The Discrete-time Internal Model Principle of Time-varying Optimization: Limitations and Algorithm Design

Gianluca Bianchin Bryan Van Scoy

***Abstract*— Time-varying optimization problems arise in a variety of engineering applications. The available information about how the problem changes in time dictates the types of algorithms that are applicable to a particular problem as well as the types of convergence guarantees that may be proven. In this paper, we study dynamic gradient-feedback algorithms for time-varying optimization in discrete time. By casting the design of such algorithms as an output regulation problem for dynamical systems, we provide necessary and sufficient conditions for the existence of a gradient-feedback algorithm that asymptotically tracks a critical trajectory of the optimization problem. When these conditions hold, we provide a design procedure to construct such an algorithm. As a fundamental limitation, we show that any algorithm that asymptotically tracks a critical trajectory needs to contain an internal model of the temporal variation, which we refer to as the *internal model principle of time-varying optimization*.**

I. Introduction

Optimization problems throughout engineering often contain parameters that vary in time, leading to the setting of time-varying optimization [1]. Applications of time-varying optimization include optimal power flow in power grids with renewable energy sources, obstacle avoidance in robotics using barrier functions [2], model predictive control [3], feature extraction in videos, magnetic resonance imaging (MRI) with high-definition video, real-time optimization for chemical and industrial processes, and others; see [1, Sections IV and V] and the references therein.

The algorithms available to solve a time-varying optimization problem depend on the available information about how the problem changes in time. For instance, suppose one has access to the optimization problem at each point in time, but has no foreknowledge as to how the problem will change at the next iteration. In this case, any method from static optimization (e.g., [1], [4]–[9]) may be applied directly to the problem at each point in time, but such algorithms can only achieve convergence to a neighborhood of a critical trajectory, where the size of the neighborhood depends on the convergence properties of the algorithm as well as how quickly the problem varies in time [10], [11].

The optimization problem, however, is not always known at each point in time (see, e.g., [12], [13]). Instead, suppose one has access to an oracle for the optimization problem (such as the gradient of the objective function) along with a model for how the optimization problem varies in time. In this case, the algorithm may exploit this information to asymptotically track a critical trajectory. This is the approach proposed in [14] for discrete-time problems and [15] for continuous-time ones. In [14], the algorithm has access to the gradient of the objective function along with knowledge of the poles of the z-transform of the time-varying parameter. Based on the internal model principle, this model of the time variation is then incorporated in the algorithm to achieve exact asymptotic tracking of the optimal trajectory. These results, however, are limited to quadratic objective functions with linear temporal variabilities.

Departing from these early works, in this paper we study discrete-time time-varying optimization problems, and we pose the following questions:

1) What is the minimal amount of information needed to design an algorithm that asymptotically tracks a critical trajectory of a time-varying optimization problem?
2) When these conditions hold, how does one design such an algorithm?

To address these questions, we cast the analysis and design of a time-varying optimization algorithm as a nonlinear output regulation problem [16], which can be studied using tools from center manifold theory for maps [17]–[19]. Our main contributions are as follows:

1) We provide necessary and sufficient conditions for a discrete-time gradient-based algorithm to asymptotically track a critical trajectory of a time-varying optimization problem (Theorems 2 and 3).
2) When these conditions hold, we provide a design procedure to construct such an algorithm (Algorithm 1). The algorithm consists of an observer combined with a function that zeros the gradient of the objective function (see Definition 2).

As a fundamental result, we show that, to achieve asymptotic tracking of a critical trajectory, it is necessary to have some knowledge of the temporal variability in the optimization and for this variability to be "observable" by the algorithm (see Thm. 2). In turn, this requires the algorithm to contain an internal model of the temporal variation; we refer to this as the *discrete-time internal model principle of time-varying optimization*, akin to its counterpart in control [20].

The rest of the paper is organized as follows. We begin by formulating the problem in §II. We consider the special case of parameter feedback optimization in §III, followed by the

G. Bianchin is with the ICTEAM institute and the Department of Mathematical Engineering (INMA) at the University of Louvain, Belgium. Email: gianluca.bianchin@uclouvain.be

B. Van Scoy is with the Dept. of Electrical and Computer Engineering, Miami University, OH 45056, USA. Email: bvanscoy@miamioh.edu

general case of dynamic feedback in §IV. Simulation results are presented in §V, and §VI summarizes our conclusions.

II. Problem setting

We consider the time-varying optimization problem:

$$\min_{x \in \mathbb{R}^n} f(x, \theta_k), \tag{1}$$

where $k \in \mathbb{N}_{\geq 0}$ denotes time or iteration, and $f : \mathbb{R}^n \times \Theta \to \mathbb{R}$, with $\Theta \subseteq \mathbb{R}^p$, is a loss function that is parametrized by the time-varying parameter vector $\theta : \mathbb{N}_{\geq 0} \to \Theta$.

A. Standing assumptions

We list hereafter the basic assumptions on which our approach to gradient-feedback theory is based.

Assumption 1 (Properties of the objective function)**.** The map $x \mapsto f(x, \theta)$ is convex and $x \mapsto \nabla_x f(x, \theta)$ is Lipschitz continuous in $\mathbb{R}^n$, for each $\theta \in \Theta$. □

Assumption 2 (Existence of an exosystem)**.** There exists a smooth (i.e., C^∞) vector field $s : \Theta \to \Theta$ and initial condition $\theta_0 \in \Theta$ such that the parameter vector satisfies

$$\theta_{k+1} = s(\theta_k), \tag{2}$$

for all $k \in \mathbb{N}_{\geq 0}$. □

Assumption 3 (Stability of the exosystem)**.** The equilibrium $\theta = 0$ of the exosystem (2) is locally Lyapunov stable. □

Convexity and smoothness are standard assumptions in optimization [21]. Assumption 2 ensures the existence of an autonomous system, called *exosystem*, describing the class of temporal variabilities of the cost taken into consideration. Notice that the vector field $s(\theta)$ may or may not be known in the applications, and that our goal hereafter is to study to what resolution (1) can be solved in relation to the available knowledge on $s(\theta)$. The class of exosystems satisfying Assumption 3 includes the (important in practice) set of systems in which every solution is periodic.

