DEES: A Class of Data-Enabled Robust Feedback Algorithms for Real-Time Optimization

Poveda, J.; Vamvoudakis, K.; Benosman, M.

TR2019-095 September 05, 2019

Abstract

The increasing availability of information-rich data sets offers an invaluable opportunity to complement and improve the performance of existing model-based feedback algorithms. Following this principle, in this paper we present a novel class of Data-Enabled Extremum Seeking (DEES) algorithms for static maps, which make use of current and recorded data in order to solve a convex optimization problem characterized by a variational inequality. The optimization dynamics synergistically combine ideas from concurrent learning and classic neuroadaptive extremum seeking in order to dispense with the assumption of requiring a persistence of excitation condition in the closed-loop system. Using analytical tools for nonlinear systems we show that for a general class of optimization dynamics it is possible to tune the parameters of the controller to guarantee convergence in finite time to an arbitrarily small neighborhood of the set of optimizers. The results are illustrated in a scalar constrained minimization problem.

IFAC Nonlinear Control Systems (NOLCOS)

DEES: A Class of Data-Enabled Robust Feedback Algorithms for Real-Time Optimization

J. I. Poveda * K. Vamvoudakis ** M. Benosman ***

* *Department of Electrical, Computer, and Energy Engineering, University of Colorado, Boulder, CO 80309 USA (e-mail: jorge.poveda@colorado.edu).*
** *School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA, 30332-0150, USA (e-mail: kyriakos@gatech.edu).*
*** *Mitsubishi Electric Research Laboratories, Cambridge, MA, 02139, USA (e-mail: m.benosman@ieee.org).*

Abstract: The increasing availability of information-rich data sets offers an invaluable opportunity to complement and improve the performance of existing model-based feedback algorithms. Following this principle, in this paper we present a novel class of Data-Enabled Extremum Seeking (DEES) algorithms for static maps, which make use of *current* and *recorded* data in order to solve a convex optimization problem characterized by a variational inequality. The optimization dynamics synergistically combine ideas from concurrent learning and classic neuro-adaptive extremum seeking in order to dispense with the assumption of requiring a persistence of excitation condition in the closed-loop system. Using analytical tools for nonlinear systems we show that for a general class of optimization dynamics it is possible to tune the parameters of the controller to guarantee convergence in finite time to an arbitrarily small neighborhood of the set of optimizers. The results are illustrated in a scalar constrained minimization problem.

Keywords: Extremum seeking, adaptive control, optimization.

1. INTRODUCTION

Recent years have seen an increase of interest in data-driven control techniques that exploit the large amounts of data available in many practical applications. However, while areas such as machine learning, reinforcement learning, and gradient-free optimization have made significant breakthroughs during the last years, there is still a need to develop data-driven algorithms with provable and certifiable robustness properties for certain applications where safety is a critical concern. This necessity has motivated the study of novel robust feedback mechanisms that combine data-driven techniques with model-based principles, i.e., learning-based controllers, see for instance Benosman (2016), Vamvoudakis et al. (2015) and the recent survey papers of Benosman (2018) and Poveda et al. (2019).

In the area of real-time optimization, extremum seeking (ES) control has emerged as a promising feedback-based methodology with certifiable stability and robustness properties. The first stability analysis of ES was presented in Krstić and Wang (2000) and Ariyur and Krstić (2003) using averaging and singular perturbation theory for ordinary differential equations. Semi-global practical results were later developed in Tan et al. (2006), and further generalized in Nešić et al. (2010); Poveda and Teel (2017a) for a broader class of optimization dynamics. Other averaging-based algorithms have been recently studied in Dürr et al. (2013); Mills and Krstić (2014); Poveda et al. (2017a, 2018); Grushkovskaya et al. (2018); Feiling et al. (2018). ES architectures based on sampled-data systems have been studied in Popovic et al. (2006); Khong et al. (2013) and in Poveda and Teel (2017b). In the context of ES with a persistence of excitation (PE) condition, the works of Guay and Zhang (2003); Guay et al. (2015); Dougherty and Guay (2017) and references therein, have developed significant results for single-agent and multi-agent optimization problems. A set-point-based relaxed PE condition with sinusoids is presented in Adetola and Guay (2007). Hybrid ES architectures that combine continuous-time and discrete-time dynamics, and that require a PE condition during the flows were also presented in Poveda et al. (2017b). These techniques have all been developed under the assumption that the controller has access only to current data. However, as shown in Chowdhary and Johnson (2010), in some adaptive controllers it is possible to relax the classic PE condition by using "sufficiently rich" recorded data. Nevertheless, to the best of our knowledge, the development and analysis of ES architectures that use current and past data concurrently during the seeking process are absent in the literature.

In this paper, we present a novel class of ES algorithms that dispense with the classic PE condition. In particular, we show that the ideas behind concurrent learning and extremum seeking control can be synergistically combined to achieve real-time optimization in settings where there exists sufficiently rich stored data. The resulting Data-

⋆ Research supported in part by the NSF CAREER CPS-1851588.

