

Accelerated Concurrent Learning Algorithms via Data-Driven Hybrid Dynamics and Nonsmooth ODEs

Daniel E. Ochoa*

DANIEL.OCHOA@COLORADO.EDU

Jorge I. Poveda*

JORGE.POVEDA@COLORADO.EDU

Anantharaman Subbaraman[†]

ANANTHARAMAN.SUBBARAMAN@US.BOSCH.COM

Gerd S. Schmidt[†]

GERDSIMON.SCHMIDT@DE.BOSCH.COM

Farshad R. Pour Safaei[†]

FARSHAD.RAMEZAN@US.BOSCH.COM

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

[†] *Bosch Research and Technology Center, Sunnyvale, CA.*

Abstract

We introduce a novel class of accelerated data-driven concurrent learning algorithms. These algorithms are suitable for the solution of high-performance system identification and parameter estimation problems with *convergence certificates* in settings where the standard persistence of excitation condition is difficult to guarantee or verify *a priori*. To achieve (uniform) fast and fixed-time convergence, the proposed algorithms exploit the existence of information-rich data sets, as well as certain non-smooth regularizations of dynamical systems that generate a family of non-Lipschitz systems modeled as data-driven ordinary differential equations (DD-ODEs) and/or data-driven hybrid dynamical systems (DD-HDS). In each scenario, we provide stability and convergence certificates via Lyapunov theory. Moreover, to illustrate the practical advantages of the proposed algorithms, we consider an online estimation problem in Lithium-Ion batteries where the satisfaction of the persistence of excitation condition is in general difficult to guarantee.

Keywords: Concurrent learning, adaptive control, hybrid systems, Lyapunov theory.

1. Introduction

In this paper, we study efficient algorithms for online parameter estimation problems which can be cast as uncertain linear parametric models of the form

$$y(t) = \phi(t)^\top \theta^*, \quad (1)$$

where $y : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is a measurable signal, $\phi : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^n$ is a uniformly bounded vector-valued regressor function, and $\theta^* \in \mathbb{R}^n$ is an unknown parameter that we want to estimate. This problem plays an important role in different areas, such as adaptive control [Ioannou and Sun \(2012\)](#), model-free optimization of dynamical systems [Poveda et al. \(2020\)](#), and reinforcement learning [Kaelbling et al. \(1996\)](#), to name just a few. To achieve online parameter estimation with convergence and robustness certificates, different feedback-based algorithms have been proposed during the last three decades; see [Boyd and Sastry \(1986\)](#);

Narendra and Annaswamy (1987, 2012). It is well-known that most of the adaptive estimation dynamics that achieve uniform convergence to the true parameter θ^* require a persistence of excitation (PE) condition in the regressor ϕ , of the form

$$\int_t^{t+T} \phi(s)\phi(s)^\top ds \succ \varepsilon I, \quad \forall t \geq t_0, \text{ where } T, \varepsilon > 0. \quad (2)$$

Indeed, in several adaptive estimation dynamics the PE condition has been shown to be *sufficient* and *necessary* to achieve (uniform) exponential convergence. This includes, the so-called gradient method Praly (2017):

$$\dot{\hat{\theta}} = -\sigma\phi(t) \left(\phi(t)^\top \hat{\theta} - y(t) \right), \quad \sigma \in \mathbb{R}_{>0}, \quad (3)$$

which has been widely used in academic and industrial applications. To relax the PE condition, the works Chowdhary and Johnson (2010), Chowdhary et al. (2012b), and Chowdhary et al. (2012a) introduced a class of concurrent learning (CL) adaptive dynamics that incorporate a sequence of recorded data $\{\phi(t_k)\}_{k=1}^N$ that is “sufficiently rich”, resulting in a data-driven ordinary differential equation (DD-ODE) of the form

$$\dot{\hat{\theta}} = -\sigma\phi(t) \left(\phi(t)^\top \hat{\theta} - y(t) \right) - \rho \sum_{k=1}^N \phi(t_k) \left(\phi(t_k)^\top \hat{\theta} - y(t_k) \right), \quad (4)$$

