Online Optimization of Linear-Time Invariant Dynamical Systems with Cost Perception

Liliaokeawawa Cothren, Ana M. Ospina
Dept. of ECEE
University of Colorado Boulder
Boulder, CO, USA
{liliaokeawawa.cothren, ana.ospina}@colorado.edu

Gianluca Bianchin
ICTEAM Institute
Université Catholique de Louvain
Louvain, Belgium
gianluca.bianchin@uclouvain.be

Emiliano Dall'Anese
Dept. of ECEE
University of Colorado Boulder
Boulder, CO, USA
emiliano.dallanese@colorado.edu

2022 56th Asilomar Conference on Signals, Systems, and Computers | 978-1-6654-5906-8/22/$31.00 ©2022 IEEE | DOI: 10.1109/IEEECONF56349.2022.10052021

Abstract—This paper considers the problem of regulating a discrete-time linear time-invariant (LTI) system to solution trajectories of a convex optimization problem, with an unknown cost. We propose a data-driven, gradient-based feedback controller that uses estimates of the cost functions obtained by a trained neural network to control the LTI system. We identify sufficient conditions to guarantee exponential input-to-state stability (ISS) of the closed loop system with respect to errors due to disturbances, temporal variability of the cost functions, and the need to use estimated costs from a neural network. Finally, we provide an illustrative numerical example in the context of online ride-share scheduling.

Index Terms—Feedback optimization, learning for control, time-varying optimization.

I. Introduction

Incorporating information from rich, perceptual data obtained by sensors or humans remains a challenge in the control of complex autonomous systems. Information may pertain to performance and safety metrics associated with the system, or perception of satisfaction, comfort, etc., of users interacting with (or affected by) the system. Recently, a number of works have addressed the problem of developing data-driven optimization methods, to concurrently learn the cost functions as the optimization algorithm is executed [1]–[4]; however, it remains an open problem how to apply similar concepts in the context of feedback optimization, where feedback controllers inspired by optimization algorithms are synthesized to regulate a dynamical system to solutions of an optimization problem.

Literature on feedback optimization includes KKT-type controllers [5]–[7] and, more recently, controllers based in first-order optimization methods as in [8]–[16]. In all of these works, the execution of feedback optimization algorithms for control systems critically relies on knowledge of the system's inputs, outputs, and cost functions. When we have incomplete knowledge of such information, controllers may be learned from neural networks [17]–[19]. In this paper, we investigate how to integrate rich data of the cost functions into controllers inspired by classical optimization algorithms for the regulation of linear, time-invariant (LTI) systems towards the solution of a convex optimization problem. Moreover, the (possibly) time-varying optimization problem includes costs associated with the system's inputs and measured outputs.

We develop a data-driven, feedback controller inspired by the projected gradient descent method. Since the time-varying cost functions must be estimated, we leverage a trained neural network (NN) that maps inputs and outputs of the system into cost function estimates. Our gradient-based controller uses these estimates of the cost via a centered-difference approximation of the gradients to generate inputs for the system. We derive sufficient conditions on the controller gain (or, step size in the gradient descent literature) to guarantee input-to-state stability (ISS) of the control loop in the sense of [20]. Further, our results illustrate that the interconnected system of the LTI system and controller tracks the optimal solution trajectory of the convex optimization problem up to an error accounting for the need to estimate the cost functions and the temporal variability of the cost functions and disturbances in the plant.

We note a similar perception-based regulation problem was considered in our prior work [21]. However, [21] developed controllers in continuous time, whereas here we consider a discrete-time setting for both the dynamical system and projected gradient-based controller. Additionally, here we extend our previous work by considering time-varying cost functions and time-varying constraints on the admissible set for the control inputs. Relative to [22], we consider the case where the cost of the problem is now known and must be learned.

II. Problem Formulation

In the following, we formalize our research problem and discuss relevant assumptions.

A. Model of the System

We consider linear time-invariant systems of the form,

$$x_{k+1} = Ax_k + Bu_k + Ew_k, \tag{1a}$$

$$y_k = Cx_k + Dw_k, \tag{1b}$$

where $k \in \mathbb{N}$ is the time index, $x : \mathbb{N} \to \mathbb{R}^{n_x}$ is the state, $u : \mathbb{N} \to \mathbb{R}^{n_u}$ is the control input, $y : \mathbb{N} \to \mathbb{R}^{n_y}$

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

is the output, $w : \mathbb{N} \to \mathbb{R}^{n_w}$ is an unknown disturbance, and the real matrices A, B, C, D and E are of appropriate dimensions. We impose the following assumption on (1).[1]

Assumption 1. The matrix $A \in \mathbb{R}^{n_x \times n_x}$ is Schur stable. Namely, for any $Q \succ 0$, there exists $P \succ 0$ such that $A^\top PA - P = -Q$. □