Enabled Extremum Seeking (DEES) algorithms are applicable to general optimization problems described as convex variational inequalities where the cost function is only accessible via measurements or function evaluations. One of the requirements that we impose on these dynamics is that the optimizing state needs to evolve in a compact set defined a priori for all time. Given that there exists several Lipschitz continuous gradient-based optimization dynamics with projection that guarantee forward invariance of compact sets, the compactness requirement does not seem to impose major restrictions on the type of optimization problems that we can study. To illustrate this point, we present a Lipschitz DEES algorithm based on projected gradient dynamics. To the best of our knowledge, the results of this paper correspond to the first ES algorithms that exploit the idea of concurrent learning using recorded data instead of dithering signals.

The rest of this paper is organized as follows: In Section 2 we present some preliminaries and definitions. In Section 3 we present the main results. Section 4 presents a numerical example, and finally Section 5 ends with the conclusions.

2. PRELIMINARIES AND NOTATION

The set of (nonnegative) real numbers is denoted by ($\mathbb{R}_{\geq 0}$) $\mathbb{R}$. The set of (nonnegative) integers is denoted by ($\mathbb{Z}_{\geq 0}$) $\mathbb{Z}$. We use $\mathbb{B}$ to denote a closed unit ball of appropriate dimension, $\rho\mathbb{B}$ to denote a closed ball of radius $\rho > 0$, and $\mathcal{X} + \rho\mathbb{B}$ to denote the union of all sets obtained by taking a closed ball of radius ρ around each point in the set $\mathcal{X}$. We use $\overline{\text{co}}\ \mathcal{X}$ to denote the closed convex hull of $\mathcal{X}$, $\overline{\mathcal{X}}$ to denote its closure, and $\text{int}(\mathcal{X})$ to denote its interior. We say that $f \in \mathcal{C}^2$ if the function f is at least twice differentiable. A constrained ODE of the form

$$\dot{x} = F(x), \quad x \in C, \tag{1}$$

is said to render a compact set $\mathcal{A}$ uniformly globally asymptotically stable (UGAS) if there exists a $\mathcal{KL}$ function β such that $|x(t)|_{\mathcal{A}} \leq \beta(|x(0)|, t)$, for all $t \in \text{dom}(x)$ and all $x(0) \in C$. If $\text{dom}(x) = [0, \infty)$, the solution x is said to be complete.

3. DATA-ENABLED EXTREMUM SEEKING

The standard ES setting in static maps aims to solve an optimization problem of the form

$$\text{minimize} \quad f(x) \quad \text{subject} \quad x \in K, \tag{2}$$

where $f : \mathbb{R}^n \to \mathbb{R}$ is an unknown continuously differentiable objective function defined on an open set containing the closed set $K \subset \mathbb{R}^n$. The objective function is assumed to be accesible only by measurements or evaluations. If f is also a convex function and the set K is convex, a point $x^* \in K$ is a solution of (2) if and only if it satisfies the variational inequality

$$(x - x^*)^\top \nabla f(x^*) \geq 0, \quad \forall\ x \in K, \tag{3}$$

where $\nabla f(x)$ is the gradient of $f(x)$, (Rockafellar and Wets, 1998, Thm 6.12). Consequently, for convex optimization problems, ES can be seen as a feedback-based data-driven algorithm designed to converge to the solutions of (3). Since, in general, it is difficult to verify a priori a convexity condition on the cost function $f(x)$, typical ES controllers assume that (3) holds only "locally". Because of this, and in order to impose some regularity conditions to problem (2), we make the following standing assumption.

Assumption 1. The function f is convex and continuously differentiable on an open set $D \supset K$, and the set K is compact, convex, and nonempty.

We denote by $\mathcal{A}$ the set of all feasible points x^* satisfying (3), that is

$$\mathcal{A} := \{x^* \in K : \ (x - x^*)^\top \nabla f(x^*) \geq 0, \ \ \forall\ x \in K\}. \tag{4}$$

By Assumption 1, this set is nonempty and compact (Facchinei and Pang, 2003, Corolloray 2.2.5).

There exist several approaches in the literature to solve the ES problem (2), including averaging-based architectures, sampled-data based controllers, stochastic algorithms, and neuro-adaptive schemes, to just name a few. In this paper, we focus on the later type of ES controllers, which exploit ideas from adaptive control and neural networks in order to obtain uniform approximations of $\nabla f(x)$ on compact sets.

3.1 Feedback Structure

In order to solve problem (2), consider an ES algorithm characterized by the following ODE:

$$\dot{\hat{w}} = F_w(\hat{w}, \phi(x), e), \tag{5a}$$

$$\dot{z} = \varepsilon F_z(\hat{w}^\top \nabla\phi(x), z), \quad z \in C_z, \tag{5b}$$

where $z := [x^\top, s^\top]^\top \in \mathbb{R}^{n+r}$ and $s \in \mathbb{R}^r$ is an auxiliary state of dimension $r \in \mathbb{Z}_{\geq 0}$. The mapping F_z is assumed to be continuous, and the set $C_z \subset \mathbb{R}^{n+r}$ is assumed to be compact. The pair (F_z, C_z) is application-dependent and designed to solve problem (3) based on the structure of K. The error signal e in (5a) is defined as

$$e = \hat{w}^\top \phi(x) - f(x), \tag{6}$$

and the state $\hat{w} \in \mathbb{R}^p$ is an auxiliary state used to estimate the ideal weights $w^* \in \mathbb{R}^N$ of a parameterization of the objective function, given by