where $\sigma \in \mathbb{R}_{\geq 0}$ and $\rho \in \mathbb{R}_{>0}$ are tunable gains. These types of algorithms have been extended in several directions to develop PE-free adaptive dynamics in the context of model-reference adaptive control Chowdhary and Johnson (2010), reinforcement learning Kamalapurkar et al. (2014), extremum seeking control Poveda et al. (2020), and general networked estimation problems Javed et al. (2021), to name just a few examples. However, by removing (or relaxing) the PE condition, these types of data-driven algorithms can also suffer from poor transient performance in terms of slow rates of convergence, especially when the matrix of recorded data is ill-conditioned. This behavior stems from the fact that systems of the form (3) or (4) can be cast as *time-varying* gradient flows for which the Hessian matrix might be degenerate whenever the PE condition is relaxed. Indeed, the slow learning rates that may emerge in CL have limited its application in practical engineering problems that require *fast* adaptation and/or estimation.

Motivated by this background, in this paper we introduce a novel class of concurrent learning algorithms able to achieve acceleration and/or fixed-time convergence properties. The dynamics make use of different types of regularization mechanisms that have been explored during the last years to design optimization algorithms and feedback controllers with high-transient performance, but which have never been studied in the context of CL. Given that the proposed dynamics are non-smooth, they are modeled either as non-Lipschitz ODEs Khalil (2002) or as hybrid dynamical systems Goebel et al. (2012). For these systems, we exploit hybrid Lyapunov-based methods to establish suitable stability and convergence properties. Moreover, to illustrate the performance of our algorithms, we study a battery estimation problem where the satisfaction of the PE condition is usually difficult to verify *a priori*, and where the standard concurrent learning algorithm of Chowdhary and Johnson

(2010) can lead to prohibitively slow convergence rates. As evidenced by the numerical experiments, the proposed fast algorithms significantly outperform the standard CL dynamics in terms of transient performance and steady state error.

The rest of this paper is organized as follows: Section 2 introduces the notation used in the paper, as well as some preliminaries on hybrid dynamical systems. Section 3 presents the main results, Section 4 presents the numerical experiments, and Section 5 ends with some conclusions. The proofs of the results are presented in the Appendix of the extended technical report [Ochoa et al. \(2020\)](#).

2. Notation and Preliminaries on Hybrid Dynamical Systems

We define $\mathbf{c}_n \in \mathbb{R}^n$ as the vector with all entries equal to $c \in \mathbb{R}$, and use $|\cdot|$ as the Euclidean norm. We use $|x|_{\mathcal{A}} := \inf_{y \in \mathcal{A}} |x - y|$ to denote the distance of a vector $x \in \mathbb{R}^n$ with respect to a closed set \mathcal{A} . To simplify notation, and given vectors $x \in \mathbb{R}^{n_1}$, $y \in \mathbb{R}^{n_2}$, and $z \in \mathbb{R}^{n_3}$, we use (x, y, z) to denote the column vector $[x^\top, y^\top, z^\top]^\top$. We use $I_n \in \mathbb{R}^{n \times n}$ to denote the identity matrix, and S^n to denote the n^{th} -Cartesian product of the set S .

2.1. Hybrid Dynamical Systems

In this paper, we will model our algorithms as Hybrid Dynamical Systems with time-varying flows, of the form (see [Goebel et al. \(2012\)](#)):

$$(x, s) \in C \times \mathbb{R}_{\geq 0}, \quad \dot{x} = F(x, s), \quad \dot{s} = 1, \quad (5a)$$

$$(x, s) \in D \times \mathbb{R}_{\geq 0}, \quad x^+ = G(x), \quad s^+ = s, \quad (5b)$$