In this work, we study the problem of designing an optimization algorithm that computes and tracks a *critical trajectory* of (1), which is a map $x^\star : \mathbb{N}_{\geq 0} \to \mathbb{R}^n$ that satisfies[1]

$$0 = \nabla_x f(x_k^\star, \theta_k), \quad \forall k \in \mathbb{N}_{\geq 0}. \tag{3}$$

Remark 1. The time variation in (1) is captured implicitly through the parameter vector θ_k. A related, yet slightly more general, problem makes the time dependency explicit:

$$\min_{x \in \mathbb{R}^n} f_0(x, k). \tag{4}$$

While Problem (1) can be cast uniquely as in (4) (by letting $f_0(x, k) = f(x, \theta_k)$ for all x and k), in general, there exists an infinite number of ways to parametrize (4) as in (1), thus leading to possible ambiguities. For instance, any $f_0(x, k)$ may be parametrized by $\theta_k = k$ (so that $f_0 \equiv f$), although this is not compatible with Assumption 3. □

[1] Existence of a critical trajectory is implied by Assumption 1.

B. Algorithm structure

Our objective is to seek an optimization algorithm that assumes no access to θ_k. Instead, as is common in first-order optimization approaches [21], we will assume only the availability of functional evaluations of the gradient function $\nabla_x f(x, \theta_k)$ at points $x \in \mathbb{R}^n$, selected by the algorithm. Formally, the optimization algorithm is described by an internal state z_k, which takes values on an open set $\mathcal{Z} \subseteq \mathbb{R}^{n_c}$ with $n_c \in \mathbb{N}_{>0}$. The optimization algorithm generates a sequence of points $x_k \in \mathbb{R}^n$ (called *exploration signal*) at which the gradient shall be evaluated, and processes functional evaluations of the gradient at these points $y_k = \nabla_x f(x_k, \theta_k)$ (called *gradient feedback signal*). Mathematically, the optimization algorithm is described by:

$$z_{k+1} = F_c(z_k, y_k), \qquad x_k = G_c(z_k), \tag{5a}$$

together with the *gradient-feedback signal*:

$$y_k = \nabla_x f(x_k, \theta_k), \tag{5b}$$

where $F_c : \mathcal{Z} \times \mathbb{R}^n \to \mathcal{Z}$ and $G_c : \mathcal{Z} \to \mathbb{R}^n$ are functions to be designed. In the remainder, we refer to (5) as a *dynamic gradient-feedback optimization algorithm*. Notice that the dynamics of the optimization algorithm (5), coupled with the time-variability generator (2), have the form of a nonlinear autonomous system:

$$z_{k+1} = F_c(z_k, y_k), \tag{6a}$$

$$y_k = \nabla_x f(G_c(z_k), \theta_k), \tag{6b}$$

$$\theta_{k+1} = s(\theta_k). \tag{6c}$$

Our objective is to design $F_c(z, y)$, $G_c(z)$, and n_c so that $y_k \to 0$ as $k \to \infty$, which ensures that x_k tracks, with zero asymptotic error, a critical trajectory $x_k^\star$ of (1).

We will assume that θ_k takes values in a neighborhood of the origin, and thus let Θ be some neighborhood of the origin of $\mathbb{R}^p$. Note that there is no loss of generality in doing so, because if θ_k takes values in the neighborhood of any other point, the former can be shifted to the origin via a change of variables without altering the critical points of (1). In what follows, we denote by $x_\circ^\star \in \mathbb{R}^n$ a point such that

$$0 = \nabla_x f(x_\circ^\star, 0), \tag{7}$$

and assume that $x_\circ^\star \in \mathbb{R}^n$ is locally unique. Moreover, we will assume that the functions $F_c(z, y)$ and $G_c(z)$ to be designed are such that

$$F_c(z_\circ^\star, 0) = z_\circ^\star, \qquad x_\circ^\star = G_c(z_\circ^\star), \tag{8}$$

for some locally unique $z_\circ^* \in \mathcal{Z}$. This ensures that the optimization algorithm (5) has an equilibrium at $z = z_\circ^\star$, and that the corresponding gradient feedback signal is identically zero at this point: $y_k = \nabla_x f(G_c(z_\circ^\star), 0) = 0$.

Definition 1. We say that (6) *asymptotically tracks a critical trajectory of* (1) *with respect to initializations in the set* $\Theta_\circ \subseteq \Theta$ if, for each initial condition (z_0, θ_0) with z_0 in some neighborhood of $z_\circ^\star$ and $\theta_0 \in \Theta_\circ$, the solution of (6) satisfies $y_k \to 0$ as $k \to \infty$. □

In practice, the initial condition θ_0 to the exosystem (2) may not be known; Definition 1 accounts for such uncertainty by allowing θ_0 to be anywhere in the set Θ_0. Observe also that, when an algorithm asymptotically tracks a critical trajectory, we have $x_k \to x_k^\star$ for some $x_k^\star$ as in (3). Namely, the exploration signal converges asymptotically to a critical trajectory. We now make our objective formal.

Problem 1 (Dynamic gradient-feedback problem)**.** Find necessary and sufficient conditions (in terms of the loss function f) for the existence of a gradient-feedback optimization algorithm as in (5) that asymptotically tracks a critical trajectory of (1) with respect to initializations in some Θ_0. When these conditions hold, derive a method to design such an algorithm. □

It is important to note that existence conditions for $F_c(z, y)$ and $G_c(z)$ shall not depend on the temporal variability of the optimization (namely, $s(\theta)$), but only on the properties of the loss function. Conversely, if such conditions were to depend on $s(\theta)$, a more general class of optimization algorithms than (5) could be constructed, implying that the algorithm formulation would not be sufficiently general.