$$f(x) = w^{\star\top} \phi(x) + \epsilon(x), \quad \forall\ x \in K, \tag{7}$$

where $w^\star \in \mathbb{R}^p$ belongs to the set of ideal weights

$$\mathbb{W}_{\phi,K} := \Big\{ w^* \in \mathbb{R}^p : |f(x) - w^{\star\top}\phi(x)| \leq \epsilon(x), x \in K \Big\},$$

which we assume to be compact, i.e., there exists $\bar{w} > 0$ such that $\mathbb{W}_{\phi,K} \subset \bar{w}\mathbb{B}$. The vector valued regressor function $\phi : K \to \mathbb{R}^p$ is assumed to be known and $\phi \in \mathcal{C}^2$. The regressor function $\phi := [\phi_1, \phi_2, \ldots, \phi_p]^\top$ should be selected such that the functions ϕ_j, $j \in \{1, \ldots, p\}$, define a complete independent basis set for $f(x)$. Typical choices of ϕ_j include quadratic functions, radial basis functions, or sigmoid functions, see Vamvoudakis and Lewis (2010).

We shall need the following technical assumption on the approximation (7).

Assumption 2. The approximation error function $\epsilon(\cdot)$ in (7) is continuously differentiable.

By computing the gradient of $f(x)$ in (7) we can obtain the following parameterized expression for ∇f:

$$\nabla f(x) = \nabla\phi(x)^\top w^\star + \nabla\epsilon(x), \quad \forall\ x \in K. \tag{8}$$

where $\nabla\phi(x)$ is the Jacobian matrix of $\phi(x)$. Since K is compact and the mappings $\nabla f(x)$ and $\nabla\varepsilon(x)$ are continuous, by the Weierstrass high-order approximation theorem (Dudley, 2002, Thm. 2.4.11), the approximation error converges to zero as the number of basis p increases, i.e., $\epsilon(x) \to 0$ and $\nabla\epsilon(x) \to 0$ as $p \to \infty$, uniformly on K. Moreover, due to Assumption 2, the compactness of K and the fact that $\phi(x)$ is C^2, we have that $\phi(x)$, $\nabla\phi(x)$, $\epsilon(x)$, and $\nabla\epsilon(x)$ are all uniformly bounded in K.

3.2 Online Learning with Persistent of Excitation

Let $\hat{f}(x)$ be an approximation of the objective function (7), defined as

$$\hat{f}(x) := \hat{w}^\top \phi(x), \quad \forall\, x \in K. \tag{9}$$

Let $\tilde{w} = \hat{w} - w^\star$. Using (6), (7) and (9) we have that

$$\begin{aligned} e(t) &= \hat{f}(x(t)) - f(x(t)) \\ &= \tilde{w}^\top \phi(x(t)) - \epsilon(x(t)). \end{aligned}$$

Consider the mapping (5a), defined as

$$F_w(\hat{w}, \phi(x), e) := -\alpha \frac{\phi(x(t))}{(1 + \phi(x(t))^\top \phi(x(t)))^2} e(t), \tag{10}$$

where $\alpha > 0$ is a tunable constant. If $\epsilon(x) = 0$, it is well-known, e.g., Ioannou and Sun (2012), Vamvoudakis and Lewis (2010), Narendra and Annaswamy (1987), that all solutions of (5a) with vector field (10) will satisfy $\hat{w}(t) \to w^\star$ as $t \to \infty$, uniformly in $(t_0, \hat{w}(0))$, if and only if the normalized regressor function $\tilde{\phi}(x(t))$, defined as

$$\tilde{\phi}(x(t)) := \frac{\phi(x(t))}{1 + \phi(x(t))^\top \phi(x(t))}, \tag{11}$$

is persistently exciting, i.e., there exists a $T > 0$ and $\gamma > 0$ such that

$$\int_t^{t+T} \tilde{\phi}(\tau)\tilde{\phi}^\top(\tau) \mathrm{d}\tau \geq \gamma I, \tag{12}$$

for all $t \geq t_0$ and all $t_0 \geq 0$. When $\epsilon(x) \neq 0$ in (10), the PE condition (12) guarantees convergence of $\hat{w}(t)$ to a neighborhood of $w^\star$ that is proportional to the upper bound of $\epsilon(x)$, see for instance Vamvoudakis and Lewis (2010).

Since the PE condition (12) needs to be satisfied for all present and future time, the PE condition can be restrictive and unfeasible for certain applications. On the other hand, as shown in Chowdhary and Johnson (2010), if the learning dynamics use current data $\phi(x(t))$ that is not PE, but which is complemented with a finite sequence of sufficiently rich stored past data $\{\phi(x(t_k))\}_{k=1}^{\bar{k}}$, it is still possible to achieve learning of the optimal weights w^*.