where $x \in \mathbb{R}^m$ is the main state of the system, s is an auxiliary state used to model the evolution of the continuous time, $F : \mathbb{R}_{\geq 0} \times \mathbb{R}^m \rightarrow \mathbb{R}^m$ is called the flow map, and $G : \mathbb{R}^m \rightarrow \mathbb{R}^m$ is called the jump map. The sets C and D , called the flow set and the jump set, respectively, condition the points in \mathbb{R}^m where the system can *flow* or *jump* via equations (5a) or (5b), respectively. In this way, the HDS can be represented by the notation $\mathcal{H} := \{C, F, D, G\}$. Systems of the form (5) can be seen as generalizations of purely continuous-time systems ($D = \emptyset$) and purely discrete-time systems ($C = \emptyset$). Solutions to HDS of the form (5) are parameterized by both a continuous-time index $t \in \mathbb{R}_{\geq 0}$, which increases continuously during flows, and a discrete-time index $j \in \mathbb{Z}_{\geq 0}$, which increments by one during jumps. Thus, the notation \dot{x} in (5a) represents the derivative $\frac{dx(t,j)}{dt}$; and x^+ in (5b) represents the value of x after an instantaneous jump, i.e., $x(t, j + 1)$. Naturally, solutions $(x, s) : \text{dom}(x, s) \rightarrow \mathbb{R}^m$ to (5) are defined on hybrid time domains. For a precise definition of hybrid time domains and solutions to HDS (5) we refer the reader to ([Goebel et al., 2012, Ch.2](#)). In some cases, our models will not depend on the auxiliary state s .

2.2. Stability and Convergence Notions

To model the different convergence properties of our algorithms, we make use of class \mathcal{KL} functions β , which are continuous functions that satisfy $\lim_{r \rightarrow 0^+} \beta(r, \nu) = 0$ for each fixed $\nu \in \mathbb{R}_{\geq 0}$, $\lim_{\nu \rightarrow \infty} \beta(r, \nu) = 0$ for each fixed $r \in \mathbb{R}_{\geq 0}$, and which are non-decreasing in its first argument, and non-increasing in the second argument. Class \mathcal{KL} functions are standard in

the feedback control; see [Khalil \(2002\)](#); [Goebel et al. \(2012\)](#). Moreover, these functions can model different types of convergence properties depending on the structure of β .

Definition 1 *Let $\mathcal{A} \subset \mathbb{R}^m$ be a closed set, and suppose every solution of (5) satisfies $|x(t, j)|_{\mathcal{A}} \leq \beta(|x(0, 0)|_{\mathcal{A}}, t + j)$, $\forall (t, j) \in \text{dom}(x)$. Then, the set \mathcal{A} is said to be:*

- (a) *Uniformly Globally Asymptotically Stable (UGAS) if β is of class \mathcal{KL} .*
- (b) *Uniformly Globally Exponentially Stable (UGES) if $\beta(r, s) = c_1 r e^{-c_2 s}$, $c_1, c_2 > 0$.*
- (c) *Uniformly Globally Finite-time Stable (UGFS) if it is UGAS and there exists a continuous function $T : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ such that $\lim_{s \rightarrow T(r)} \beta(r, s) = 0$.*
- (d) *Uniformly Globally Fixed-time Stable (UGFXS) if it is UGFS and there exists $T^* > 0$ such that $T(r) < T^*$ for all $r \in \mathbb{R}_{\geq 0}$.*

Note that all the properties listed in Definition 1 are stronger than standard convergence notions used in offline optimization or estimation algorithms. In particular, UGAS implies not only convergence in the standard limiting sense, but also *uniform* global stability (in the sense of Lyapunov) and *uniform* global attractivity.