III. The parameter-feedback problem

To address Problem 1, we first study a simpler problem that allows us to derive the necessary framework to tackle our objectives in their generality. To this end, in place of the dynamic optimization algorithm (5), we begin by considering an algebraic optimization algorithm of the form:

$$x_k = H_c(\theta_k), \tag{9}$$

where $H_c : \Theta \to \mathbb{R}^n$ is a mapping to be designed; we will require that H_c is of class C^0 and satisfies the fixed-point condition $x_\circ^\star = H_c(0)$ (cf. (8)). Because of the explicit dependence on θ_k, we will refer to (9) as a *static parameter-feedback* optimization algorithm. Our objective is to design the map H_c so that the composition of (2), (5b), and (9):

$$\begin{aligned} y_k &= \nabla_x f(H_c(\theta_k), \theta_k), \\ \theta_{k+1} &= s(\theta_k), \end{aligned} \tag{10}$$

tracks, with zero asymptotic error, a critical trajectory of (1). For this framework, we reformulate Problem 1 as follows.

Problem 2 (Static parameter-feedback problem)**.** Find necessary and sufficient conditions (in terms of the loss function f) for the existence of a parameter-feedback optimization algorithm as in (9) that asymptotically tracks a critical trajectory of (1) with respect to initializations in some Θ_0. When these conditions hold, derive a method to design such an algorithm. □

Solvability of the static parameter-feedback problem will depend on the existence of a function that zeros the gradient.

Definition 2 (Mapping zeroing the gradient)**.** A mapping $H_c : \Theta \to \mathbb{R}^n$ *zeros the gradient* at the point $\theta \in \Theta$ if

$$0 = \nabla_x f(H_c(\theta), \theta). \tag{11}$$

Moreover, H_c zeros the gradient *on a set* $\Theta_\circ \subseteq \Theta$ if (11) holds for all $\theta \in \Theta_\circ$. □

The following definition is instrumental.

Definition 3 (Limit point and limit set)**.** A point $\theta_\omega \in \Theta$ is said a *limit point with respect to the initialization* $\theta_\circ \in \Theta$ if there exists a sequence $\{k_i\}_{i \in \mathbb{N}_{\geq 0}}$, with $k_i \to \infty$ as $i \to \infty$, such that the trajectory of (2) with $\theta_0 = \theta_\circ$ satisfies $\theta_{k_i} \to \theta_\omega$ as $i \to \infty$. For $\theta_\circ \in \Theta$, let $\Omega(\theta_\circ)$ denote the set of all limit points (i.e., for all sequences $\{k_i\}_{i \in \mathbb{N}_{\geq 0}}$) of (2) with respect to the initialization $\theta_\circ$. Given $\Theta_\circ \subseteq \theta$, the set $\Omega(\Theta_\circ) := \cup_{\theta_\circ \in \Theta_\circ} \Omega(\theta_\circ)$ is called the *limit set with respect to initializations in* $\Theta_\circ$ [22]. □

Intuitively, $\Omega(\Theta_\circ)$ denotes the set of all limit points (equilibria, limit cycles, etc.) that can be reached by the exosystem when initialized at points in $\Theta_\circ$. Notice also that, by Assumption 3, $\Omega(\Theta_\circ)$ is contained in some neighborhood of the origin of $\mathbb{R}^p$. For example, if the exosystem (2) is linear and the origin is a Lyapunov stable equilibrium, then $\Omega(\Theta_\circ)$ is some neighborhood of the origin, whose radius depends on the radius of the initialization set $\Theta_\circ$. The following result characterizes all parameter-feedback optimization algorithms that achieve asymptotic tracking of a critical trajectory.

Theorem 1 (Parameter-feedback algorithm characterization)**.** Let Assumptions 1, 2, and 3 hold, and let $\Theta_\circ \subseteq \Theta$. The parameter-feedback algorithm (10) asymptotically tracks a critical trajectory of (1) with respect to initializations in $\Theta_\circ$ if and only if the mapping H_c zeros the gradient on $\Omega(\Theta_\circ)$.

□

Proof. We show that, for each initialization $\theta_0 \in \Theta_\circ$, the parameter-feedback algorithm asymptotically tracks a critical trajectory of (1) with respect to θ_0 if and only H_c zeros the gradient on $\Omega(\theta_0)$, from which the result follows.

Lyapunov stability of the exosystem (Assumption 3) implies that the trajectory is bounded and therefore has a subsequence that converges to some limit point $\theta_\omega \in \Omega(\theta_0)$. Then by definition, there exists an increasing sequence k_i such that the trajectory θ_{k_i} converges to θ_ω as $i \to \infty$. By continuity of the gradient (Assumption 1) and that of H_c,

$$\lim_{i \to \infty} y_{k_i} = \lim_{i \to \infty} \nabla_x f(H_c(\theta_{k_i}), \theta_{k_i}) = \nabla_x f(H_c(\theta_\omega), \theta_\omega).$$

When $y_k \to 0$ as $k \to \infty$, the left-hand side is zero, which implies that H_c zeros the gradient on θ_ω. Since this holds for any limit point, H_c zeros the gradient on $\Omega(\theta_0)$.

Now suppose H_c zeros the gradient on $\Omega(\theta_0)$. The right-hand side of the above equation is then zero, which implies the existence of a sequence k_i such that $y_{k_i} \to 0$ as $i \to \infty$. Since this holds for any limit point $\theta_\omega \in \Omega(\theta_0)$, any convergent subsequence of y_k converges to zero. Moreover, y_k is bounded due to Lipschitz continuity of the gradient, so y_k also converges to zero as $k \to \infty$. □

Intuitively, the theorem states that the parameter-feedback algorithm asymptotically tracks a critical trajectory if and only if the mapping H_c is chosen so that it zeros the gradient on the limit set of the exosystem, $\Omega(\Theta_\circ)$. Once a mapping $H_c(\theta)$ zeroing the gradient is known, a parameter-feedback optimization algorithm solving Problem 2 is given by (9).