3.3 Online Learning with Recorded Data

To establish the main ideas behind the DEES algorithms, let $k \in \{1, 2, \ldots, \bar{k}\}$ denote the index of a stored data point x_k, i.e., $x_k = x(t_k)$, and let $\phi(x_k)$ be the regressor vector evaluated at that point. We still denote by $e(t)$ in (10) the estimation error corresponding to the data collected at the current time t, but we now also introduce an estimation error associated with the data previously collected at time t_k, given by

$$e(t_k, t) = \hat{f}(x(t_k)) - f(x(t_k)) \tag{13a}$$

$$= \tilde{w}(t)^\top \phi(x_k) - \epsilon(x_k). \tag{13b}$$

for all $t_k \in \{1, 2, \ldots, \bar{k}\}$. Note that the estimation error $\tilde{w}$ still depends on the current time t.

Definition 3. The sequence of stored data $\{\phi(x_k)\}_{k=1}^{k=\bar{k}}$ is said to be $\bar{k}$-*sufficiently rich* if the following inequality is satisfied

$$\sum_{k=1}^{\bar{k}} \tilde{\phi}(x_k)\tilde{\phi}_i(x_k)^\top \succ 0, \tag{14}$$

with $\tilde{\phi}(x_k) := \frac{\phi(x(t_k))}{1+\phi(x(t_k))^\top \phi(x(t_k))}$. □

According to Definition 3, the sequence of stored data is $\bar{k}$-sufficiently rich if its elements form a basis for the parameterized uncertainty during the window of discrete time $\{1, 2, \ldots, \bar{k}\}$. Indeed, by defining the regressor matrix

$$\phi^{mem} := [\phi(x_1), \phi(x_2), \ldots, \phi(x_k)]^\top, \tag{15}$$

the condition $\text{rank}(\phi^{mem}) = p$ is sufficient to satisfy (14).

Remark 4. Unlike the PE condition (12), which applies to the past and future behavior of $\phi(t)$, the condition (14) only needs to be verified for *past* data. This data can be obtained by performing repetitive experimental tests, or by exciting the system during an initial *finite* amount of time. This approach exploits information-rich data sets that are available in several applications, e.g., transportation systems, robotics, energy systems, etc. □

Using data that satisfies condition (14), we can consider learning dynamics (5a) that dispense with the PE condition. To streamline the presentation of the algorithm, and with some abuse of notation, we will use $t_0 = t$ to denote the current time, and we define

$$e(t_0, t) := e(t). \tag{16}$$

Using this notation, we replace the mapping F_w in (10) by the data-enabled mapping

$$F_w(\hat{w}, \phi(x), e) = -\alpha \sum_{k=0}^{\bar{k}} \frac{\phi(x(t_k))}{(\phi(x(t_k))^\top \phi(x(t_k)) + 1)^2} e(t_k, t). \tag{17}$$

Lemma 5. Suppose that Assumptions 1-2 hold and that the signals $\phi(x(t_k))$ in (17) are $\bar{k}$-sufficiently rich. Then, for each pair $(\bar{\nu}, c) \in \mathbb{R}^2_{>0}$ such that $\bar{\nu} < \sqrt{2c}$ there exists a sufficiently large $p^* \in \mathbb{Z}_{>0}$ such that for each $\lceil p \rceil > p^*$ there exists a UGAS compact set $\mathcal{A}_{p,z} \subset \mathbb{W}_{\phi,K} + \bar{\nu}\mathbb{B} \times C_z$ for the dynamics (5) with $\varepsilon = 0$ and (5a) restricted to the compact set $\mathbb{W}_{\phi,K} + \sqrt{2c}\mathbb{B}$. □

Proof: The proof of the following Lemma follows similar ideas as the proofs in Chowdhary and Johnson (2010) and Vamvoudakis and Lewis (2010). We divide the proof in three main steps.

Step 1: Let $\tilde{w} := \hat{w} - w^\star$, and consider the error dynamics

$$\begin{aligned} \dot{\tilde{w}} = &-\alpha \sum_{k=0}^{\bar{k}} \tilde{\phi}(x(t_k))\tilde{\phi}(x(t_k))^\top \tilde{w} \\ &+ \alpha \sum_{k=0}^{\bar{k}} \frac{\tilde{\phi}(x(t_k))}{\phi(x(t_k))^\top \phi(x(t_k)) + 1} \epsilon(x_k), \end{aligned} \tag{18}$$

(19)

which can be written as

$$\dot{\tilde{w}} = -\alpha\tilde{\phi}(x(t))\tilde{\phi}(x(t))^\top\tilde{w} - \alpha\sum_{k=1}^{\bar{k}}\tilde{\phi}(x(t_k))\tilde{\phi}(x(t_k))^\top\tilde{w}$$
$$+ \alpha\sum_{k=0}^{\bar{k}}\frac{\tilde{\phi}(x(t_k))}{\phi(x(t_k))^\top\phi(x(t_k))+1}\epsilon(x_k).$$

Define $P(t)$ as

$$P(t) := \tilde{\phi}(x(t))\tilde{\phi}(x(t))^\top + \sum_{k=1}^{\bar{k}}\tilde{\phi}(x(t_k))\tilde{\phi}(x(t_k))^\top. \quad (20)$$

Using that $\{\tilde{\phi}(x(t_k))\}_{k=1}^{\bar{k}}$ is $\bar{k}$-sufficiently rich, that $x(t_k)$ is constrained to a compact set, and that $\phi(\cdot)$ is continuous, there exists $\delta_1, \delta_2 > 0$ such that

$$\delta_2 I_p \succ \sum_{k=1}^{\bar{k}}\tilde{\phi}(x(t_k))\tilde{\phi}(x(t_k))^\top \succ \delta_1 I_p, \quad (21)$$

and since $\tilde{\phi}(x(t))\tilde{\phi}(x(t))^\top$ in (20) is symmetric, positive semidefinite, and uniformly bounded, there exists $\delta_3 > 0$ such that

$$\delta_3 I_p \succ P(t) \succ \delta_1 I_p, \quad (22)$$

for all $t \geq t_0$ and all $t_0 \geq 0$.