3. Accelerated Adaptive Concurrent Learning Dynamics

To describe the dynamics considered in this paper for the estimation of θ^* in (1), let the mappings $\Psi : \mathbb{R} \rightarrow \mathbb{R}^n$ and $B : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be defined as

$$\Psi(t) = \frac{\phi(t)}{(1 + \phi(t)^\top \phi(t))^2}, \quad \text{and} \quad B(\hat{\theta}) := \sum_{k=1}^N \Psi(t_k) \left(\phi(t_k)^\top \hat{\theta} - y(t_k) \right). \quad (6)$$

Using (6) and (1), the DD-ODE (4) can be written as a time-invariant dynamical system of the form

$$(\hat{\theta}, s) \in \mathbb{R}^n \times \mathbb{R}_{\geq 0}, \quad \dot{\hat{\theta}} = -\sigma A(s, \hat{\theta}) - \rho B(\hat{\theta}), \quad \dot{s} = 1, \quad (7)$$

where $A(s, \hat{\theta}) := \Psi(s)(\phi(s)^\top \hat{\theta} - y(s))$. Taking system (7) as a benchmark, we will construct four different data-driven CL dynamics that will achieve (uniform) global asymptotic convergence, exponential convergence, finite-time convergence, and fixed-time convergence, respectively, to the true parameter $\mathcal{A}_0 := \{\theta^*\}$. The convergence of these dynamics will depend on the ‘‘richness’’ properties of the available recorded data, a notion that is captured by a finite-time version of persistence of excitation; see [Astrom and Wittenmark \(1989\)](#).

Assumption 3.1 *Let $\{\phi(t_k)\}_{k=1}^N$ be a sequence of recorded data. Then, the matrix $D := [\phi(t_1), \phi(t_2), \dots, \phi(t_N)] \in \mathbb{R}^{n \times N}$ satisfies $\text{rank}(D) = p$.*

Sequences of data satisfying Assumption 3.1 are said to be *sufficiently-rich* (SR). The following lemma provides an equivalent (and instrumental) characterization of SR data.

Lemma 2 *Let $\{\phi(t_k)\}_{k=1}^N$ be a sequence of recorded data, and let $P := \sum_{k=1}^N \frac{\phi(t_k)\phi(t_k)^\top}{(1 + \phi(t_k)^\top \phi(t_k))^2}$. Then $\{\phi(t_k)\}_{k=1}^N$ is SR if and only if there exists $\gamma \in \mathbb{R}_{> 0}$ such that $P \succeq \gamma I_n$.*

We call the constant γ the *level of richness* of the data $\{\phi(t_k)\}_{k=1}^N$.

3.1. Data-Driven Accelerated Hybrid Dynamics with Periodic Restarting

The first dynamical system that we consider is inspired by Nesterov’s ODEs studied in the context of accelerated optimization; see [Su et al. \(2016\)](#) and [Wibisono et al. \(2016\)](#). Such algorithms can induce suitable acceleration properties by incorporating dynamic momentum, emulating in continuous time the acceleration properties of Nesterov’s accelerated optimization algorithm; see [Nesterov \(2004\)](#). However, unlike the results of [Su et al. \(2016\)](#) and [Wibisono et al. \(2016\)](#), in the setting of CL we are also interested in establishing suitable robustness properties that are relevant in applications where noisy measurements are unavoidable. Such robustness properties can be obtained by endowing the dynamics with discrete-time restarting mechanisms that persistently reset the momentum coefficient/state of the dynamics. The combination of continuous-time and discrete-time dynamics leads to a hybrid regularization of the Nesterov’s discrete-time algorithm which, for time-invariant problems, has been modeled as a HDS of the form (5) in [Poveda and Li \(2019\)](#), [Ochoa et al. \(2019\)](#), and [Ochoa et al. \(2021\)](#). Based on this setting, the *hybrid accelerated concurrent learning* (HACL) dynamics that we consider in this paper are modeled by a HDS with state $x := (\hat{\theta}, p, \tau)$, where $\hat{\theta}$ is the estimation state, $p \in \mathbb{R}^n$ is the momentum state, and $\tau \in \mathbb{R}_{>0}$ is a resetting state. The dynamics are given by