Remark 2 (Knowledge of $\Omega(\Theta_\circ)$)**.** Notice that, when our goal is to design a parameter-feedback algorithm as in (9), such goal can be accomplished without an exact knowledge of the limit set $\Omega(\Theta_\circ)$. Indeed, it follows from the Lyapunov stability assumption (cf. Assumption 3) and the sufficiency part of the proof of Theorem 1 that if H_c zeros the gradient on some subset of the origin of $\mathbb{R}^p$ containing $\Omega(\Theta_\circ)$, then the choice (9) ensures that $y_k \to 0$ as $k \to \infty$. □

We conclude this section by illustrating the design procedure for parameter-feedback algorithms on a quadratic problem.

Example 1. Consider an instance of (1) with quadratic cost and time-variability that depends linearly on θ_k:

$$f(x, \theta_k) = \tfrac{1}{2} x^\mathsf{T} R x + x^\mathsf{T} Q \theta_k, \tag{12}$$

with $R \in \mathbb{S}^n$ and $Q \in \mathbb{R}^{n \times p}$. Notice that (12) admits a critical point for arbitrary θ_k if and only if $\operatorname{Im} Q \subseteq \operatorname{Im} R$, in which case $x_k^\star$ is unique. In this case, designing an optimization algorithm amounts to finding x_k such that we regulate to zero the signal:

$$y_k = \nabla_x f(x_k, \theta_k) = R x_k + Q \theta_k.$$

Applying Theorem 1 requires finding a linear transformation $H_c(\theta) = H_c \theta, H_c \in \mathbb{R}^{n \times p}$, such that $0 = (R H_c + Q)\theta$ for all θ in some neighborhood of the origin. Using $\operatorname{Im} Q \subseteq \operatorname{Im} R$, we can choose $H_c = -R^\dagger Q$, where $R^\dagger$ is the pseudo-inverse of R; observe that this choice for $H_c(\theta)$ zeros the gradient globally in $\mathbb{R}^p$. By substituting into (10), we have

$$y_k = R H_c \theta_k + Q \theta_k = 0, \quad \forall k \in \mathbb{N}_{\geq 0}.$$

Namely, the gradient is identically zero at all times. Interestingly, this behavior originates for two reasons: (i) θ_k is measurable at each k, and (ii) $H_c(\theta)$ obtained for this particular problem zeros the gradient on the entire $\mathbb{R}^p$ (not just some limit set of the trajectories of θ). When one of these two properties fails (as in Section IV, shortly below), this behavior can no longer be expected. □

IV. The dynamic gradient-feedback problem

In this section, we tackle the dynamic gradient-feedback problem (Problem 1).

A. Fundamental results

We begin with the following instrumental characterization.

Theorem 2 (Gradient-feedback algorithm characterization)**.** Suppose Assumptions 1, 2, and 3 hold, assume that $F_c(z, y)$ and $G_c(z)$ are such that the equilibrium $z = z_\circ^\star$ of

$$z_{k+1} = F_c(z_k, \nabla_x f(G_c(z_k), 0)),$$

is locally exponentially stable. The gradient-feedback optimization algorithm (6) asymptotically tracks a critical trajectory of (1) with respect initializations in $\Theta_\circ$ if and only if there exists a C^2 mapping $z = \sigma(\theta)$ with $\sigma(0) = z_\circ^\star$, defined on $\Omega(\Theta_\circ)$, which satisfies:

$$\sigma(s(\theta_\omega)) = F_c(\sigma(\theta_\omega), 0), \tag{13a}$$

$$0 = \nabla_x f(G_c(\sigma(\theta_\omega)), \theta_\omega). \tag{13b}$$

for all limit points $\theta_\omega \in \Omega(\Theta_\circ)$. □

Proof. The coupled dynamics (6) have the form:

$$\begin{aligned} z_{k+1} &= (A_c + B_c R M) z_k + B_c Q \theta_k + \chi(z_k, \theta_k), \\ \theta_{k+1} &= S \theta_k + \psi(\theta_k), \end{aligned} \tag{14}$$

for some mappings $\chi(z, \theta)$ and $\psi(\theta)$ that vanish at the fixed point along with their first-order derivatives, where the following matrices are Jacobians evaluated at the fixed point:

$$\begin{aligned} A_c &= \left[\frac{\partial F_c}{\partial z}\right]_{(z,y)=(z_\circ^\star,0)}, \quad & B_c &= \left[\frac{\partial F_c}{\partial y}\right]_{(z,y)=(z_\circ^\star,0)}, \\ R &= \left[\frac{\partial \nabla_x f}{\partial x}\right]_{(x,\theta)=(x_\circ^\star,0)}, \quad & M &= \left[\frac{\partial G_c}{\partial z}\right]_{z=z_\circ^\star}, \\ Q &= \left[\frac{\partial \nabla_x f}{\partial \theta}\right]_{(x,\theta)=(x_\circ^\star,0)}, \quad & S &= \left[\frac{\partial s}{\partial \theta}\right]_{\theta=0}. \end{aligned} \tag{15}$$

By assumption, the eigenvalues of the matrix $A_c + B_c R M$ are located on the open unit disc. Then by [17, Theorem 6], (14) has a center manifold at $(z_\circ^\star, 0)$, which is the graph of a mapping $z = \sigma(\theta)$ with $\sigma(\theta)$ satisfying (see [17, Eq. (2.8.4)]

$$\sigma(s(\theta)) = F_c(\sigma(\theta), \nabla_x f(G_c(\sigma(\theta)), \theta)). \tag{16}$$

Similar to the parameter-feedback case (Thm. 1), we show that, for each initialization $\theta_0 \in \Theta_0$, the gradient-feedback algorithm asymptotically tracks a critical trajectory of (1) with respect to θ_0 if and only the mapping σ satisfies (13) for all limit points $\theta_\omega \in \Omega(\theta_0)$, from which the result follows.