Step 2: Let $\rho(x)$ be given by

$$\rho(x) := \alpha\sum_{k=0}^{\bar{k}}\frac{\tilde{\phi}(x(t_k))}{[\phi(x(t_k))^\top\phi(x(t_k))+1]}\epsilon(x_k). \quad (23)$$

Consider the quadratic Lyapunov function

$$V(\tilde{w}) = 0.5\tilde{w}^\top\tilde{w}, \quad (24)$$

which satisfies

$$\dot{V} \leq -\delta_3|\tilde{w}|^2 + \tilde{w}^\top\rho(x), \quad (25)$$

where we used the upper bound in (22) and the definition of P in (20). If the approximation error $\epsilon(x)$ is zero in (7), we have that $\rho(x) = 0$ and $\tilde{w}(t)$ converges exponentially fast to zero. Moreover, the level sets $L_c := \{\tilde{w} \in \mathbb{R}^p : V(\tilde{w}) \leq c\}$ are positive invariant for each $c > 0$. On the other hand, when $\rho(x) \neq 0$, by the definition of the entries ρ in (23), the fact that x is constrained to a compact set, the continuity of $\phi(\cdot)$, and the approximation properties of the regressions in (7), for any $\nu > 0$ there exists a sufficiently large $p^* \in \mathbb{Z}_{>0}$ such that for all $\lceil p\rceil \geq p^*$ the residual term satisfies $|\rho(x)| < \nu$. Thus, for $\lceil p\rceil \geq p^*$ equation (25) satisfies

$$\begin{aligned}\dot{V} &\leq -\delta_a|\tilde{w}|^2 + \nu|\tilde{w}|,\\ &= -(1-\theta)\delta_a|\tilde{w}|^2 - \theta\delta_a|\tilde{w}|^2 + \nu|\tilde{w}|,\\ &\leq -(1-\theta)\delta_a|\tilde{w}|^2, \ \ \forall\ |\tilde{w}| \geq \bar{\nu},\ \theta \in (0,1), \end{aligned} \quad (26)$$

where $\bar{\nu} := (\theta\delta_a)^{-1}\nu$. Combining inequalities (24) and (26) we get ultimate boundedness of $\tilde{w}(t)$ with residual set proportional to ν.

Step 3: Finally, since Step 2 implies that for each $\tilde{w}(0) \in L_c$ and each $\nu > 0$ there exists a sufficiently large $p^* > 0$ such that for each $\lceil p\rceil \geq p^*$ there exists a $T > 0$ such that $[\tilde{w}(t), z(t)]^\top \in \bar{\nu}\mathbb{B} \times C_z$ for all $t \geq T$, by (Goebel et al., 2012, Corollary 7.7), there exists an asymptotically stable set $\mathcal{A}_{p,z} \subset \mathbb{W}_{\phi,K} + \bar{\nu}\mathbb{B} \times C_z$ for the dynamics (5) with $\varepsilon = 0$ and $\hat{w}$ restricted to a compact set. ■

After characterizing a data-enabled learning mechanism for the online estimation of the gradient ∇f, we can proceed to design the optimization dynamics to solve problem (3).

3.4 Robust Gradient-Based Optimization Dynamics

The optimization dynamics (5b) are designed under the assumption that $\nabla f(x) = \nabla\phi(x)^\top\hat{w}$. In particular, to solve the VI problem (3), we consider optimization dynamics with state $z = [x^\top, s^\top]^\top \in \mathbb{R}^{n+r}$ given by

$$\dot{z} = F_z(\nabla f(x), z), \quad z \in C_z, \quad (27)$$

where the function F_z and the set C_z are designed to satisfy the following Assumption.

Assumption 6. The dynamics (27) satisfy the following:

(a) The mapping F_z is continuous with respect to both arguments.
(b) The set C_z satisfies $C_z := K \times S$, where $S \subset \mathbb{R}^r$ is a compact set.
(c) There exists a nonempty compact set $\mathcal{S} \subset S$ such that the set $\mathcal{A} \times \mathcal{S}$ is UGAS.
(d) There exists an $\delta > 0$ such that for each measurable function $e : \mathbb{R}_{\geq 0} \to \mathbb{R}^n$ satisfying $\sup_{t\geq 0}|e(t)| \leq \delta$, the perturbed system

$$\dot{z} = F_z(\nabla f(z) + e, z), \quad z \in C_z, \quad (28)$$

generates complete solutions from each $z(0) \in C_z$. □

In words, Assumption 6 asks that system (27) solves problem (2) under the assumption of perfect knowledge of the gradient, and generates complete solutions under vanishing perturbations acting on the gradient. In some cases, item (d) can be relaxed and complete solutions are only required from compact subsets of C_z.

