paper. The proofs of the main results are reported in the Appendix of the extended version of our paper in Cothren et al. (2021)¹.

2. Problem Formulation

We consider continuous-time linear dynamical systems described by:

$$\dot{x} = Ax + Bu + Ew_t, \quad y = Cx + Dw_t, \quad (1)$$

where $x : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^n$ is the state, $u : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^m$ is the input, $y : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^p$ is the output, $w_t : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^q$ is an unknown and time-varying exogenous input or disturbance, and A, B, C, D , and E are matrices of appropriate dimensions. We make the following assumptions on (1).

Assumption 1 *The matrix A is Hurwitz stable; namely, for any $Q \in \mathbb{R}^{n \times n}$, $Q \succ 0$, there exists $P \in \mathbb{R}^{n \times n}$, $P \succ 0$, such that $A^\top P + PA = -Q$.* $\square$

Assumption 2 *The function $t \mapsto w_t$ is locally absolutely continuous.* $\square$

Under Assumption 1, for given vectors $u_{\text{eq}} \in \mathbb{R}^m$ and $w_{\text{eq}} \in \mathbb{R}^q$, (1) has a unique stable equilibrium point $x_{\text{eq}} = -A^{-1}(Bu_{\text{eq}} + Ew_{\text{eq}})$; see, e.g., (Khalil, 2002, Theorems 4.5 and 4.6). Moreover, at equilibrium, the relationship between system inputs and outputs is given by the algebraic map $y_{\text{eq}} = Gu_{\text{eq}} + Hw_{\text{eq}}$, where $G := -CA^{-1}B$ and $H := D - CA^{-1}E$. Assumption 2 characterizes how the exogenous inputs can vary over time.

We consider the problem of developing a feedback controller, inspired by online optimization methods, to regulate (1) to the solutions of the following time-dependent optimization problem:

$$u_t^* \in \arg \min_{\bar{u} \in \mathbb{R}^m} \phi(\bar{u}) + \psi(G\bar{u} + Hw_t), \quad (2)$$

1. **Notation.** We denote by $\mathbb{N}, \mathbb{N}_{>0}, \mathbb{R}, \mathbb{R}_{>0}$, and $\mathbb{R}_{\geq 0}$ the set of natural numbers, the set of positive natural numbers, the set of real numbers, the set of positive real numbers, and the set of non-negative real numbers, respectively. For vectors $x \in \mathbb{R}^n$ and $u \in \mathbb{R}^m$, $\|x\|$ denotes the Euclidean norm of x and $(x, u) \in \mathbb{R}^{n+m}$ denotes their vector concatenation. If $x, u \in \mathbb{R}^n$, then $(x^\top, u^\top) = [x^\top; u^\top] \in \mathbb{R}^{2 \times n}$ denotes the matrix with rows given by $x^\top$ and $u^\top$. For a symmetric matrix $W \in \mathbb{R}^{n \times n}$, $W \succ 0$ denotes that W is positive definite and $W \succeq 0$ denotes that W is positive semidefinite. Moreover, we let $\bar{\lambda}(W)$ and $\underline{\lambda}(W)$ denote the largest and smallest eigenvalues of W , respectively. For a continuously differentiable function $\phi : \mathbb{R}^n \rightarrow \mathbb{R}$, we denote its gradient by $\nabla \phi(x) \in \mathbb{R}^n$.

for all $t \in \mathbb{R}_{\geq 0}$, where $\phi : \mathbb{R}^m \rightarrow \mathbb{R}$ and $\psi : \mathbb{R}^p \rightarrow \mathbb{R}$ are cost functions associated with the system's inputs and outputs, respectively. The optimization problem (2) formalizes an equilibrium selection problem for which the objective is to select an optimal input u_t^* for the system (1) (and, consequently, the corresponding steady-state output $y_t^* = Gu_t^* + Hw_t$) that minimizes the cost specified by the the loss functions $\phi(\cdot)$ and $\psi(\cdot)$. We note that, since the cost function is parametrized by w_t , the solutions of (2) are also time-varying, and thus define optimal trajectories (the sub-script t is utilized to emphasize the temporal variability of w_t and, consequently, that of u_t^*).

Remark 1 (Relationship with output regulation) *The problem (2) formalizes an optimal regulation problem with steady-state constraints similar to the well-established output-regulation problem Davison (1976); with respect to the classical framework, in our setting the optimal trajectories are not generated by an exosystems (i.e., a known autonomous linear model) but instead are specified as the solution of a convex optimization problem.* $\square$

In this work, we focus on a setting where the exogenous input w_t is *unknown* and the cost functions $u \mapsto \phi(u)$ and $y \mapsto \psi(y)$ are *unknown* as explained in Section 1. In this setup, the output regulation problem tackled in this paper is summarized as follows.

Problem 1 *Design a data-driven output-feedback controller for (1) that learns the cost functions $u \mapsto \phi(u)$ and $y \mapsto \psi(y)$ from infrequent functional evaluations while concurrently driving the inputs and outputs of (1) to the time-varying optimizer of (2) up to an error that accounts for the functional estimation errors and the temporal variability of the unknown disturbance.*

Although unknown, we impose the following regularity assumptions on the cost functions.