Assumption 1 guarantees that for any fixed $u \in \mathbb{R}^{n_u}$ and $w \in \mathbb{R}^{n_w}$, (1) admits a unique equilibrium point of the state and output, given by,

$$x^* = (I - A)^{-1} Bu + (I - A)^{-1} Ew, \tag{2a}$$

$$y^* = Gu + Hw, \tag{2b}$$

where I is the identity matrix, $G := C(I - A)^{-1} B$ is the steady state transfer function from control inputs to outputs, and $H := D + C(I - A)^{-1} E$ is the steady state gain from disturbances to system outputs. Finally, we impose the following restrictions on the control inputs and disturbances.

Assumption 2. At time $k \in \mathbb{N}$, the set of admissible control inputs $\mathcal{U}_k \subset \mathbb{R}^{n_u}$ is compact and convex. The unknown disturbances are assumed to be bounded; i.e., the set of admissible disturbances $\mathcal{W}$ is compact. □

Notably, the admissible set $\mathcal{U}_k$ given in Assumption 2 is possibly time-varying at each time step k. The restriction to such a set is often a common choice in several applications due to physical or operational constraints for actuators. In the following, we state our optimization problem and formalize our controller design.

B. Target Optimization Problem

We focus on driving, at every time index $k \in \mathbb{N}$, the system (1) to the solution of the following time-varying optimization problem:

$$u_k^* \in \arg\min_{\bar{u} \in \mathcal{U}_k} \phi(\bar{u}, \theta_{\phi,k}) + \psi(G\bar{u} + Hw_k, \theta_{\psi,k}) \tag{3}$$

[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. For vectors $x \in \mathbb{R}^n$ and $u \in \mathbb{R}^m$, $\|x\|$ denotes the Euclidean norm of x, $\|x\|_\infty$ denotes the supremum norm, and $(x, u) \in \mathbb{R}^{n+m}$ denotes their vector concatenation; $x^\top$ denotes transposition, and x_i denotes the i-th element of x. For a matrix $A \in \mathbb{R}^{n\times m}$, $\|A\|$ is the induced 2-norm and $\|A\|_\infty$ the supremum norm. $A \succ 0$ means that the matrix A is positive definite. $\Pi_{\mathcal{U}_k}(z)$ denotes the Euclidean projection of $z \in \mathbb{R}^n$ onto $\mathcal{U}_k \subset \mathbb{R}^n$ at time $k \in \mathbb{N}_{>0}$; or, $\Pi_{\mathcal{U}_k}(z) := \arg\min_{u\in\mathcal{U}_k} \|u - z\|^2$. For continuously differentiable $\phi : \mathbb{R}^n \to \mathbb{R}$, $\nabla\phi(x) \in \mathbb{R}^n$ denotes its gradient.

Partial ordering. The first orthant partial order on $\mathbb{R}^n$ is denoted as $\preceq$ and it is defined as follows: for any $x, z \in \mathbb{R}^n$, we say that $x \preceq z$ if $x_i \leq z_i$ for $i = 1, \ldots, n$. We say that a function $\phi : \mathbb{R}^n \to \mathbb{R}^n$ is monotone if for any $x, z \in \mathbb{R}^n$ such that $x \preceq \phi$ we have that $\phi(x) \leq \phi(z)$. Finally, the interval $[x, z]$, for some $x, z \in \mathbb{R}^n$, is defined as $[x, z] = \{w \in \mathbb{R}^n : x \preceq w \preceq z\}$.

Set covering. Let $\mathcal{Q}, \mathcal{Q}_s \subset \mathbb{R}^n$, with $\mathcal{Q}$ compact. We say that $\mathcal{Q}_s$ is an ϱ-cover of $\mathcal{Q}$, for some $\varrho > 0$, if for any $x \in \mathcal{Q}$ there exists a $z \in \mathcal{Q}_s$ such that $\|x - z\|_\infty \leq \varrho$. We say that $\mathcal{Q}_s$ is an ϱ-cover of $\mathcal{Q}$ "with respect to the partial order $\preceq$," for some $\varrho > 0$, if for any $x \in \mathcal{Q}$ there exists $w, z \in \mathcal{Q}_s$ such that $x \in [w, z]$ and $\|w - z\|_\infty \leq \varrho$ [23].

where $\phi : \mathcal{U} \times \Omega_\phi \to \mathbb{R}$ and $\psi : \mathbb{R}^{n_y} \times \Omega_\psi \to \mathbb{R}$ are cost functions parametrized by $\theta_{\phi,k} \in \Omega_\phi$ and $\theta_{\psi,k} \in \Omega_\psi$ associated with the systems inputs and outputs, respectively. Further, we assume that Ω_ϕ and Ω_ψ are compact sets. Our choice of such cost functions allows us to consider functions whose parameters vary in time.