3.5 Main Result

Having characterized the data-enabled learning dynamics (17) and the optimization dynamics (27) that comprise the DEES algorithms (5), we are ready to present the main result of this paper.

Theorem 7. (Convergence of DEES Algorithms) Suppose that Assumptions 1, 2 and 6 hold, and that the sequence of data $\{\phi(x(t_k))\}_{k=1}^{\bar{k}}$ in (17) is $\bar{k}$-sufficiently rich. Then, for each pair $\Delta > \nu > 0$ there exists a $p^* \in \mathbb{Z}_{>1}$ such that for $\lceil p\rceil > p^*$ there exists $\varepsilon^* \in \mathbb{R}_{>0}$ such that for each $\varepsilon \in (0, \varepsilon^*)$ there exists a $T_{\nu,\varepsilon} \in \mathbb{R}_{>0}$ such that the x-component of every solution of the DEES dynamics (5) with $|\hat{w}(0)|_{\mathbb{W}_{\phi,K}} \leq \Delta$ satisfies

$$x(t) \in \mathcal{A} + \nu\mathbb{B}, \quad (29)$$

for all $t \geq T_{\nu,\varepsilon}$. □

Proof: The proof Theorem 7 makes use of tools from singular perturbation theory, robustness results for well-posed systems, and Ω-limit sets. In particular, the closed-loop system is given by

$$\dot{\hat{w}} = -\alpha\sum_{k=0}^{\bar{k}}\frac{\phi(x(t_k))}{[\phi(x(t_k))^\top\phi(x(t_k))+1]^2}\left(\tilde{w}^\top\phi(x_k) - \epsilon(x_k)\right) \quad (30a)$$

$$\dot{z} = \varepsilon F_z\left(\hat{w}^\top\nabla\phi(x_k), z\right), \quad z \in C_z. \quad (30b)$$

Let $\tau = \varepsilon t$ be a new time scale. When p is sufficiently large this system can be seen as a perturbed version of the nominal system in the τ-time scale

$$\varepsilon \frac{d\hat{w}}{d\tau} = -\alpha \sum_{k=0}^{\bar{k}} \frac{\phi(x(t_k))}{\left[\phi(x(t_k))^\top \phi(x(t_k)) + 1\right]^2} \left(\tilde{w}^\top \phi(x_k)\right) \tag{31a}$$

$$\frac{dz}{d\tau} = F_z\left(\hat{w}^\top \nabla\phi(x_k), z\right), \quad z \in C_z. \tag{31b}$$

When $\varepsilon > 0$ is sufficiently small and $\hat{w}$ is restricted to evolve in the compact set $\mathbb{W}_{\phi,K} + \rho\mathbb{B}$, system (31) is in singular-perturbation form, see Wang et al. (2012). The *boundary layer dynamics* are given by system (17), which, by the proof of Lemma 5, generates trajectories $\hat{w}$ that converge to $\mathbb{W}_{\phi,K}$ whenever p is sufficiently large. Thus, the *slow dynamics* correspond precisely to the gradient dynamics (27), which by item (c) in Assumption 6 renders UGAS the set $\tilde{\mathcal{A}} := \mathcal{A} \times \mathcal{S}$. Using singular perturbation results e.g., (Wang et al., 2012, Thm. 1), we obtain that for each $\nu/2 > 0$ there exists $\varepsilon^* > 0$ such that for each $\varepsilon \in (0, \varepsilon^*)$ there exists a $T_{\varepsilon,\nu} > 0$ such that if $|\hat{w}(0)| \leq \Delta$, the solutions of (31b) satisfy $z(\tau) \in \tilde{\mathcal{A}} + 0.5\nu\mathbb{B}$ for all $\tau \geq T$. The convergence result for the original system (30) follows now directly by taking ρ sufficiently large and by robustness results for ODEs of the form (1) with a continuous mapping F and a closed set C (Goebel et al., 2012, Thm. 7.21). ■

Remark 8. Unlike the classic ES architectures, the DEES algorithms do not require the injection of a dithering signal to the nominal state x. Instead, they rely on the dynamics (17) with rich recorded data, and a time-scale separation that can be achieved by selecting ε sufficiently small. □

4. APPLICATION: REAL-TIME OPTIMIZATION WITH LIPSCHITZ PROJECTION

We consider an optimization problem (2) where the cost function $f(x)$ is only accessible by measurements, and Assumption 1 holds. In particular, we consider the following DEES algorithm:

$$\dot{w} = -\alpha \sum_{k=0}^{\bar{k}} \frac{\phi(x(t_k))}{\left[\phi(x(t_k))^\top \phi(x(t_k)) + 1\right]^2} e(t_k, t) \tag{32a}$$

$$\dot{x} = \varepsilon\left[-x + \mathrm{P}_K(x - \hat{w}^\top \nabla\phi(x))\right], \tag{32b}$$