Assumption 3 *The function $u \mapsto \phi(u)$ is continuously-differentiable, convex, and ℓ_u -smooth, for some $\ell_u \geq 0$; namely, $\exists \ell_u \geq 0$ such that $\|\nabla\phi(u) - \nabla\phi(u')\| \leq \ell_u\|u - u'\|$ holds $\forall u, u' \in \mathbb{R}^m$.* $\square$

Assumption 4 *The function $y \mapsto \psi(y)$ is continuously-differentiable, convex, and ℓ_y -smooth, for some $\ell_y \geq 0$; namely, $\exists \ell_y \geq 0$ such that $\|\nabla\psi(y) - \nabla\psi(y')\| \leq \ell_y\|y - y'\|$ holds $\forall y, y' \in \mathbb{R}^p$.* $\square$

Assumption 5 *For any $w \in \mathbb{R}^q$, the composite cost $u \mapsto \phi(u) + \psi(Gu + Hw)$ is μ_u -strongly convex, with $\mu_u > 0$.* $\square$

Assumptions 3-4 imply that $u \mapsto \phi(u) + \psi(Gu + Hw_t)$ is ℓ -smooth, with $\ell := \ell_u + \|G\|^2\ell_y \geq 0$. Two implications follow from Assumption 5: (i) the optimizer u_t^* is unique, and (ii) $\phi(u) + \psi(Gu + Hw_t)$ satisfies the Polyak-Łojasiewicz (PL) inequality as shown in Karimi et al. (2016); namely,

$$\|\nabla\phi(u) + G^\top \nabla\psi(Gu + Hw_t)\|^2 \geq 2\mu_u(\phi(u) + \psi(Gu + Hw_t) - \phi(u_t^*) - \psi(Gu_t^* + Hw_t)),$$

holds for all $u \in \mathbb{R}^m$. Regarding the functions $\phi(u)$ and $\psi(y)$, we make the following assumptions.

Assumption 6 *The function $u \mapsto \phi(u)$ admits the representation $\phi(u) = \sum_{i=1}^{N_b} \alpha_i b_i(u)$, for some $N_b \in \mathbb{N}_{>0} \cup \{+\infty\}$, where for all $i \in \{1, \dots, N_b\}$, $b_i : \mathbb{R}^m \rightarrow \mathbb{R}$ are continuously differentiable basis functions and $\alpha_i \in \mathbb{R}$ are fitting parameters.* $\square$

Assumption 7 *The function $y \mapsto \psi(y)$ admits the representation $\psi(y) = \sum_{i=1}^{M_b} \rho_i d_i(y)$, for some $M_b \in \mathbb{N}_{>0} \cup \{+\infty\}$, where for all $i \in \{1, \dots, M_b\}$, $d_i : \mathbb{R}^p \rightarrow \mathbb{R}$ are continuously differentiable basis functions and $\rho_i \in \mathbb{R}$ are fitting parameters.* $\square$

In compact form, we denote $\alpha := (\alpha_1, \dots, \alpha_N)$, $\rho := (\rho_1, \dots, \rho_M)$, $b(u) := (b_1(u), \dots, b_N(u))$, and $d(u) := (d_1(u), \dots, d_M(u))$. Moreover, we let ∇b denote the Jacobian of $b(u)$, and ∇d the Jacobian of $d(u)$. We illustrate the above assumptions through the following two examples.

Example 1 (Non-parametric models). *A representation as in Assumptions 6-7 can be obtained by using tools from Reproducing Kernel Hilbert Spaces, in which a function ϕ defined over a measurable space is estimated via interpolation based on symmetric, positive definite kernel functions (Hastie et al., 2009, Chapter 5), Bazerque and Giannakis (2013). Additional non-parametric models utilize orthonormal basis functions such as polynomials, or can leverage radial basis functions and multilayer feed-forward networks Hornik et al. (1989).*

Example 2 (Convex parametric models). *Consider the cost $\phi(u) = \frac{1}{2}u^\top \Upsilon u + v^\top u + r$, where $\Upsilon \in \mathbb{R}^{m \times m}$, $\Upsilon \succeq 0$, $v \in \mathbb{R}^m$, $r \in \mathbb{R}$. Taking as an example $m = 2$, $\Upsilon = [\Upsilon_{ij}]$, $v = (c_1, c_2)$ and $u = (u_1, u_2)$, the function admits the representation in Assumption 6 with $N_b = 6$, where $b_1(u) = 1$, $b_2(u) = u_1$, $b_3(u) = u_2$, $b_4(u) = u_1^2/2$, $b_5(u) = u_1 u_2$, and $b_6(u) = u_2^2/2$, and $\alpha = (r, v_1, v_2, \Upsilon_{11}, \Upsilon_{12}, \Upsilon_{22})$ (with the constraint $\Upsilon_{12} = \Upsilon_{21}$).*

In our learning strategy, we consider a finite number of basis functions as $\tilde{\phi}(u) = \sum_{i=1}^N \alpha_i b_i(u)$ and $\tilde{\psi}(y) = \sum_{j=1}^M \rho_j d_j(y)$, where $N < \infty$ and $M < \infty$; in particular, we assume that $N \leq N_b$ and $M \leq M_b$. Accordingly, we will consider two cases:

(Case 1) The functions ϕ and ψ are represented by a finite number of basis functions (i.e., $N_b < \infty$ and $M_b < \infty$), and we utilize $N = N_b$ and $M = M_b$.

(Case 2) At least one of the functions ϕ or ψ is represented by an infinite number of basis functions (i.e., $N_b = \infty$ and/or $M_b = \infty$). In this case, we utilize $N < \infty$ and $M < \infty$.