The problem (3) formalizes an equilibrium selection problem for which we select an optimal input for the plant - which is at equilibrium - that minimizes the combined cost $\phi_k + \psi_k$. Addiotionally, the problem (3) is parameterized by the unknown disturbance w_k. This means that the solutions of (3) are also parameterized by w_k. Accordingly, solutions of (3) are actually characterized by the pair (u_k^*, w_k). We impose the following assumptions on the costs in (3).

Assumption 3. The following hold:

1) The map $u \to \nabla\phi(u; \theta_\phi)$ is ℓ_u-Lipschitz continuous with $\ell_u \geq 0$ for all $u \in \mathcal{U}$, for all $\theta_\phi \in \Omega_\phi$.
2) The map $y \to \nabla\psi(y; \theta_\psi)$ is ℓ_y-Lipschitz continuous with $\ell_y \geq 0$ for all $y \in \mathbb{R}^{n_y}$, for all $\theta_\psi \in \Omega_\psi$.
3) The map $u \to \phi(u; \theta_\phi) + \psi(Gu + Hw_k; \theta_\psi)$ is μ-strongly convex with $\mu > 0$ for all $w_k \in \mathbb{R}^{n_w}$, $\theta_\phi \in \Omega_\phi$, $\theta_\psi \in \Omega_\psi$. □

Assumption 3 guarantees that the gradient of the composite cost at equilibrium $\phi(u; \theta_\phi) + \psi(Gu + Hw_k; \theta_\psi)$ is ℓ-smooth with $\ell := \ell_u + \|G\|^2 \ell_y > 0$. Recall that $\theta_{\phi,k}$ and $\theta_{\psi,k}$ are the parameters of the cost functions at time k (and they may change over time); hereafter, for notational simplicity, we will use the short-hand notation:

$$\phi_k(u_k) = \phi(u_k, \theta_{\phi,k}), \tag{4a}$$

$$\psi_k(u_k) = \psi(Gu_k + Hw_k, \theta_{\psi,k}). \tag{4b}$$

In the following section, the design a data-assisted gradient-based control method.

C. Implicit Solution Tracking

We consider controllers of the form,

$$u_{k+1} = \Pi_{\mathcal{U}_k} \{u_k - \eta\Psi_k(u_k, y_k)\}, \tag{5}$$

where $\Psi_k(u_k, y_k) := \nabla\phi_k(u_k) + G^\top \nabla\psi_k(y_k)$ and with $\eta > 0$ a gain (also referred to as the step size in gradient methods). The controller (5) includes system outputs y_k rather than the system equilibrium $Gu_k + Hw_k$, which circumvents the need to measure the disturbance w_k. Importantly, this controller critically requires exact knowledge of the system output y_k and the gradients $\nabla\phi_k$ and $\nabla\psi_k$.

In this work, we consider the case of estimating the cost functions ϕ_k and ψ_k using trained NNs; these estimated functions are denoted by $\hat{\phi}_k$ and $\hat{\psi}_k$. To guarantee well-posedness of training, we assume that the networks are trained over the compact sets $\mathcal{U}_{\text{train}} := \bigcup_{k\in\mathbb{N}_{>0}} \mathcal{U}_k \times \Omega_\phi$ and $\mathcal{Y}_{\text{train}} \times \Omega_\psi$, respectively. Note that in general we may not have $\mathcal{U}_{\text{train}}$ compact, but for the purposes of our setting (namely, within engineering applications), this set may be assumed to be compact. The set $\mathcal{Y}_{\text{train}}$ is compact since $\mathcal{U}_{\text{train}} \times \mathcal{W}$ is compact and the output map is continuous [24]. The proposed algorithm is described next.

Algorithm 1 Optimization with NN Cost Perception

Training

Given: $\{(u_k^{(i)}, \theta_{\phi,k}^{(i)}), \phi_k(u_k^{(i)})\}_{i=1}^N, \{(y_k^{(i)}, \theta_{\psi,k}^{(i)}), \psi_k(y_k^{(i)})\}_{i=1}^M$

Obtain:

$$\hat{\phi}_k \leftarrow \text{NN-learning}(\{(u_k^{(i)}, \theta_{\phi,k}^{(i)}, \phi_k(u_k^{(i)}))\}_{i=1}^N)$$
$$\hat{\psi}_k \leftarrow \text{NN-learning}(\{(y_k^{(i)}, \theta_{\psi,k}^{(i)}, \psi_k(y_k^{(i)}))\}_{i=1}^M)$$

Gradient-based Feedback Control

Given: $x_0 \in \mathcal{X}_0$, $u_0 \in \mathcal{U}_0$, NN maps $\hat{\phi}_k, \hat{\psi}_k$.