where

$$\mathrm{P}_K(x) = \mathrm{argmin}_{u\in K}\|x-u\|_2. \tag{33}$$

The projection dynamics (32b) are Lipschitz continuous since $P_K(x)$ satisfies the non-expansive property $|P_K(x) - P_K(y)| \leq |x - y|$. Also, dynamics (32b) render forward invariant the set K (Xia and Wang, 2000, Thm. 3.2). Moreover, when $\hat{w}^\top \nabla\phi(x) = \nabla f(x)$ small perturbations acting on the gradient do not affect the forward invariance of K since $P_K(x - \nabla f(x) + e) \in K$ for any e. Additionally, since by Assumption 1 the function f is convex, we have that $\nabla f(x)$ is a monotone gradient mapping (Rockafellar and Wets, 1998, Thm. 12.17). Thus, by (Gao, 2003, Thm. 3), the monotonicity and Lipschitz continuity of $\nabla f(x)$, plus the convexity and closedness of K, imply the convergence of the solutions of the projected gradient dynamics $\dot{x} = \varepsilon\left[-x + P_K(x - \nabla f(x))\right]$ to the set $\mathcal{A}$. Since K and $\mathcal{A}$ are compact, and the vector field in (32b) is Lipschitz

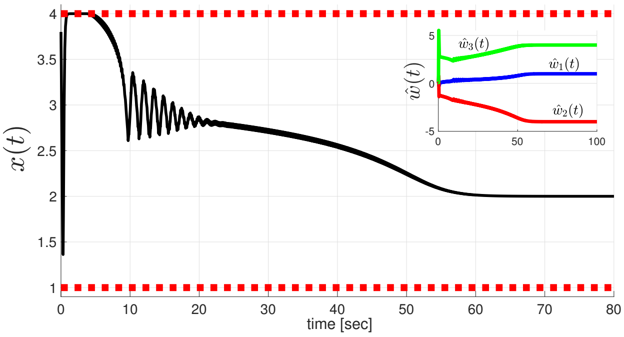

Fig. 1. Evolution in time of the state $x(t)$ with PE condition (12) on the regressor vector $\phi(t)$. The dotted lines describe the limits of the set K.

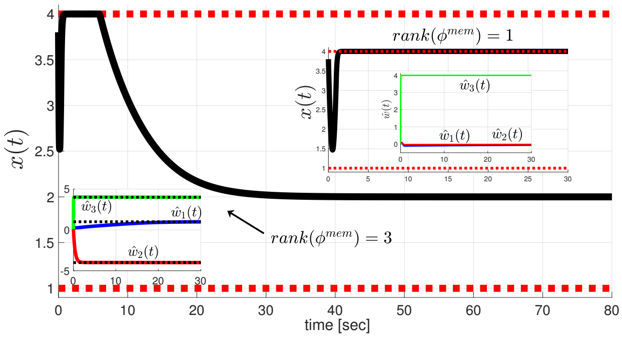

Fig. 2. Evolution in time of the state $x(t)$ with no PE condition and using memory-based ES with recorded sequence $\{\phi(x_k)\}$.

continuous, the convergence is uniform in K. Therefore, Assumption 6 is satisfied.

Figures 1 and 2 show simulations of the solutions of the dynamics (32) with and without memory for the case when

$$f(x) = (x-2)^2, \quad K = [1, 4]. \tag{34}$$

For the simulation shown in Figure 1 no recorded data was used and a dithering signal was injected into the closed-loop system to satisfy the standard PE condition (12). On the other hand, in Figure 2 we show two simulations where the DEES algorithm (32) was used. The algorithm uses a vector-valued regressor $\phi : \mathbb{R} \to \mathbb{R}^3$, given by $\phi(x) = [x^2, x, 1]^\top$. For the case when $rank(\phi^{mem}) = 3$ condition (14) is satisfied and the DEES converges to the optimizer of the cost function. On the other hand, when $rank(\phi^{mem}) < 3$, the inset shows that the algorithm does not converge to x^*. The recorded data $\{\phi(x_k)\}_k$ was generated by exciting the state during 5 seconds and sampling the state every 0.1 seconds. The red dotted lines describe the limits of the set K.

5. CONCLUSION

We presented a novel class of data-enabled ES algorithms that exploit information-rich data sets and avoid the injection of persistent dithering signals in the closed-loop system. The proposed scheme uses recorded data during the learning phase concurrently with current data in order to guarantee convergence to an ε-neighborhood of a convex optimization problem where the mathematical form of the cost function is unknown. Our results are general enough to be applied to different types of optimization dynamics that evolve on compact sets. We anticipate that our results can also be extended to settings where the cost function is generated by a stable dynamical system, to optimization problems with slowly varying cost functions, as well as to distributed optimization problems in multi-agent systems.

REFERENCES

Adetola, V. and Guay, M. (2007). Parameter convergence in adaptive extremum-seeking control. *Automatica*, 43, 105–110.

Ariyur, K.B. and Krstić, M. (2003). *Real-Time Optimization by Extremum-Seeking Control*. Wiley.

Benosman, M. (2016). *Learning-Based Adaptive Control: An Extremum Seeking Approach - Theory and Applications*. Butterworth-Heinemann, Cambridge, MA.

Benosman, M. (2018). Model-based vs data-driven adaptive control: An overview. *International Journal of Adaptive Control and Signal Processing*, 753–776.

Chowdhary, G. and Johnson, E. (2010). Concurrent learning for convergence in adaptive control without persistence of excitation. *49th IEEE Conference on Decision and Control*, 3674–3679.