We remark that the scenario in (Case 2) can be of interest also in cases where N_b or M_b are finite but sufficiently large making the learning technique computationally intractable. In this case, in our framework we will approximate the two functions by utilizing $N \ll N_b$ and/or $M \ll M_b$. Finally, we note that an instance of (Case 1) is the convex parametric model in Example 2. For the latter, for a general nonparametric model as in Example 1, due to the Weierstrass high-order approximation theorem, the error in the representation of the function over a compact set goes to zero with the increasing of N_b and M_b ; however, a limited number of basis functions may be selected for model complexity considerations (see, e.g., Hastie et al. (2009) and Bazerque and Giannakis (2013)).

3. Gradient-flow Controller with Concurrent Learning

To address Problem 1, we propose a concurrent learning and online optimization scheme in which, at every time t , we: (i) process a newly available functional evaluation $(u, \phi(u))$ (resp. $(y, \psi(y))$) via a learning method to determine an estimate $\hat{\alpha}_t$ of the vector of parameters α (resp. to determine an estimate $\hat{\rho}_t$ of the vector ρ) at time t ; and (ii) use an approximate gradient-flow parametrized by the current estimates $(\hat{\alpha}_t, \hat{\rho}_t)$ to steer (1) towards the optimal solution of the problem (2).

For the learning procedure, we utilize both noisy functional evaluations received during the operation of the algorithm as well as from historical data. Accordingly, let $\hat{\phi}_{t_k} = \phi(u_{t_k}) + \varepsilon_{t_k}$ be the noisy functional evaluation received at time $t_k \in \mathbb{R}_{\geq 0}$, $k \in \mathbb{N}$, at the point $u_{t_k} \in \mathbb{R}^m$, in the presence of measurement noise $\varepsilon_{t_k} \in \mathbb{R}^m$. (By convention, we let $t_k = 0$, $k \in \mathbb{N}$, if the k -th




Figure 1: Block diagram of the proposed gradient flow controller with concurrent learning.

pair $\{(\hat{\phi}_{t_k}, u_{t_k})\}$ is derived from historical data.) Similarly, we let $\hat{\psi}_{t_j} = \psi(y_{t_j}) + v_{t_j}$ be a noisy functional evaluation of $\psi(y)$, received at time $t_j \in \mathbb{R}_{\geq 0}$, $j \in \mathbb{N}$, in the presence of output noise $v_{t_j} \in \mathbb{R}^p$. We consider learning methods that yield estimates of α and ρ at time t based on the data $\{\hat{\phi}_{t_k}, u_{t_k}\}_{t_k \leq t}$ and $\{\hat{\psi}_{t_j}, y_{t_j}\}_{t_j \leq t}$, respectively.

Framework 1: Data-based online gradient-flow controller

Input: $\hat{\alpha}_{t_0}, \hat{\rho}_{t_0}$ based on recorded data; gain $\eta > 0$, number of basis functions N, M .

For all $t \geq t_0$:

Learning

if new data point $(\hat{\phi}(u_t), u_t)$ obtained **then**

$$\hat{\alpha}_t \leftarrow \text{parameter-learning}(\{(\hat{\phi}_{t_k}, u_{t_k})\}_{t_k \leq t}) \quad (3a)$$

if new data point $(\hat{\psi}(y_t), y_t)$ obtained **then**

$$\hat{\rho}_t \leftarrow \text{parameter-learning}(\{(\hat{\psi}_{t_j}, y_{t_j})\}_{t_j \leq t}) \quad (3b)$$

end

Feedback Control

$$\dot{x} = Ax + Bu + Ew_t, \quad y = Cx + Dw_t$$

$$\dot{u} = -\eta(\nabla b(u)^\top \hat{\alpha}_t + G^\top \nabla d(y)^\top \hat{\rho}_t). \quad (3c)$$

The proposed scheme is illustrated in Figure 1, and it is described by the pseudo-code in Framework 1 (with t_0 denoting the initial time). According to the proposed method, an initial estimate of the parameters α and ρ is obtained using recorded data, where $\text{parameter-learning}(\cdot)$ is a map that represents an estimation step; examples of learning methods will be provided in Section 3.2. Since for any time $t \geq t_0$ the parameters are updated only when a new functional evaluation is available, we let $t \mapsto \hat{\alpha}_t$ and $t \mapsto \hat{\rho}_t$ be piece-wise constant right-continuous functions that represent the most up-to-date estimates of α and ρ , respectively. We utilize a gradient-flow controller, as shown in (3c), where the gain $\eta > 0$ induces a time-scale separation between the plant and the controller. Finally, we notice that since the vector field characterizing (3c) is piece-wise continuous in time, the initial value problem (3c) always admits a local solution that is unique (Khalil, 2002, Thm 3.1).

3.1. Main results

In this section, we characterize the transient performance of Framework 1. To this end, let

$$z := (u - u_t^*, x - x_t^*), \quad x_t^* = -A^{-1}Bu_t^* - A^{-1}Ew_t, \quad (4)$$

denote the error between the state of (3c) and the equilibrium of (2), respectively. In the following, we provide sufficient conditions on η and on the estimation errors $\alpha - \hat{\alpha}_t$, $\rho - \hat{\rho}_t$ so that the interconnection between the plant and the data-based controller (3c) is exponentially stable.

The first result is stated for (*Case 1*), where the functions ϕ and ψ are represented by a finite number of basis functions, and we set precisely $N = N_b$ and $M = M_b$.