Lyapunov stability of the exosystem (Assumption 3) implies that the exosystem trajectory θ_k is bounded and therefore has a convergent subsequence. Then by definition, there exists an increasing sequence k_i such that the trajectory θ_{k_i} converges to some limit point $\theta_\omega \in \Omega(\theta_0)$ as $i \to \infty$. Furthermore, the closed-loop system is locally exponentially stable (by assumption), so $z_k \to z_\circ^*$ as $k \to \infty$. By continuity of G_c, we then also have that $x_k \to x_\circ^*$ as $k \to \infty$. Then using continuity of the gradient (Assumption 1), we have that

$$\lim_{i \to \infty} y_{k_i} = \lim_{i \to \infty} \nabla_x f(x_{k_i}, \theta_{k_i}) = \nabla_x f(x_\circ^*, \theta_\omega). \tag{17}$$

When $y_k \to 0$ as $k \to \infty$, the left-hand side is zero, which implies that (13b) holds on θ_ω. Since this holds for any limit point, (13b) holds on $\Omega(\theta_0)$. Equation (13a) then follows from the center manifold (16).

Now suppose the conditions (13) hold for all $\theta_\omega \in \Omega(\theta_0)$. The right-hand side of (17) is then zero, which implies the existence of a sequence k_i such that $y_{k_i} \to 0$ as $i \to \infty$. Since this holds for any limit point, any convergent subsequence of y_k also converges to zero. Moreover, y_k is bounded due to Lipschitz continuity of the gradient, so y_k also converges to zero as $k \to \infty$. □

The two conditions in (13) fully characterize the class of optimization algorithms that achieve asymptotic tracking of a critical trajectory. In words, a gradient-feedback optimization algorithm (5) tracks a critical trajectory if and only if, for some mapping σ, the composite function $G_c \circ \sigma$ zeros the gradient on the limit set of the exosystem (see (13b)), and the controller $F_c(z, y)$ is algebraically related to the exosystem $s(\theta)$ as given by (13a). Notice that, by Theorem 1, the former condition implies that

$$x_k = G_c(\sigma(\theta_k)) \tag{18}$$

is a parameter-feedback optimization algorithm for (1). We remark that (13) constitute the discrete-time counterpart of the continuous-time setting [15, Thm. 4].

Remark 3 (The internal model principle)**.** We interpret (13a) as the *(discrete-time) internal model principle of time-varying optimization,* as it expresses the requirement that any optimization algorithm that achieves asymptotic tracking must include an internal model of the exosystem. □

It is important to note that, by Theorem 2, the exosystem state θ and that of the optimization z must be related, in the limit set of the exosystem, by the relationship:

$$z_k = \sigma(\theta_k). \tag{19}$$

Intuitively, (19) is interpreted as the existence of a change of coordinates between the state of the exosystem and that of the optimization algorithm; see [15] for a discussion.

Remark 4 (Special cases)**.** An important special case is obtained when both σ and G_c are the identity operators; in this case, the internal model condition (13a) simplifies to

$$s(\theta) = F_c(\theta, 0),$$

which states that the controller vector field $F_c(z, y)$ must coincide with that of the exosystem $s(\theta)$ on the limit set of the exosystem. In this case, (19) gives $z_k = \theta_k$, meaning that the controller state z_k and that of the exosystem θ_k must coincide on the limit set. □

While Theorem 2 provides a full characterization of all gradient-feedback algorithms that achieve tracking, it remains to address under what conditions on the loss function $f(x, \theta)$ such an algorithm is guaranteed to exist. To address this question, we first require the following definition.

Definition 4 (Exponential detectability [23])**.** A Lyapunov stable dynamical system

$$\theta_{k+1} = f(\theta_k, x_k), \qquad\qquad y_k = h(\theta_k, x_k),$$

where f and h are smooth mappings such that $f(0, 0) = 0$ and $h(0, 0) = 0$, is *exponentially detectable from* y if there exists a dynamical system

$$\hat{\theta}_{k+1} = g(\hat{\theta}_k, y_k, x_k), \tag{20}$$

where g is a smooth mapping with $g(0, 0, 0) = 0$ such that: (i) if $\theta_0 = \hat{\theta}_0$, then $\theta_k = \hat{\theta}_k$ for all $k \in \mathbb{N}_{\geq 0}$, and (ii) there exists an open neighborhood Θ_1 of the origin such that $\|\hat{\theta}_k - \theta_k\| \leq M a^k \|\hat{\theta}_0 - \theta_0\|$ for all $k \in \mathbb{N}_{\geq 0}$, $\|\hat{\theta}_0 - \theta_0\| \in \Theta_1$, and for some positive constants M and $0 < a < 1$. In this case, (20) is called a local exponential observer [23]. □

Although we impose no restriction on how θ influences the loss $f(x, \theta)$, we require the following.

Assumption 4 (Detectability of the exosystem)**.** The exosystem (2) is exponentially detectable from the gradient-feedback signal (5b). □

Intuitively, since θ_k is unmeasurable, the temporal variability of the cost can only be evaluated through measurements of y_k. Contrarily, undetectability of the exosystem corresponds to a redundant description of the exogenous signal: if some modes of the exosystem did not influence the gradient, then y_k would be independent of those modes and they can thus be removed without altering the problem.

We are now equipped to address the conditions on f under which a gradient-feedback algorithm is guaranteed to exist.

Theorem 3 (Existence of gradient-feedback algorithms)**.** Suppose Assumptions 1, 2, 3, and 4 hold. There exists a gradient-feedback optimization algorithm that solves Problem 1 if and only if there exists a mapping $H_c : \Theta \to \mathbb{R}^n$ that zeros the gradient on the limit set of (2) with respect to its initialization. □

Proof. (Only if) By Theorem 2, there exists a mapping $z = \sigma(\theta)$ such that (13b) holds. Then, the gradient condition (11) holds with $H_c(\theta) = G_c(\sigma(\theta))$.