For $k \geq 0$:

$$x_{k+1} = Ax_k + Bu_k + Ew_k, \; y_k = Cx_k + Dw_k \tag{6a}$$
$$\hat{\phi}_{g,k}(u_k) = \sum_{i=1}^{n_u} \frac{1}{2\varepsilon} \left(\hat{\phi}_k(u_k + \varepsilon b_i) - \hat{\phi}_k(u_k - \varepsilon b_i) \right) b_i \tag{6b}$$
$$\hat{\psi}_{g,k}(y_k) = \sum_{i=1}^{n_x} \frac{1}{2\varepsilon} \left(\hat{\psi}_k(y_k + \varepsilon d_i) - \hat{\psi}_k(y_k - \varepsilon d_i) \right) d_i \tag{6c}$$
$$u_{k+1} = \Pi_{\mathcal{U}_k}\Big\{ u_k - \eta \left(\hat{\phi}_k(u_k) + G^\top \hat{\psi}_k(y_k) \right) \Big\} \tag{6d}$$

In Algorithm 1, $\epsilon > 0$, b_i is the i^{th} canonical vector of $\mathbb{R}^{n_u}$, and d_i is the i^{th} canonical vector of $\mathbb{R}^{n_y}$; moreover, the operator NN-learning$(\cdot)$ refers to a training procedure for the NNs. The controller (6d) implements estimated gradients obtained via a centered difference approximation of the NN functions $\hat{\phi}_k$ and $\hat{\psi}_k$.

III. Main Results

For our analysis, we rewrite (6) as

$$x_{k+1} = Ax_k + Bu_k + Ew_k, \tag{7a}$$
$$y_k = Cx_k + Dw_k, \tag{7b}$$
$$u_{k+1} = \Pi_{\mathcal{U}_k} \{u_k - \eta \Psi_k(u_k, y_k) + e_k\} \tag{7c}$$

where (7a) is the plant and (7c) is the controller with e_k the gradient error defined as:

$$e_k := \eta \left(\Psi_k(u_k, y_k) - \hat{\phi}_{g,k}(u_k) - \hat{\psi}_{g,k}(y_k) \right). \tag{8}$$

Our main result is stated in terms of the following errors in the states of (6):

$$\omega_k := \begin{bmatrix} \|x_k - x_k^*\| \\ \|u_k - u_k^*\| \end{bmatrix}. \tag{9}$$

The main result is stated next.

Theorem 1. Let Assumptions 1-3 be satisfied. Let the controller gain be such that:

$$\eta \in \left(0, \min \left\{ \frac{2\mu}{\ell^2}, \frac{\bar{\lambda}(A) c_P + 1}{\ell_{x_u} \ell_y \|G\| \|C\|} \right\} \right), \tag{10}$$

where $\mu, \ell, \ell_{x_u}, \ell_y > 0$ are described in Assumptions 3. Then, there exists constants $r_{M_0} > 0$, $c_{M_0} \in [0, 1)$, $\bar{M} > 0$, and $\bar{N} > 0$ such that the following holds:

$$\|\omega_{k+1}\| \leq r_{M_0} (c_{M_0})^{k+1} \|\omega_0\| + r_{M_0} \frac{c_{M_0}}{1 + c_{M_0}} \left(\bar{M} \|\bar{\nu}\| + \bar{N} \sup_{0 \leq s \leq k} \|e_k\| \right), \tag{11}$$

where $\bar{\nu} := [\sup_{0 \leq s \leq k} \|w_{s+1} - w_s\|, \sup_{0 \leq s \leq k} \|u_{s+1}^* - u_s^*\|]^\top$.

The proof for Theorem 1 is given in the extended version on ArXiv. Theorem 1 guarantees that if Assumptions 1-3 are satisfied and a suitably sized controller gain η is chosen, then the interconnected system exponentially converges to the solution trajectory of the associated optimization problem up to a neighborhood characterized by error terms. The first error term, given by $\|\bar{\nu}\|$, corresponds to the temporal variability of the unknown disturbance w and the optimizer, u^*. Importantly, the time-variability of the optimizer u^* is influenced by both the time-varying cost functions and the unknown disturbance w; hence, both error terms within $\|\bar{\nu}\|$ are necessary. The second error term, given by $\sup_{0 \leq s \leq k} \|e_k\|$, corresponds to the difference between the nominal gradients and the estimated gradients that the controller (7c) obtained via a trained neural network. Moreover, the subsequent corollaries specifically tailor the bond of $\sup_{0 \leq s \leq k} \|e_k\|$ to different types of neural networks.