Theorem 1 (Control bound for finite number of basis functions) *Let Assumptions 1-7 be satisfied, and assume that $N = N_b < \infty$ and $M = M_b < \infty$. Let $z(t)$ be defined as in (4), with $(u(t), x(t))$ the state of (3c). Suppose that the learning errors satisfy,*

$$0 < \epsilon := \ell_u \sup_{t_0 \leq \tau \leq t} \|\alpha - \hat{\alpha}_\tau\| + \ell_y \|G\|^2 \sup_{t_0 \leq \tau \leq t} \|\rho - \hat{\rho}_\tau\| < \frac{c_0}{c_3}, \quad (5)$$

where $c_0 := s \min \{2\mu_u \eta, \underline{\lambda}(Q)/\bar{\lambda}(P)\}$, $c_3 := \eta \max \{2\ell\mu_u^{-1}, 4c_1^{-1}\|PA^{-1}B\|\}$, $c_1 := \min\{(1 - \theta)\mu_u/2, \theta\underline{\lambda}(P)\}$, $\theta := \ell_y\|G\|\|C\|/(\ell_y\|G\|\|C\| + 2\|PA^{-1}B\|)$, and $s \in (0, 1)$. Suppose further that the controller gain satisfies,

$$0 < \eta < \frac{(1-s)^2 \underline{\lambda}(Q)}{(2-s)2\|PA^{-1}B\|\ell_y\|G\|\|C\|}. \quad (6)$$

Then there exists $\kappa_1, \kappa_2, \kappa_3 > 0$ such that the error $z(t)$ satisfies

$$\|z(t)\| \leq \kappa_1 e^{-\frac{1}{2}a(t-t_0)} \|z(t_0)\| + \kappa_2 \int_{t_0}^t e^{-\frac{1}{2}a(t-\tau)} \Delta(\tau) d\tau + \kappa_3 \int_{t_0}^t e^{-\frac{1}{2}a(t-\tau)} \|\dot{w}_\tau\| d\tau, \quad (7)$$

for all $t \geq t_0 \geq 0$, where $a := c_0 - \epsilon c_3$ and $\Delta(\tau) := \|\nabla b(u_\tau^*)\| \|\alpha - \hat{\alpha}_\tau\| + \|\nabla d(y_\tau^*)\| \|\rho - \hat{\rho}_\tau\|$.

Detailed expressions for the constants $\kappa_1, \kappa_2, \kappa_3$ and the proof of Theorem 1 are provided in the Appendix of the extended version in Cothren et al. (2021). Theorem 1 asserts that if the worst-case estimation error for the parameters of the cost functions, captured by $\sup_{t_0 \leq \tau \leq t} \|\alpha - \hat{\alpha}_\tau\|$ and $\sup_{t_0 \leq \tau \leq t} \|\rho - \hat{\rho}_\tau\|$, is such that ϵ satisfies the condition (5), then a sufficiently-small choice of the controller gain η guarantees exponential convergence of the state of (3c) to a neighborhood of the optimizer u_t^* of (2) and the corresponding state x_t^* . In particular, the error $z(t)$ is ultimately bounded by two terms: the first depends on the error $\Delta(t)$, which accounts for the estimation error in the function parameters, and the second depends on temporal variability of w_t (which affects the dynamics of the plant (1) making the optimizer of (2) time varying). The condition on ϵ also suggests prerequisites on the “richness” of the recorded data in the sense that they must yield a sufficiently small estimation error.

It is important to notice that $a = c_0 - \epsilon c_3$, which characterizes the rate of exponential decay, is proportional to smallest value between the rate of the convergence of the open-loop plant (found in c_0 as the ratio $\underline{\lambda}(Q)/\bar{\lambda}(P)$ (Khalil, 2002, Chapter 4, Theorem 4.10)) and the strong convexity parameter μ_u , characterizing the cost function. Moreover, the rate of convergence a is proportional to the controller gain η (as described by c_0 and c_3), and inversely proportional to the worst-case estimation errors of the parameters of the cost, $\sup_{t_0 \leq \tau \leq t} \|\alpha - \hat{\alpha}_\tau\|$ and $\sup_{t_0 \leq \tau \leq t} \|\rho - \hat{\rho}_\tau\|$.

Remark 2 (Asymptotic behavior and input-to-state stability) Two important implications follow from the statement of Theorem 1 as subcases.

- If $\lim_{t \rightarrow \infty} \Delta(t) = 0$ and $\lim_{t \rightarrow \infty} \dot{w}_t = 0$, then (7) guarantees that $\lim_{t \rightarrow \infty} z(t) = 0$, namely, the state of (3c) converges (exactly) to the optimizer of (2). See (Khalil, 2002, Lemma 9.6).
- If $\text{ess sup}_{t_0 \leq \tau \leq t} \|\Delta(\tau)\| < \infty$ and $\text{ess sup}_{t_0 \leq \tau \leq t} \|\dot{w}_\tau\| < \infty$, then (7) guarantees that

$$\|z(t)\| \leq \kappa_1 e^{-\frac{1}{2}a(t-t_0)} \|z(t_0)\| + 2a^{-1} (\kappa_2 \text{ess sup}_{t_0 \leq \tau \leq t} \|\Delta(\tau)\| + \kappa_3 \text{ess sup}_{t_0 \leq \tau \leq t} \|\dot{w}_\tau\|). \quad (8)$$