(If) We will prove this claim by constructing a gradient-feedback algorithm that achieves $y_k \to 0$ as $k \to \infty$.

By Assumption 4, there exists a neighborhood N of the origin and a dynamical system

$$\hat{\theta}_{k+1} = g(\hat{\theta}_k, y_k), \tag{21}$$

such that $\hat{\theta}_k \to \theta_k$ exponentially, for any $\|\hat{\theta}_0 - \theta_0\| \in N$. Consider then the optimization algorithm

$$F_c(z, y) = g(z, y), \qquad\qquad G_c(z) = H_c(z), \tag{22}$$

where $H_c(z)$ is as in (11). The claim then follows by application of Theorem 2 with σ the identity operator. □

Interestingly, the theorem shows that existence of a dynamic gradient-feedback algorithm is equivalent to that of a static parameter-feedback one. Although this may seem surprising (since a parameter-feedback algorithm has access to more information: θ_k can be measured, while this is not the case in the gradient-feedback case), it is worth noting that in the gradient-feedback construction, the dynamic state of the controller z_k acts as an alternative representation (i.e., in different coordinates) of the exosystem state θ_k (see (19)) and the exploration signal x_k acts as a parameter-feedback algorithm (see (18)). With this interpretation, it is evident that the two problems shall be equivalent, provided that θ_k can be estimated from the available measurements y_k (as done in the construction (22)).

Remark 5 (Algorithm design)**.** Once the mapping $H_c(\theta)$ as in (11) and a local exponential observer $g(\theta, y)$ as in (21) are available, algorithm (22) solves the gradient-feedback problem. □

We illustrate the applicability of the results in this section in the following example.

Example 2. Consider the quadratic problem studied in Example 1, and assume that the exosystem follows the linear model $\dot{\theta} = S\theta$ for some $S \in \mathbb{R}^{p\times p}$. According to Theorem 3, an optimization algorithm given by

$$z_{k+1} = A_c z_k + B_c y_k, \quad x_k = G_c z_k, \quad y_k = R x_k + Q\theta, \tag{23}$$

where $A \in \mathbb{R}^{n_c\times n_c}$, $B_c \in \mathbb{R}^{n_c\times n}$, $G_c \in \mathbb{R}^{n\times n_c}$, achieves asymptotic tracking if and only if there exists a linear transformation $\Sigma \in \mathbb{R}^{n_c\times p}$ such that:

$$0 = (\Sigma S - A_c\Sigma)\theta, \tag{24a}$$

$$0 = (RG_c\Sigma + Q)\theta, \tag{24b}$$

for all θ in the limit set of the exosystem. □

B. Design of gradient-feedback optimization algorithms

Theorem 3 provides an explicit technique to synthesize gradient-feedback algorithms (see Remark 5). Yet, its application remains challenging as one needs to design an exponential observer for the state of the exosystem. We next show that this process can be accomplished by having access only to first-order information on the exosystem. We begin by presenting an instrumental lemma; its statement hinges on the following notation (which we recall from (15)):

$$Q = \left[\frac{\partial \nabla_x f}{\partial \theta}\right]_{(x,\theta)=(x_\circ^\star,0)}, \qquad S = \left[\frac{\partial s}{\partial \theta}\right]_{\theta=0}.$$

Lemma 4 (First-order detectability of exosystem)**.** There is an exponential observer for (4) if and only if the pair (Q, S) is detectable. □

Proof. The claim follows directly from [23, Cor. 3.4]. □

Harnessing this tool, a technique to design an exponential observer of the exosystem is presented in Algorithm 1. Here, a linear Luenberger observer is used to estimate the exosystem state (see line 4), and a parameter feedback algorithm is then applied to the estimated state to regulate the gradient to zero (precisely, $G_c(z)$ is designed following the approach of Theorem 1 (see line 3).

Remark 6 (Alternative observer structures)**.** Instead of a Luenberger observer, alternative dynamic observers could be considered in Line 4 of the algorithm to achieve different asymptotic or transient properties of the resulting gradient-feedback optimization algorithm. □

Algorithm 1: Gradient-feedback algorithm design

Data: $s(\theta)$, $\nabla_x f(x,\theta)$, $H_c(\theta)$ satisfying (13), Jacobian matrices Q and S in (15)

1 $n_c \leftarrow n$;
2 $L \leftarrow$ any matrix such that $S - LQ$ is Schur stable;
3 $G_c(z) \leftarrow H_c(z)$;
4 $F_c(z,y) \leftarrow s(z) + L(y - \nabla_x f(H_c(z), z))$;

Result: $F_c(z,y)$, $G_c(z)$, and n_c that solve Problem 1

We illustrate the applicability of (13) on a quadratic problem in the following example.

Example 3. Consider the quadratic problem from Example 2. A direct application of Algorithm 1 gives:

$$A_c = S, \qquad B_c = L, \qquad G_c = -R^\dagger Q,$$

where L is any matrix such that $S - LQ$ is Schur stable; notice that this choice satisfies (24) with $\Sigma = I$. □

Remark 7 (Tracking accuracy vs internal model fidelity)**.** In general, the tracking accuracy of the optimization algorithm will depend on the fidelity of the internal model as well as the asymptotic behavior of the exosystem. This property is discussed in detail in [15, Sec. 5] for quadratic problems. We stress that *this is not only a limitation of our approach, but of any algorithm seeking exact tracking of a critical trajectory.* In this sense, the internal model principle proved in Theorem 2 (see also Remark 3) provides a *fundamental limitation* that should be carefully considered when designing optimization algorithms for time-varying problems. □