Corollary 1 (Feedforward NN). Consider the system (7). Let Assumptions 1-3 be satisfied and let $\eta > 0$ be chosen so that (10) is satisfied. Suppose that the feedforward networks $\hat{\phi}_{F,k}$ and $\hat{\psi}_{F,k}$ approximate the costs ϕ_k and ψ_k over the sets $\mathcal{U}_{\text{train}} := \bigcup_{k \in \mathbb{N}_{>0}} \mathcal{U}_k$ and $\mathcal{Y}_{\text{train}} \subset \mathbb{R}^{n_y}$, respectively. Then the error (9) satisfies (11) with e_k as,

$$\|e_k\| \leq \delta_{F,u} + n_u \varepsilon^{-1} \sup_{u_k \in \mathcal{U}_{\text{train}}} |\phi_k(u_k) - \hat{\phi}_k(u_k)| + \|G\| \delta_{F,y} + n_y \varepsilon^{-1} \|G\| \sup_{y_k \in \mathcal{Y}_{\text{train}}} |\psi_k(y_k) - \hat{\phi}_k(y_k)|, \tag{12}$$

where $\delta_{F,u} := \sup_{u_k \in \mathcal{U}_{\text{train}}} \|\nabla \phi_k(u_k) - \sum_{i=1}^{n_u} \frac{1}{2\varepsilon} (\phi_k(u_k + \varepsilon b_i) - \phi_k(u_k - \varepsilon b_i)) b_i\|$ and $\delta_{F,y} := \sup_{y_k \in \mathcal{Y}_{\text{train}}} \|\nabla \psi_k(y_k) - \sum_{i=1}^{n_y} \frac{1}{2\varepsilon} (\psi_k(y_k + \varepsilon d_i) - \psi_k(y_k - \varepsilon d_i)) d_i\|$.

In Corollary 1, $\delta_{F,u}$ and $\delta_{F,y}$ are bounds on the centered difference approximations for ϕ_k and ψ_k. Further, the bound (12) explicitly characterizes the neighborhood of convergence of the optimizer in terms of the uniform approximation ability of the feedforward net and the accuracy of the centered difference method. The proof of Corollary 1 is given in [21, Prop. 3]. In the following, we identify an alternative bound for the error e_k if residual networks are used instead. To do so, we must consider the lifted counterparts $\tilde{\phi}_k$ and $\tilde{\psi}_k$ for the functions ϕ_k and ψ_k given by,

$$\tilde{\phi}_k = \iota_\phi(z) \circ \phi_k, \; \iota_\phi(z) = (z, 0, \ldots, 0), \; z \in \mathbb{R}, \tag{13a}$$
$$\tilde{\psi}_k = \iota_\psi(z) \circ \psi_k, \; \iota_\psi(z') = (z', 0, \ldots, 0), \; z' \in \mathbb{R}, \tag{13b}$$

where $\iota_\phi : \mathbb{R} \to \mathbb{R}^{n_u}$ and $\iota_\psi : \mathbb{R} \to \mathbb{R}^{n_y}$ are injections. Now, we characterize results for e_k if we leverage residual networks $\hat{\phi}_{R,k}$ and $\hat{\psi}_{R,k}$ to estimate $\tilde{\phi}_k$ and $\tilde{\psi}_k$.

Corollary 2 (Residual NN). Consider the system (7). Let Assumptions 1-3 be satisfied and let $\eta > 0$ be chosen so that (10) is satisfied. Suppose that the residual networks $\hat{\phi}_{R,k}$ and $\hat{\psi}_{R,k}$ approximate the costs $\tilde{\phi}_k$ and $\tilde{\psi}_k$ over the sets $\mathcal{U}_{\text{train}} := \bigcup_{k\in\mathbb{N}_{>0}} \mathcal{U}_k$ and $\mathcal{Y}_{\text{train}} \subset \mathbb{R}^{n_y}$, respectively. Let the set of training points $\mathcal{U}_{\text{train},s}$ (resp., $\mathcal{Y}_{\text{train},s}$) be a ρ_u-cover (ρ_y-cover) of $\mathcal{U}_{\text{train}}$ ($\mathcal{Y}_{\text{train}}$) with respect to the partial order $\preceq$ for some $\rho_u > 0$ ($\rho_y > 0$). Suppose that the NNs can be decomposed as $\hat{\phi}_{R,k} = m_u + A_u$ (resp., $\hat{\psi}_{R,k} = m_y + A_y$) where $m_u : \mathbb{R}^{n_u} \to \mathbb{R}^{n_u}$ and $m_y : \mathbb{R}^{n_y} \to \mathbb{R}^{n_y}$ are monotone, and A_u, A_y are linear functions. Then, the error (9) satisfies (11) with e_k as,