It follows that the bound (8) guarantees input-to state stability of (3c) (in the sense of Sontag and Wang (1997); Angeli et al. (2003); Sontag (2022)) with respect to the inputs Δ and $\dot{w}_t$. $\square$

We now consider (Case 2), where the estimation procedure is affected by a truncation error; that is, when only $N \ll N_b$ and $M \ll M_b$ basis functions are utilized by the controller. Accordingly, consider the approximated functions $\tilde{\phi}(u) = \sum_{i=1}^N \alpha_i b_i(u)$ and $\tilde{\psi}(y) = \sum_{j=1}^M \rho_j d_j(y)$, and define the two truncation errors for $\phi(u)$ and $\psi(y)$ as $e_\phi(u) := \sum_{i=N+1}^{N_b} b_i(u) \alpha_i$ and $e_\psi(y) := \sum_{j=M+1}^{M_b} d_j(y) \rho_j$, respectively. We make the following regularity assumptions:

Assumption 8 The approximated functions $\tilde{\phi}(u)$ and $\tilde{\psi}(y)$ are ℓ_u^N - and ℓ_y^M -smooth for constants $\ell_u^N \geq 0$ and $\ell_y^M \geq 0$, respectively.

Assumption 9 The truncation terms $e_\phi(u)$ and $e_\psi(y)$ have a Lipschitz-continuous gradient with constants $\ell_u^e \geq 0$ and $\ell_y^e \geq 0$, respectively.

Assumption 8 is satisfied if the basis functions $\{b_i\}_{i=1}^N$ and $\{d_i\}_{i=1}^M$ are strongly smooth. Assumption 9 is a technical condition that is required in our proof; nevertheless, this condition is satisfied for parametric convex models and several non-parametric models; see Hastie et al. (2009); Bazerque and Giannakis (2013), and the recent results of Poveda et al. (2021).

We next characterize the controller transient performance in the presence of truncation errors.

Theorem 2 (Stability with function approximation error) Let Assumptions 1-9 be satisfied, and assume that $N < N_b$ and $M < M_b$ basis functions are utilized in the learning of ϕ and ψ , respectively. Suppose that

$$0 < \epsilon' := \ell_u^N \sup_{t_0 \leq \tau \leq t} \|\alpha - \hat{\alpha}_\tau\| + \ell_y^M \|G\|^2 \sup_{t_0 \leq \tau \leq t} \|\rho - \hat{\rho}_\tau\| + \ell_u^e + \ell_y^e \|G\|^2 < \frac{c_0}{c_3}, \quad (9)$$

where s, c_0, c_1, c_3, θ are given in Theorem 1 and η satisfies (6). Then there exists $\kappa_1, \kappa_2, \kappa_3 > 0$ such that the error (4) satisfies

$$\|z(t)\| \leq \kappa_1 e^{-\frac{1}{2}a'(t-t_0)} \|z(t_0)\| + \kappa_2 \int_{t_0}^t e^{-\frac{1}{2}a'(t-\tau)} \Xi(\tau) d\tau + \kappa_3 \int_{t_0}^t e^{-\frac{1}{2}a'(t-\tau)} \|\dot{w}_\tau\| d\tau, \quad (10)$$

for all $t \geq t_0 \geq 0$, where $a' = c_0 - \epsilon' c_3$ and $\Xi(\tau) := \|\nabla b(u_t^*)\| \|\alpha - \hat{\alpha}_\tau\| + \|G\| \|\nabla d(y_t^*)\| \|\rho - \hat{\rho}_\tau\| + \|\nabla e_\phi(u_t^*)\| + \|G\| \|\nabla e_\psi(y_t^*)\|$.

Detailed expressions for the constants $\kappa_1, \kappa_2, \kappa_3$ are provided in the extended version in Cothren et al. (2021). By comparison with the definition of $\Delta(t)$, Theorem 1, $\Xi(t)$ simultaneously accounts for estimation error and truncation error. We also notice that similar claims to those in Remark 2 can be derived in this case.

3.2. Parameter learning

We present some methods that can be utilized in the step parameter-learning($\cdot$) in Framework 1.

Least Squares Estimator. Consider the function $\phi(u)$. Suppose that at a given time $t \in \mathbb{R}_{\geq 0}$, $K > 0$ data points $(\hat{\phi}_{t_k}, u_{t_k})$ are available, with $\{t_k\}_{k=1}^{K-1}$ the time instants where data points are received. Let $\hat{\Phi}_t := (\hat{\phi}_{t_1}, \dots, \hat{\phi}_{t_K})$ be a vector collecting the functional evaluations received up to time t , and $B_t \in \mathbb{R}^{K \times N}$ be a matrix with rows the regression vectors $\{b(u_{t_k})^\top, k = 1, \dots, K\}$. Given these data points, the ordinary least squares (LS) method determines a solution to the following optimization problems; see, e.g., [Kay \(1993\)](#); [Hastie et al. \(2009\)](#); [Beck \(2017\)](#):

$$\text{if } K \geq N: \quad \hat{\alpha}_t = \arg \min_{\alpha \in \mathbb{R}^N} \|\hat{\Phi}_t - B_t \alpha\|^2, \quad \text{if } K < N: \quad \hat{\alpha}_t = \arg \min_{\alpha \in \mathbb{R}^N, \text{ s.t.: } \hat{\Phi}_t = B_t \alpha} \|\alpha\|^2.$$

The above optimization problems admit a unique closed-form solution given by $\hat{\alpha}_t = B_t^\dagger \hat{\Phi}_t$, where $B_t^\dagger$ denotes the Moore-Penrose inverse of B_t . Moreover, the resulting approximation error admits a closed-form expression given by $\|B_t \hat{\alpha}_t - \hat{\Phi}_t\|^2 = \|(I - B_t B_t^\dagger) \hat{\Phi}_t\|^2$, which can be interpreted by noting that $(I - B_t B_t^\dagger)$ is the orthogonal projector onto the null space of $B_t^\top$.