Dougherty, S. and Guay, M. (2017). An extremum-seeking controller for distributed optimization over sensor networks. *IEEE Transactions on Automatic Control*, 62(2), 930–933.

Dudley, R.M. (2002). *Real Analysis and Probability*. Cambridge Uniersity Press.

Dürr, H., Stanković, M., Ebenbauer, C., and Johansson, K.H. (2013). Lie bracket approximation of extremum seeking systems. *Automatica*, 49, 1538–1552.

Facchinei, F. and Pang, J. (2003). *Finite-Dimensional Variational Inequalities and Complementary Problems*. Springer.

Feiling, J., Koga, S., Krstić, M., and Oliveira, T.R. (2018). Gradient extremum seeking for static maps with actuation dynamics governed by diffusion pdes. *Automatica*, 95, 197–206.

Gao, X. (2003). Exponential stability of globally projected dynamic systems. *IEEE Transactions on Neural Networks*, 14(2), 426–431.

Goebel, R., Sanfelice, R.G., and Teel, A.R. (2012). *Hybrid Dynamical Systems: Modeling, Stability, and Robustness*. Princeton University Press.

Grushkovskaya, V., Zuyev, A., and Ebenbauer, C. (2018). On a class of generating vector fields for the extremum seeking problem: Lie bracket approximation and stability properties. *Automatica*, 94, 151–160.

Guay, M., Vandemeulen, I., Dougherty, S., and Mclellan, P.J. (2015). Distributed extremum-seeking control over networks of dynamic agents. *In proc. of American Control Conferencel*, 159–164.

Guay, M. and Zhang, T. (2003). Adaptive extremum seeking control of nonlinear dynamic systems with parametric uncertainties. *Automatica*, 39, 1283–1293.

Ioannou, P.A. and Sun, J. (2012). *Robust Adaptive Control*. Dover Publications Inc., Mineola, NY.

Khong, S.Z., Nešić, D., Tan, Y., and Manzie, C. (2013). Unified framework for sampled-data extremum seeking control: Global optimisation and multi-unit systems. *Automatica*, 49, 2720–2733.

Krstić, M. and Wang, H.H. (2000). Stability of extremum seeking feedback for general nonlinear dynamic systems. *Automatica*, 36(4), 595–601.

Mills, G. and Krstić, M. (2014). Constrained extremum seeking in 1-dimension. *In Proc. of 53rd IEEE Conf. Decision Control*, 2654–2659.

Narendra, K.S. and Annaswamy, A. (1987). Persistent excitation in adaptive systems. *International Journal of Control*, 45(1), 127–160.

Nešić, D., Tan, Y., Moase, W., and Manzie, C. (2010). A unifying approach to extremum seeking: Adaptive schemes based on estimation of derivatives. *Proc. of 49th IEEE Conference on Decision and Control.*

Popovic, D., Jankovic, M., Magner, S., and Teel, A. (2006). Extremum seeking methods for optimization of variable cam timing engine operation. *IEEE Transactions on control systems technology*, 14, 398–407.

Poveda, J.I., Benosman, M., Sanfelice, R.G., and Teel, A.R. (2018). A hybrid adaptive feedback law for robust obstacle avoidance and coordination in multiple vehicle systems. *In proc. of American Control Conferece*, 616–621.

Poveda, J.I., Benosman, M., and Teel, A.R. (2017a). Distributed extremum seeking in multi-agent systems with arbitrary switching graphs. *IFAC World Congress*, 50(1), 735–740.

Poveda, J.I., Benosman, M., and Teel, A.R. (2019). Hybrid online learning control in networked multiagent systems: A survey. *International Journal of Adaptive Control and Signal Processing*, 33(2), 228–261.

Poveda, J.I. and Teel, A.R. (2017a). A framework for a class of hybrid extremum seeking controllers with dynamic inclusions. *Automatica*, 76, 113–126.

Poveda, J.I. and Teel, A.R. (2017b). A robust event-triggered approach for fast sampled-data extremization and learning. *IEEE Transactions on Automatic Control*, 62(10), 4949–4964.

Poveda, J.I., Vamvoudakis, K.G., and Benosman, M. (2017b). A neuro-adaptive architecture for extremum seeking control using hybrid learning dynamics. *In proc. of American Control Conferece. Boston, MA.*, 542–547.

Rockafellar, R.T. and Wets, R.J. (1998). *Variational Analysis*. Springer.

Tan, Y., Nešić, D., and Mareels, I. (2006). On non-local stability properties of extremum seeking controllers. *Automatica*, 42(6), 889–903.

Vamvoudakis, K., Antsaklis, P., Dixon, W., Hespanha, J., and Lewis, F. (2015). Autonomy and machine intelligence in complex systems: A tutorial. *In Proc. of American Contr. Conf.*, 5062–5079.

Vamvoudakis, K.G. and Lewis, F.L. (2010). Online actor-critic algorithm to solve the continuous-time infinite horizon optimal control problem. *Automatica*, 46(5), 878–888.

Wang, W., Teel, A.R., and Nešić, D. (2012). Analysis for a class of singularly perturbed hybrid system via averaging. *Automatica*, 48, 1057–1068.

Xia, Y.S. and Wang, J. (2000). On the stability of globally projected dynamical systems. *Journal of Optimization Theory and Applications*, 106(1), 129–150.