$$\begin{aligned} \|e_k\| \leq \delta_{F,u} + \|G\|\delta_{F,y} + \\ + n_u^{3/2}\varepsilon^{-1}(3\delta_{u,tr} + 2\omega_\phi(\rho_u) + 2\|A_u\|_\infty) \\ + n_y^{3/2}\varepsilon^{-1}\|G\|(3\delta_{y,tr} + 2\omega_\psi(\rho_x) + 2\|A_x\|_\infty), \end{aligned} \tag{14}$$

where $\delta_{F,u}$ and $\delta_{F,y}$ are given in Corollary 1, ω_ϕ and ω_ψ are moduli of continuity for $\tilde{\phi}_k$ and $\tilde{\psi}_k$, and

$$\delta_{u,tr} := \sup_{u_k \in \mathcal{U}_{\text{train},s}} \|\tilde{\phi}_k(u_k) - \hat{\phi}_{k\,R,k}(u_k)\|_\infty,$$

$$\delta_{y,tr} := \sup_{y_k \in \mathcal{Y}_{\text{train},s}} \|\tilde{\psi}_k(y_k) - \hat{\psi}_{k\,R,k}(y_k)\|_\infty.$$

Corollary 2 clearly identifies the error e_k if a residual NN is used. In contrast to Corollary 1, the use of a residual NN characterizes the error in terms of the set of samples $\mathcal{U}_{\text{train},s}$ and $\mathcal{Y}_{\text{train},s}$. The proof of Corollary 2 is given in [21, Prop. 3-4].

IV. Numerical Simulations

We illustrate the performance of the proposed controller in the context of a ride-service scheduling application. By borrowing the setup in [22], we consider a ride service provider (RSP) that seeks to maximize its fleet utilization profit by dispatching electric vehicles to serve ride requests from customers. The area of interest is modeled as a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ represent a region, and the edge $(i,j) \in \mathcal{E}$ indicates that a ride from node $i \in \mathcal{V}$ to node $j \in \mathcal{V}$ is allowed. As a simple example, we consider $n = 4$ regions, and the time is discretized as $\Delta = 5\,\text{min}$. The elasticity of the price u_k^{ij} from region i to j at time k is modeled as follows:

$$u_k^{ij} = \frac{p_{\max}^{ij}}{\theta^{ij}}\left(1 - \frac{d_k^{ij}}{\delta^{ij}}\right), \tag{15}$$

where $\delta_k^{ij} \in \mathbb{R}_{\geq 0}$ is the demand of rides from region i to region j at time $k \in \mathbb{Z}_{>0}$, $d_k^{ij} \in \mathbb{R}_{\geq 0}$ denotes the demand from region i to j at time k, $\theta^{ij} \in [0,1]$ represents the steepness of elasticity, and $p_{\max}^{ij} \in \mathbb{R}_{>0}$ is an upper limit on prices from i to j. The states $x_k^i \in \mathbb{R}_{\geq 0}$ denote the idle-vehicle occupancy of the fleet in region i at time k. We denote by $a_{ij} \in \mathbb{R}_{\geq 0}$ the fraction of unoccupied vehicles that travel from i to j at every time step. The travel times can vary over time and we model them by a Boolean variable as: $\sigma_k^{ij,\tau} = 1$ if travel time form i to j at time k is τ slots and, 0 otherwise, for all $i,j \in \mathcal{V}$ and k, $\tau \in \mathbb{Z}_{\geq 0}$.

Therefore, the discrete dynamic associated with the occupancy of idle vehicles in each region i is:

$$\begin{aligned} x_{k+1}^i = x_k^i - \sum_{j\in\mathcal{V}} a_{ij}x_k^i + \sum_{j\in\mathcal{V}} a_{ji}x_k^j \\ - \sum_{j\in\mathcal{V}} d_k^{ij} + \underbrace{\sum_{j\in\mathcal{V}} \sum_{\tau=k-T}^{k-1} \sigma_\tau^{ij,k-\tau} d_\tau^{ji} + g_k^i}_{:=w_k^i} \end{aligned} \tag{16}$$

where $d_k^{ij} = \delta_k^{ij}\left(1 - \theta^{ij}\frac{u_k^{ij}}{p_{\max}^{ij}}\right)$.