A similar procedure can be utilized for the function ψ ; in particular, we let $\hat{\Psi}_t := (\hat{\psi}_{t_1}, \dots, \hat{\psi}_{t_K})$ be a vector collecting the functional evaluations of ψ received up to time t , and we define the matrix $D_t := (d(y_{t_1})^\top, \dots, d(y_{t_K})^\top)$. Then, the LS yields the estimate $\hat{\rho}_t = D_t^\dagger \hat{\Psi}_t$.

Recursive Least Squares. To avoid the computation of the Moore-Penrose inverse, one can utilize the recursive LS approach; we refer the reader to [Ljung \(1999\)](#) and [Kay \(1993\)](#) an overview and for the main equations of the recursive LS. See also [Notarnicola et al. \(2021\)](#).

Ridge Regression. Given the data points $(\hat{\Psi}_t, B_t)$, the ridge regression involves the solution of the optimization problem $\min_{\alpha \in \mathbb{R}^N} \|\hat{\Phi}_t - B_t \alpha\|^2 + \lambda_t \|\alpha\|^2$, where $\lambda_t > 0$ is a tuning parameter. While this criterion was proposed to alleviate the singularity of $B_t^\top B_t$ when $K < N$ in [Hoerl and Kennard \(1970\)](#), the regularization $\lambda_t \|\alpha\|^2$ can be shown to impose a penalty on the norm of α ; in fact, for a given $\lambda_t > 0$, there exists $\nu_t > 0$ such that the solution of the ridge regression is equivalent to the LS with the constraint $\|\alpha\| \leq \nu_t$ as explained by, e.g., [Hastie et al. \(2009\)](#). The ridge regression problem admits a unique closed-form solution given by $\hat{\alpha}_t = (B_t^\top B_t + \lambda_t I)^{-1} B_t^\top \hat{\Phi}_t$, where I_N denotes the identity matrix. Similarly, upon receiving the K -th data point, the estimate of the vector $\hat{\rho}_t$ can be updated as $\hat{\rho}_t = (D_t^\top D_t + \lambda_t I)^{-1} D_t^\top \hat{\Psi}_t$, where I_N denotes the identity matrix. To avoid the matrix inversion, a recursive strategy via the Woodbury matrix identity can be adopted.

Sparse Linear Regression. To select the basis functions that provide a parsimonious representation of the function, one could utilize a sparse linear regression method. This amounts to solving the problem $\min_{\alpha \in \mathbb{R}^N} \|\hat{\Phi}_t - B_t \alpha\|^2 + \lambda_t \|\alpha\|_1$, where $\lambda_t > 0$ is a tuning parameter that promotes sparsity of the vector α as explained in [Tibshirani \(1996\)](#). The solution of the problem can be found in closed form, where the i -th entry of $\hat{\alpha}_t$ is given by $\hat{\alpha}_{i,t} = \max\{|z_{i,t}| - \lambda_t, 0\} \text{sgn}(z_{i,t})$, with $z_t = (B_t^\top B_t)^{-1} B_t^\top \hat{\Phi}_t$ and where $\text{sgn}(\cdot)$ is the sign function.

4. Numerical Verification

In this section, we numerically verify the proposed Framework 1 in two cases: (i) with constant disturbance w_t , and (ii) with time-varying disturbance. As an illustrative example, we utilize the LS estimator described in Section 3.2. We consider the cost functions $\phi(u) = \frac{1}{2} u^\top \Upsilon u + v^\top u + r$ and $\psi(y) = \frac{1}{2} \|Gu + Hw_t - \xi\|^2$, where $\Upsilon \in \mathbb{R}^{m \times m}$ is a symmetric positive-definite matrix, $v \in \mathbb{R}^m$ is positive element-wise, $r \in \mathbb{R}$ is positive, and $\xi \in \mathbb{R}^p$ is a reference output of the system.




Figure 2: Evolution of the error $\|z(t)\|$, and theoretical bounds provided in Theorem 1. Left: constant disturbance, $\dot{w}_t = 0$. Right: time-varying disturbance, $\dot{w}_t \neq 0$. Vertical black dotted lines represent time instants where functional evaluations are received.

Since the functions are convex and quadratic, we consider the basis expansion in Example (2) (see also Notarnicola et al. (2021)); notice that to get orthonormal basis functions, we remove the lower-triangular part of Υ and we impose that $\Upsilon_{ij} = \Upsilon_{ji}$, for $i \neq j$. For illustration purposes, we consider a case where we estimate $\phi(u)$ (while $\psi(y)$ is known), and we consider the case $m = 4$. We generate some of the matrices of the plant and of the cost functions using normally distributed random variables; the values of the matrices are reported in our extended version of the paper in Cothren et al. (2021). We utilize four data points $\{(u, \phi(u))\}$ as recorded data; as such, the LS is under-determined at the start of the algorithm. During the execution of the algorithm, we simulate the arrival of new data points by using a Poisson clock. The gain of the controller satisfies the condition outlines in Theorem 1.