V. Simulation results

In this section, we illustrate our approach through numerical simulations. We consider the following instance of (1):

$$\min_{x\in\mathbb{R}} f(x,\theta_k) := \frac{1}{2}(x - \theta_k^{(1)})^2 + \kappa \log(1 + e^{\mu x}), \tag{25}$$

where $f : \mathbb{R}\times\Theta \to \mathbb{R}, \Theta = \mathbb{R}^2$, and we utilized the vector notation $\theta_k = (\theta_k^{(1)}, \theta_k^{(2)})$ (the choice to use $\Theta = \mathbb{R}^2$ instead of $\Theta = \mathbb{R}^1$ will be discussed shortly below). (25) models a logistic regression problem with a time-varying regularization term. Intuitively, an optimizer of (25) is a point that tracks the time-varying signal $\theta_k^{(1)}$, while seeking to avoid large values of x, which are penalized by the logistic term. For our experiments, we choose $\mu = 0.5$ and $\kappa = 1$. The function $f(x,\theta)$ satisfies Assumption 1; in particular, the cost is strongly convex in x (since $\nabla_{xx} f(x,\theta) = 1 + \kappa\mu^2 \frac{\exp(\mu x)}{[1+\exp(\mu x)]^2} \geq 1$), and thus the optimizer is unique for each θ. We let $\theta_k^{(1)} = \cos(\omega k)$, which can be generated by a two-dimensional linear exosystem $\theta_{k+1} = S\theta_k$ (hence the choice $\Theta = \mathbb{R}^2$). For our simulations, we generate matrix S by discretizing a continuous-time linear system with state matrix $S_{\text{ct}} = [0, 1; \omega^2, 0]$ with $\omega = 0.2$, yielding $S = [0.9801, 0.9933; -0.0397, 0.9801]$.

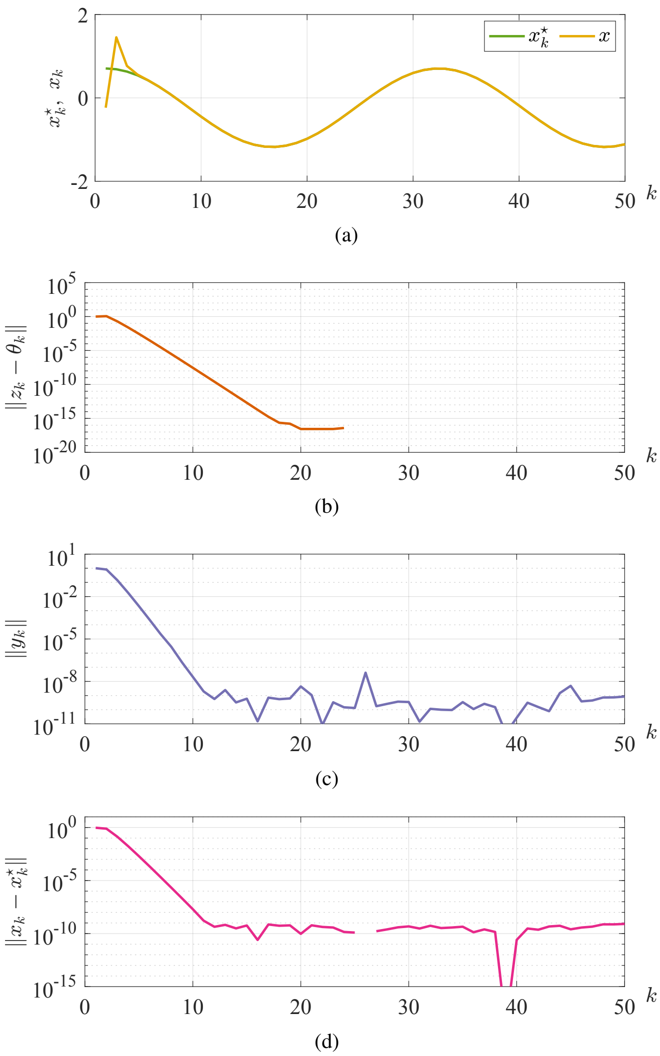

Fig. 1: Simulation results illustrating the behavior of an algorithm synthesized using Algorithm 1 for the problem (25). The proposed algorithm successfully computes the time-varying optimizer of (25) with exponential rate of convergence. In (b), absence of a line means that the value of the timeseries is numerically zero. See Section V for a discussion.

In Fig. 1, we plot four relevant time series illustrating the behavior of the optimization algorithm applied to this problem. We applied Algorithm 1 with $Q = [-1, 0]$, where we chose the observer gain L such that the spectral radius (the maximum eigenvalue modulus) of $S - LQ$ is 0.1. Moreover, a mapping zeroing the gradient has been computed numerically, yielding $H_c(\theta) = (0.9819 \cdot \theta^{(1)} - 0.2469, 0)$. From the numerical simulations, we can conclude the following: (i) from Fig. 1(b), we see that $z_k \to \theta_k$ exponentially, and thus z_k is a local exponential observer for θ_k; (ii) from Fig. 1 (c), we see that $\|y_k\| \to 0$ exponentially, and thus the algorithm converges to a critical point of (25); more precisely: (iii) in Fig. 1(a) and (d), we see that $\|x_k - x_k^\star\| \to 0$, and thus the algorithm converges to the time-varying optimizer of (25).

In Fig. 2, we plot the error $\|x_k - x_k^\star\|$ of the algorithm proposed here and compare it with the prediction-correction algorithm proposed in [10]. For this simulation, we used

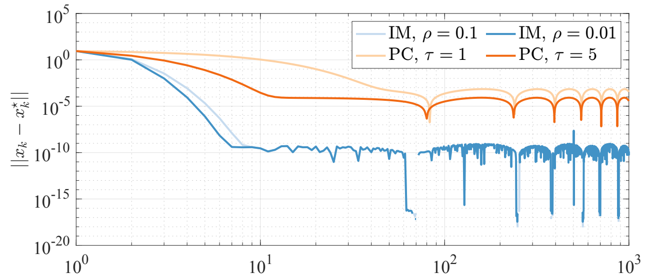

Fig. 2: Comparison between the approach proposed here (labed IM for Internal Model in the plot) and the Prediction Correction (labeled PM in the plot) algorithm [10] for the problem (25). In the plot, ρ denotes the spectral radius of the observer for θ, and τ the horizon of the prediction step [10]. Even by employing large prediction horizons, the approach proposed here outperformed [10] for this problem.