$$x \in C := \left\{ x \in \mathbb{R}^{2n+2} : \tau \in [\delta, \Delta] \right\}, \quad \begin{pmatrix} \dot{x} \\ \dot{s} \end{pmatrix} = \begin{pmatrix} \dot{\hat{\theta}} \\ \dot{p} \\ \dot{\tau} \\ \dot{s} \end{pmatrix} = F(x, s) := \begin{pmatrix} \frac{2}{\tau} (p - \hat{\theta}) \\ -2k\tau \left(\sigma A(s, \hat{\theta}) + \rho B(\hat{\theta}) \right) \\ \frac{1}{2} \\ 1 \end{pmatrix}, \quad (8a)$$

$$x \in D := \left\{ x \in \mathbb{R}^{2n+2} : \tau = \Delta \right\}, \quad \begin{pmatrix} x^+ \\ s^+ \end{pmatrix} = \begin{pmatrix} \hat{\theta}^+ \\ p^+ \\ \tau^+ \\ s^+ \end{pmatrix} = G(x) := \begin{pmatrix} \hat{\theta} \\ (1-q)p + q\hat{\theta} \\ \delta \\ s \end{pmatrix}, \quad (8b)$$

where $k \in \mathbb{R}_{>0}$ is a tunable gain, $\infty > \Delta > \delta > 0$ are tunable parameters that describe how frequently the algorithm resets, and $q \in \{0, 1\}$ is a Boolean variable that characterizes the resetting policy of the algorithm. In particular, when $q = 0$ the HACL only resets the coefficient τ , whereas when $q = 1$ the algorithm also resets the momentum state p . By construction, the discrete-time updates of the system are periodic and separated by intervals of flow of duration $2(\Delta - \delta)$. To guarantee suitable convergence properties, we will impose the following “data-driven” condition on the parameters (δ, Δ) and the gains (k, ρ) .

Assumption 3.2 *The tunable parameters (δ, Δ, k) satisfy $2k\rho(\Delta^2 - \delta^2) > \frac{1}{\gamma}$, where $\gamma > 0$ is given by Lemma 2.*

The following theorem, which is the first main result of this paper, characterizes the convergence properties of the HACL dynamics. Throughout the paper, we use $\tilde{\theta} := \hat{\theta} - \theta^*$ to denote the estimation error.

Theorem 3.1 *Suppose that Assumptions 3.1 and 3.2 hold. Then, every maximal solution of system (8) has an unbounded time domain, and the closed set $\mathcal{A} := \mathcal{A}_0 \times \mathcal{A}_0 \times [\delta, \Delta] \times \mathbb{R}_{\geq 0}$ is UGAS. Moreover, for each compact set of initial conditions $K \subset C \cup D$, the following convergence properties hold for all $(t, j) \in \text{dom}(x, s)$:*

(a) If $q = 0$, then for each $j \in \mathbb{Z}_{\geq 0}$ there exists $\beta_j \in \mathbb{R}_{>0}$ such that each trajectory of the system satisfies the bound

$$|\tilde{\theta}(t, j)|^2 \leq \frac{\beta_j}{k\rho\tau(t, j)^2}. \quad (9)$$

(b) If $q = 1$, the set \mathcal{A} is UGES, and each trajectory of the system satisfies the bound

$$|\tilde{\theta}(t, j)| \leq k_0 \tilde{\gamma}^j |\tilde{\theta}(0, 0)|, \quad \text{where } \tilde{\gamma} = \sqrt{\frac{1}{k\rho\Delta^2} \left(\frac{1}{2\gamma} + k\rho\delta^2 \right)} \in (0, 1), \quad k_0 = \frac{\Delta}{\delta}. \quad (10)$$

The result of Theorem 3.1 establishes two main convergence properties: item (a) establishes that the estimation error decreases at a rate of approximately $O(1/\tau^2)$ during flows, where τ increases linearly with time. On the other hand, item (b) establishes exponential convergence with a convergence rate adjustable via the values of $(\delta, \Delta, k, \rho)$. In this case, information-rich data sets ($\gamma \gg 1$) lead to faster rates of convergence. Optimal restarting periods, similar to those studied in O’Donoghue and Candes (2013) and Poveda and Li (2021), can also be derived for system (8).