Figure 2 illustrates the evolution of the error $z(t)$, defined in (4), as well as the theoretical bound provided in equation (7) of Theorem 1. Figure 2(a) illustrated the case where the disturbance is constant; new functional evaluations are received at the times marked with vertical black dotted lines. The theoretical bound depends on the estimation error, and it exhibits step changes when a new data point is received; the theoretical bound appears to be tight in the numerical simulation.

Figure 2(b) illustrated the case where the disturbance is time varying; in this case, the bound is affected by both the estimation error and $\|\dot{w}_t\|$. In both cases, the numerical trajectory exhibits an exponential convergence up to an asymptotic error. The asymptotic error is affected by the error in the estimation of the parameters and the variability of the disturbance.

5. Conclusions

We proposed a data-enabled gradient-flow controller to regulate an LTI dynamical system to the minimizer of an unknown functions. The controller is aided by a learning method that estimates the unknown costs from functional evaluations; to this end, appropriate basis expansion representations (either parametric or non-parametric) are utilized. We established sufficient conditions on the estimation error and the controller gain to ensure that the error between the optimizer of the problem and the state of the closed-loop system is ultimately bounded; the error bound accounts for the functional estimation errors and the temporal variability of the unknown disturbance. Future works will look at learning methods such as concurrent learning dynamics and neural networks.

Acknowledgments

This work was supported by the National Science Foundation (NSF) through the Awards CMMI 2044946 and CAREER 1941896.

References

- D. Angeli, E. D. Sontag, and Y. Wang. Input-to-state stability with respect to inputs and their derivatives. *International Journal of Robust and Nonlinear Control: IFAC-Affiliated Journal*, 13 (11):1035–1056, 2003.
- K. B. Ariyur and M. Krstić. *Real-time optimization by extremum-seeking control*. John Wiley & Sons, 2003.
- J. Bazerque and G. B. Giannakis. Nonparametric basis pursuit via sparse kernel-based learning: A unifying view with advances in blind methods. *IEEE Signal Processing Magazine*, 30(4): 112–125, 2013.
- A. Beck. *First Order Methods in Optimization*. MOS-SIAM Series on Optimization, Tel-Aviv University, Tel-Aviv, Israel, 2017.
- G. Bianchin, J. I. Poveda, and E. Dall’Anese. Online optimization of switched LTI systems using continuous-time and hybrid accelerated gradient flows. *arXiv preprint arXiv:2008.03903*, 2020.
- G. Bianchin, J. Cortés, J. I. Poveda, and E. Dall’Anese. Time-varying optimization of LTI systems via projected primal-dual gradient flows. *IEEE Trans. on Control of Networked Systems*, 2021a.
- G. Bianchin, E. Dall’Anese, J. I. Poveda, D. Jacobson, E. J. Carlton, and A. G. Buchwald. Planning a return to normal after the COVID-19 pandemic: Identifying safe contact levels via online optimization. *arXiv preprint arXiv:2109.06025*, 2021b.
- D. D. Bourgin, J. C. Peterson, D. Reichman, S. J. Russell, and T. L. Griffiths. Cognitive model priors for predicting human decisions. In *International Conference on Machine Learning*, pages 5133–5141, 2019.
- F. D. Brunner, H.-B. Dürr, and C. Ebenbauer. Feedback design for multi-agent systems: A saddle point approach. In *IEEE Conf. on Decision and Control*, pages 3783–3789, 2012.
- G. Chowdhary and E. Johnson. Concurrent learning for convergence in adaptive control without persistency of excitation. In *IEEE Conf. on Decision and Control*, pages 3674–3679. IEEE, 2010.
- G. Chowdhary, T. Yucelen, M. Mühlegg, and E. N. Johnson. Concurrent learning adaptive control of linear systems with exponentially convergent bounds. *International Journal of Adaptive Control and Signal Processing*, 27(4):280–301, 2013.
- M. Colombino, E. Dall’Anese, and A. Bernstein. Online optimization as a feedback controller: Stability and tracking. *IEEE Trans. on Control of Network Systems*, 7(1):422–432, 2020.

- 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 (extended version). *arXiv preprint*, 2021. <https://arxiv.org/abs/2112.01652>.
- E. Davison. The robust control of a servomechanism problem for linear time-invariant multivariable systems. *IEEE Trans. on Automatic Control*, 21(1):25–34, 1976.
- A. M Devraj and S. P Meyn. Zap Q-learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 2232–2241, 2017.
- F. Fabiani, A. Simonetto, and P. J. Goulart. Learning equilibria with personalized incentives in a class of nonmonotone games. *arXiv preprint arXiv:2111.03854*, 2021.
- F. Gündling, F. Hopp, and K. Weihe. Efficient monitoring of public transport journeys. *Public Transport*, 12(3):631–645, 2020.
- T. Hastie, R. Tibshirani, and J. Friedman. *The elements of statistical learning*. Springer, 2009.
- A. Hauswirth, S. Bolognani, G. Hug, and F. Dörfler. Timescale separation in autonomous optimization. *IEEE Trans. on Automatic Control*, 66(2):611–624, 2021a.
- A. Hauswirth, S. Bolognani, G. Hug, and F. Dörfler. Optimization algorithms as robust feedback controllers. *arXiv preprint arXiv:2103.11329*, 2021b.
- K. Hirata, J. P. Hespanha, and K. Uchida. Real-time pricing leading to optimal operation under distributed decision makings. In *Proc. of American Control Conf.*, Portland, OR, June 2014.
- A. E. Hoerl and R. W. Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970.
- K. Hornik, M. Stinchcombe, and H. White. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.
- M. Jin and J. Lavaei. Stability-certified reinforcement learning: A control-theoretic perspective. *IEEE Access*, 8:229086–229100, 2020.
- A. Jokic, M. Lazar, and P. P.-J. Van Den Bosch. On constrained steady-state regulation: Dynamic KKT controllers. *IEEE Trans. on Automatic Control*, 54(9):2250–2254, 2009.
- B. Karg and S. Lucia. Stability and feasibility of neural network-based controllers via output range analysis. In *IEEE Conference on Decision and Control*, pages 4947–4954, 2020.
- H. Karimi, J. Nutini, and M. Schmidt. Linear convergence of gradient and proximal-gradient methods under the Polyak-Łojasiewicz condition. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 795–811. Springer, 2016.
- S. M. Kay. *Fundamentals of statistical signal processing: estimation theory*. Prentice-Hall, Inc., 1993.
- H. K. Khalil. *Nonlinear Systems; 3rd ed.* Prentice-Hall, Upper Saddle River, NJ, 2002.