Then, the RSP's maximization problem at every k can be expressed as:

$$\begin{aligned} \max_{u,\,x} \sum_{i\in\mathcal{V}}\sum_{j\in\mathcal{V}} \underbrace{u^{ij}\delta^{ij}\left(1 + \frac{\theta^{ij}c^{ij}}{p_{\max}^{ij}}\right) - {u^{ij}}^2\frac{\theta^{ij}\delta^{ij}}{p_{\max}^{ij}}}_{:=\phi^{ij}(u^{ij})} - c^{ij}\delta^{ij} - \rho\|x\|^2 \\ \text{s. to: } 0 = -\sum_{j\in\mathcal{V}} a_{ij}x^i + \sum_{j\in\mathcal{V}} a_{ji}x^j - \sum_{j\in\mathcal{V}} d^{ij} + w^i, \\ x^i \geq 0, \quad \forall i,j,\in\mathcal{V}, \end{aligned} \tag{17}$$

where $u = [u^{ij}]$, $x = [x^i]$ for all $i,j \in \mathcal{V}$, $c^{ij} \in \mathbb{R}_{>0}$ is the cost of routing vehicles from i to j, $\rho\|x\|^2$ with $\rho \in \mathbb{R}_{>0}$ describes the RSP's objective of maximizing fleet utilization, and $\phi^{ij}(u^{ij})$ is an unknown function that depends of the price elasticity, u^{ij}.

We solve the optimization problem (17) via Algorithm 1. Figure 1 presents the percentage of idle vehicle occupancy x^i per region i when (a) the true function $\phi^{ij}(u^{ij})$ is used, (b) $\hat{\phi}^{ij}(u^{ij})$ is learned via NN. We find that the percentage of idle vehicle occupancy stabilizes asymptotically around 10% for all regions, where errors persist due to the learning error and the need for a finite difference approximation. Similarly, Figure 2 shows the behavior of the control variables u_{ij} for both cases. Finally, in Figure 3 we quantify the relative error of the control inputs with respect to the case where the true cost function versus the estimated cost functions are used, and we find a satisfying convergence to a neighborhood within 0.5% relative error.

V. Conclusion

We have designed a projected gradient based feedback controller that regulates a physical system's outputs to the solution of a constrained, convex optimization problem. We have also augmented our controller with a regression-based NN to learn the cost functions using a sufficient amount of historical data. We established guarantees for the asymptotic convergence of the system up to a neighborhood based on the learning error and external disturbances using suitable bounds for the controller gain.

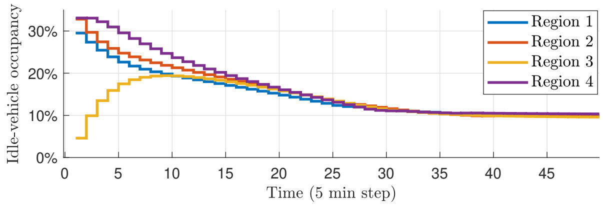

(a)

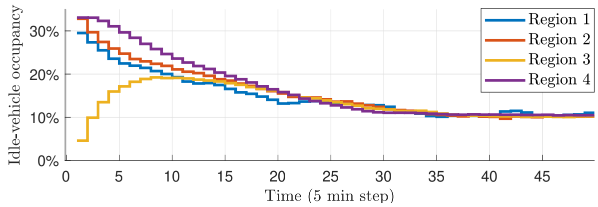

(b)

Fig. 1. States per region (a) True gradient (b) Estimated gradient using feedforward neural networks.

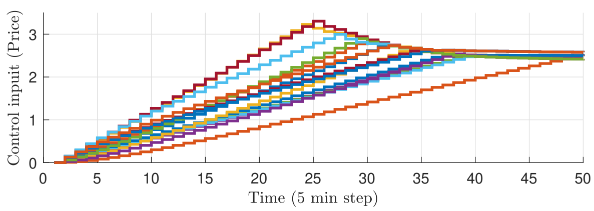

(a)

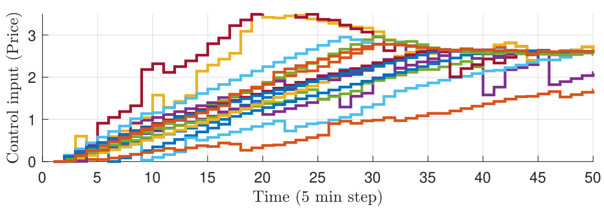

(b)

Fig. 2. Control inputs (a) True gradient (b) Estimated gradient using feedforward neural networks.

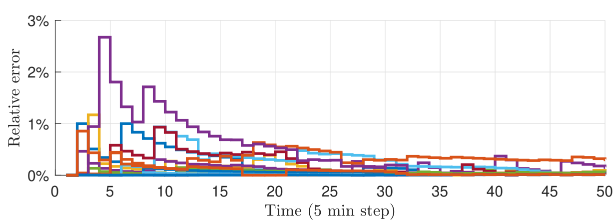

Fig. 3. Control input relative error.

References

[1] A. Simonetto, E. Dall'Anese, J. Monteil, and A. Bernstein, "Personalized optimization with user's feedback," Automatica, vol. 131, p. 109767, 2021.