3.2. Data-Driven Accelerated Hybrid Dynamics with Adaptive Restarting

The HACL dynamics (8) implement a periodic restarting mechanism that is coordinated by the state τ . In this subsection, we now consider an alternative approach based on *adaptive restarting*, where the momentum state is reset whenever a certain state-dependent condition is satisfied. Such type of mechanisms have been studied in the optimization literature; see Su et al. (2016), O’Donoghue and Candes (2013), Teel et al. (2019). However, in the context of CL, these types of mechanisms have remained mostly unexplored. To study this case, we introduce the function $H(\hat{\theta}, p) = \frac{|\hat{\theta}|_P^2}{2} + \frac{1}{2}|p|^2$, where $\tilde{\theta} := \hat{\theta} - \theta^*$, and where $|\cdot|_P : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ is a data-induced norm defined as $|u|_P^2 := u^\top P u$ for all $u \in \mathbb{R}^n$, with P as defined in Lemma 2. We then consider the *Hybrid Hamiltonian Concurrent Learning* (HHCL) algorithm, with $x = (\hat{\theta}, p, \tau) \in \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}_{\geq 0}$, and dynamics given by

$$x \in C := C_0 \times [0, \Delta], \quad \begin{pmatrix} \dot{\hat{\theta}} \\ \dot{p} \\ \dot{\tau} \end{pmatrix} = \begin{pmatrix} 0 & k\rho I_p & 0 \\ -k\rho I_p & 0 & 0 \\ 0 & 0 & k\rho \end{pmatrix} \begin{pmatrix} \frac{\partial H}{\partial p} \\ \frac{\partial H}{\partial \hat{\theta}} \\ 1 \end{pmatrix}, \quad (11a)$$

$$x \in D := (C_0 \times \{\Delta\}) \cup (D_0 \times [0, \Delta]), \quad \begin{pmatrix} \hat{\theta}^+ \\ p^+ \\ \tau^+ \end{pmatrix} = \begin{pmatrix} I_p & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \hat{\theta} \\ p \\ \tau \end{pmatrix}, \quad (11b)$$

where $\Delta := \frac{n\pi}{2k\rho\sqrt{\gamma}}$, γ is the level of richness of the data, and where

$$C_0 := \left\{ (\hat{\theta}, p) : \langle B(\hat{\theta}), p \rangle \leq 0 \right\}, \quad D_0 := \left\{ (\hat{\theta}, p) : \langle B(\hat{\theta}), p \rangle = 0 \ \& \ |p|^2 \geq |B(\hat{\theta})|^2 / \bar{\lambda} \right\},$$

with $\bar{\lambda} \geq \lambda_{\max}(P)$. Given that $\frac{\partial H}{\partial \hat{\theta}} = B(\hat{\theta})$ and $\frac{\partial H}{\partial p} = p$, the construction of the sets C_0 and D_0 indicate that system (11) is allowed to flow whenever there is no increase in the *potential* energy of the data-induced Hamiltonian function H .

Remark 3.1 *The role of the timer τ in system (11) is to guarantee the existence of an initial reset after an interval of flow of duration $\Delta > 0$. Once this reset has occurred, the update $p^+ = 0$ will guarantee that the next reset of the system will happen before $\tau = \Delta$, i.e., due to the condition $|p|^2 \geq |B(\hat{\theta})|^2 / \bar{\lambda}$ for all $x \in D$. Such types of bounds on the reset times have also been used in Teel et al. (2019) for standard optimization problems. However, in the context of CL, their application is new. In particular, note that the parameter Δ in the jump set D is now data-dependent.*

The following theorem is the second main result of this paper.