- M. Krstic and H.-H. Wang. Stability of extremum seeking feedback for general nonlinear dynamic systems. *Automatica*, 36(4):595–602, 2000.
- L. S. P. Lawrence, Z. E. Nelson, E. Mallada, and J. W. Simpson-Porco. Optimal steady-state control for linear time-invariant systems. In *IEEE Conf. on Decision and Control*, pages 3251–3257, December 2018.
- L. S. P. Lawrence, J. W. Simpson-Porco, and E. Mallada. Linear-convex optimal steady-state control. *IEEE Trans. on Automatic Control*, 2021. (To appear).
- L. Ljung. *System identification: theory for the user*. 2nd edition Prentice-Hall, Upper Saddle River, NJ, 1999.
- W. Luan, D. Sharp, and S. LaRoy. Data traffic analysis of utility smart metering network. In *IEEE Power & Energy Society General Meeting*, pages 1–4, 2013.
- X. Luo, Y. Zhang, and M. M. Zavlanos. Socially-aware robot planning via bandit human feedback. In *International Conf. on Cyber-Physical Systems*, pages 216–225. IEEE, 2020.
- M. Marchi, J. Bunton, B. Gharesifard, and P. Tabuada. Safety and stability guarantees for control loops with deep learning perception. *IEEE Control Systems Letters*, 6:1286–1291, 2022.
- S. Menta, A. Hauswirth, S. Bolognani, G. Hug, and F. Dörfler. Stability of dynamic feedback optimization with applications to power systems. In *Annual Conf. on Communication, Control, and Computing*, pages 136–143, 2018.
- S. Munir, J. A. Stankovic, C.-J. M. Liang, and S. Lin. Cyber physical system challenges for human-in-the-loop control. In *8th International Workshop on Feedback Computing*, 2013.
- I. Notarnicola, A. Simonetto, F. Farina, and G. Notarstefano. Distributed personalized gradient tracking with convex parametric models, 2021.
- A. M. Ospina, A. Simonetto, and E. Dall’Anese. Personalized demand response via shape-constrained online learning. In *IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids*, pages 1–6, 2020.
- J. I. Poveda, M. Benosman, and K. G. Vamvoudakis. Data-enabled extremum seeking: a cooperative concurrent learning-based approach. *International Journal of Adaptive Control and Signal Processing*, 35(7):1256–1284, 2021.
- A. Robey, Hu H., L. Lindemann, H. Zhang, D. Dimarogonas, S. Tu, and N. Matni. Learning control barrier functions from expert demonstrations. In *IEEE Conf. on Decision and Control*.
- A. Simonetto, E. Dall’Anese, J. Monteil, and A. Bernstein. Personalized optimization with user’s feedback. *Automatica*, 131:109767, 2021.
- E. D. Sontag. Remarks on input to state stability on open sets, motivated by perturbed gradient flows in model-free learning. *Systems & Control Letters*, 161:105138, 2022.
- E. D. Sontag and Y. Wang. Output-to-state stability and detectability of nonlinear systems. *Systems & Control Letters*, 29(5):279–290, 1997.

- P. Tabuada and B. Ghahsifard. Universal approximation power of deep residual neural networks via nonlinear control theory. *arXiv preprint arXiv:2007.06007*, 2020.
- A. Taylor, A. Singletary, Y. Yue, and A. Ames. Learning for safety-critical control with control barrier functions. In *Learning for Dynamics and Control*, pages 708–717. PMLR, 2020.
- A. R. Teel and D. Popovic. Solving smooth and nonsmooth multivariable extremum seeking problems by the methods of nonlinear programming. In *American Control Conference*, volume 3, pages 2394–2399, 2001.
- R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society*, 58(1):267–288, 1996.
- E. van Nunen, F. Esposto, A. K. Saberi, and J.-P. Paardekooper. Evaluation of safety indicators for truck platooning. In *IEEE Intelligent Vehicles Symposium*, pages 1013–1018, 2017.
- H. Yin, P. Seiler, and M. Arcak. Stability analysis using quadratic constraints for systems with neural network controllers. *IEEE Trans. on Automatic Control*, 2021. (Early Access).
- T. Zheng, J. Simpson-Porco, and E. Mallada. Implicit trajectory planning for feedback linearizable systems: A time-varying optimization approach. In *American control Conference*, pages 4677–4682, 2020.