[2] 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.

[3] I. Notarnicola, A. Simonetto, F. Farina, and G. Notarstefano, "Distributed personalized gradient tracking with convex parametric models," 2021.

[4] A. M. Ospina, A. Simonetto, and E. Dall'Anese, "Time-varying optimization of networked systems with human preferences," IEEE Transactions on Control of Network Systems, pp. 1–12, 2022.

[5] A. Jokic, M. Lazar, and P. P.-J. Van Den Bosch, "On constrained steady-state regulation: Dynamic KKT controllers," IEEE Trans. on Automatic control, vol. 54, no. 9, pp. 2250–2254, 2009.

[6] K. Hirata, J. P. Hespanha, and K. Uchida, "Real-time pricing leading to optimal operation under distributed decision makings," in American Control Conf., 2014.

[7] F. D. Brunner, H.-B. Dürr, and C. Ebenbauer, "Feedback design for multi-agent systems: A saddle point approach," in IEEE Conference on Decision and Control, 2012, pp. 3783–3789.

[8] M. Colombino, E. Dall'Anese, and A. Bernstein, "Online optimization as a feedback controller: Stability and tracking," IEEE Trans. On Control of Networked Systems, vol. 7, no. 1, pp. 422–432, 2020.

[9] T. Zheng, J. Simpson-Porco, and E. Mallada, "Implicit trajectory planning for feedback linearizable systems: A time-varying optimization approach," in American Control Conference, 2020, pp. 4677–4682.

[10] A. Hauswirth, S. Bolognani, G. Hug, and F. Dörfler, "Timescale separation in autonomous optimization," IEEE Trans. on Automatic Control, vol. 66, no. 2, pp. 611–624, 2021.

[11] L. S. Lawrence, J. W. Simpson-Porco, and E. Mallada, "Linear-convex optimal steady-state control," IEEE Transactions on Automatic Control, vol. 66, no. 11, pp. 5377–5384, 2020.

[12] G. Belgioioso, D. Liao-McPherson, M. H. de Badyn, S. Bolognani, J. Lygeros, and F. Dorfler, "Sampled-data online feedback equilibrium seeking: Stability and tracking," in IEEE Conference on Decision and Control, dec 2021.

[13] A. Agarwal, J. W. Simpson-Porco, and L. Pavel, "Game-theoretic feedback-based optimization," IFAC-PapersOnLine, vol. 55, no. 13, pp. 174–179, 2022.

[14] M. Nonhoff and M. A. Müller, "An online convex optimization algorithm for controlling linear systems with state and input constraints," in 2021 American Control Conference (ACC). IEEE, 2021, pp. 2523–2528.

[15] N. Agarwal, E. Hazan, and K. Singh, "Logarithmic regret for online control," Advances in Neural Information Processing Systems, vol. 32, 2019.

[16] E. Minasyan, P. Gradu, M. Simchowitz, and E. Hazan, "Online control of unknown time-varying dynamical systems," Advances in Neural Information Processing Systems, vol. 34, pp. 15 934–15 945, 2021.

[17] B. Karg and S. Lucia, "Stability and feasibility of neural network-based controllers via output range analysis," in IEEE Conference on Decision and Control, 2020, pp. 4947–4954.

[18] H. Yin, P. Seiler, and M. Arcak, "Stability analysis using quadratic constraints for systems with neural network controllers," 2021, (Early Access).

[19] M. Marchi, J. Bunton, B. Gharesifard, and P. Tabuada, "Safety and stability guarantees for control loops with deep learning perception," IEEE Control Systems Letters, vol. 6, pp. 1286–1291, 2022.

[20] Z.-P. Jiang and Y. Wang, "Input-to-state stability for discrete-time nonlinear systems," Automatica, vol. 37, no. 6, pp. 857–869, 2001.

[21] L. Cothren, G. Bianchin, and E. Dall'Anese, "Online optimization of dynamical systems with deep learning perception," IEEE Open Journal of Control Systems, vol. 1, pp. 306–321, 2022.

[22] G. Bianchin, M. Vaquero, J. Cortes, and E. Dall'Anese, "Online stochastic optimization for unknown linear systems: Data-driven synthesis and controller analysis," 2021. [Online]. Available: https://arxiv.org/abs/2108.13040

[23] M. Marchi, B. Gharesifard, and P. Tabuada, "Training deep residual networks for uniform approximation guarantees," in Conference on Learning for Dynamics and Control, ser. Proceedings of Machine Learning Research, vol. 144, 2021, pp. 677–688.

[24] W. Rudin, Principles of mathematical analysis. McGraw-hill New York, 1976, vol. 3.