$\omega = 0.02$. The prediction-correction algorithm has been implemented following [10, Algorithm 1] with stepsize $\gamma = 0.2$. Numerically, we are led to conclude that, for this problem, our approach outperforms the prediction-correction algorithm both in convergence rate and in asymptotic precision. The difference in performance can be further appreciated by varying the spectral radius of $S - LQ$ for the exogenous signal observer in the set $\{0.1, 0.01\}$ and by varying the horizon of prediction $\tau \in \{1, 5\}$ in [10]. As expected, reducing the spectral radius of the observer and increasing the prediction horizon improve both the rate of convergence and the asymptotic precision of the two algorithms. In both cases, however, the prediction-correction method is outperformed by the approach in this work.

VI. Conclusions

The main result of this work is a fundamental result in time-varying optimization which states that any algorithm that asymptotically tracks a critical trajectory must embed an internal model of the time variation. We exploited this result to provide a design procedure to construct algorithms for time-varying optimization. The proposed approach relies on an exponential observer to estimate the temporal variability of the problem, combined with an algorithm design that zeros the gradient. Possible extensions include the use of other observers to influence the properties of the resulting algorithm, and application of the methodology to structured time-varying problems arising from particular applications.

References

[1] A. Simonetto, E. Dall'Anese, S. Paternain, G. Leus, and G. B. Giannakis, "Time-varying convex optimization: Time-structured algorithms and applications," *Proceedings of the IEEE*, vol. 108, no. 11, pp. 2032–2048, 2020.

[2] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *18th European Control Conference*, 2019. doi: 10.23919/ECC.2019.8796030. pp. 3420–3431.

[3] J. Rawlings, D. Mayne, and M. Diehl, *Model Predictive Control: Theory, Computation, and Design*. Nob Hill Publishing, 2017. ISBN 978-0-9759377-3-0. [Online]. Available: https://books.google.com/books?id=MrJctAEACAAJ

[4] M. Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent," in *International conference on machine learning*, 2003, pp. 928–936.

[5] E. Hazan, A. Agarwal, and S. Kale, "Logarithmic regret algorithms for online convex optimization," *Machine Learning*, vol. 69, no. 2, pp. 169–192, 2007.

[6] E. C. Hall and R. M. Willett, "Online convex optimization in dynamic environments," *IEEE Journal of Selected Topics in Signal Processing*, vol. 9, no. 4, pp. 647–662, 2015.

[7] S. M. Fosson, "Centralized and distributed online learning for sparse time-varying optimization," *IEEE Transactions on Automatic Control*, vol. 66, no. 6, pp. 2542–2557, 2020.

[8] S. Shalev-Shwartz, "Online learning and online convex optimization," *Foundations and Trends in Machine Learning*, vol. 4, no. 2, pp. 107–194, 2012.

[9] R. Dixit, A. S. Bedi, R. Tripathi, and K. Rajawat, "Online learning with inexact proximal online gradient descent algorithms," *IEEE Transactions on Signal Processing*, vol. 67, no. 5, pp. 1338–1352, 2019.

[10] A. Simonetto, A. Mokhtari, A. Koppel, G. Leus, and A. Ribeiro, "A class of prediction-correction methods for time-varying convex optimization," *IEEE Transactions on Signal Processing*, vol. 64, no. 17, pp. 4576–4591, 2016. doi: 10.1109/TSP.2016.2568161.

[11] T. Ni, X.-Z. Xie, and W.-Z. Gu, "Discrete-time Euler-smoothing methods for time-varying convex constrained optimization," *Journal of the Franklin Institute*, vol. 361, no. 10, p. 106898, 2024. doi: 10.1016/j.jfranklin.2024.106898.

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

[13] G. Bianchin, J. Cortés, J. I. Poveda, and E. Dall'Anese, "Time-varying optimization of LTI systems via projected primal-dual gradient flows," *IEEE Transactions on Control of Network Systems*, vol. 9, no. 1, pp. 474–486, Mar. 2022. doi: 10.1109/TCNS.2021.3112762.

[14] N. Bastianello, R. Carli, and S. Zampieri, "Internal model-based online optimization," *IEEE Transactions on Automatic Control*, 2023.

[15] G. Bianchin and B. V. Scoy, "The internal model principle of time-varying optimization," *arXiv preprint*, Aug. 2024, arXiv:2407.08037.

[16] B. Castillo, S. Di Gennaro, S. Monaco, and D. Normand-Cyrot, "Nonlinear regulation for a class of discrete-time systems," *Systems & Control Letters*, vol. 20, no. 1, pp. 57–65, 1993. doi: 10.1016/0167-6911(93)90087-M.

[17] J. Carr, *Applications of centre manifold theory*. Springer-Verlag, 1981, vol. 35. doi: 10.1007/978-1-4612-5929-9.

[18] S. Wiggins, *Ordinary Differential Equations: A Dynamical Point of View*. World Scientific, 2023.

[19] G. Osipenko, *Center Manifolds*. New York, NY: Springer New York, 2011, pp. 48–62. ISBN 978-1-4614-1806-1

[20] B. A. Francis and W. M. Wonham, "The internal model principle of control theory," *Automatica*, vol. 12, no. 5, pp. 457–465, 1976.

[21] E. Hazan, "Introduction to online convex optimization," *Foundations and Trends in Optimization*, vol. 2, no. 3-4, pp. 157–325, 2016.

[22] G. D. Birkhoff, *Dynamical systems*. American Mathematical Society: New York, 1927.

[23] W. Lin and C. I. Byrnes, "Design of discrete-time nonlinear control systems via smooth feedback," *IEEE Transactions on Automatic Control*, vol. 39, no. 11, pp. 2340–2346, 1994.