Theorem 3.2 *Suppose that Assumption 3.1 holds. Then, system (11) renders the set $\mathcal{A}_H = \mathcal{A}_0 \times \{0\} \times [0, \Delta]$ UGES, and every solution has an unbounded time-domain and satisfies the bound*

$$|\tilde{\theta}(t, j)| \leq \sqrt{\frac{2c_0}{\gamma}} \min \left\{ 1, e^{-\frac{\alpha}{2}(t-\Delta)} \right\} |\tilde{\theta}(0, 0)| \quad (12)$$

for all $(t, j) \in \text{dom}(x)$, where $\alpha = \frac{1}{\Delta} \ln \left(1 + \frac{\gamma}{\bar{\lambda}} \right)$, $c_0 > 0$.

3.3. Finite-Time and Fixed-Time Concurrent Learning Dynamics

While the hybrid CL dynamics (8) and (11) can induce sublinear and linear convergence rates, the convergence properties of the algorithms are still of *asymptotic* nature, i.e., $\theta(t) \rightarrow \theta^*$ only as $t \rightarrow \infty$. In this subsection, we consider a different class of learning dynamics able to achieve exact convergence to the true parameter θ^* in a *finite* amount of time. Moreover, in some cases this finite time can be upper bounded by a constant independent of the initial conditions of the estimate $\hat{\theta}$, which leads to *fixed-time* convergence guarantees.

In particular, to achieve finite time convergence we consider the *Finite-Time Concurrent Learning* (FTCL) dynamics modeled by the following non-smooth DD-ODE with $s \in \mathbb{R}_{\geq 0}$:

$$\hat{\theta} \in C := \mathbb{R}^n, \quad \dot{\hat{\theta}} = -k \frac{\sigma A(s, \hat{\theta}) + \rho B(\hat{\theta})}{|\rho B(\hat{\theta})|^{\frac{1}{2}}}, \quad \dot{s} = 1, \quad (13)$$

where $(k, \sigma, \rho) \in \mathbb{R}_{>0} \times \mathbb{R}_{\geq 0} \times \mathbb{R}_{>0}$ are tunable gains, and where the pair (A, B) is defined in (6) and (7). Note that this system is not Lipschitz continuous. However, under Assumption 3.1 and the uniform boundedness of the regressor ϕ , the vector field is everywhere continuous in $\hat{\theta}$ due to the structure of A and B . Thus, existence of solutions is guaranteed. For this system we establish the following result, which is the third main contribution of the paper.

Theorem 3.3 *Suppose that Assumption 3.1 holds. Then, system (13) renders the set \mathcal{A}_0 UGFTS, every solution has an unbounded time domain, and the settling time function satisfies $T(\hat{\theta}(0)) \leq \frac{2}{k\rho\gamma} |P| \sqrt{|\hat{\theta}(0) - \theta^*|}$.*

The result of Theorem 3.3 guarantees that $\hat{\theta}(t) = \theta^*$ for all $t \geq T(\hat{\theta}(0))$, where $T(\hat{\theta}(0))$ depends on the initial conditions of the estimate $\hat{\theta}(0)$, as well as the level of richness γ of the data. To remove the dependence on $\hat{\theta}(0)$ we can further consider a class of *Fixed-Time Concurrent Learning* (FXCL) dynamics, modeled by the following nonsmooth DD-ODE:

$$\hat{\theta} \in C := \mathbb{R}^n, \quad \dot{\hat{\theta}} = -k \frac{\sigma A(s, \hat{\theta}) + \rho B(\hat{\theta})}{|B(\hat{\theta})|^a} - k \frac{\sigma A(s, \hat{\theta}) + \rho B(\hat{\theta})}{|B(\hat{\theta})|^{-a}}, \quad \dot{s} = 1, \quad (14)$$

where $(k, \sigma, \rho) \in \mathbb{R}_{>0} \times \mathbb{R}_{\geq 0} \times \mathbb{R}_{>0}$ are tunable gains, and $a \in (0, 1)$ is a tunable exponent. System (14) is not Lipschitz continuous. However, under Assumption 3.1 and the uniform boundedness of the regressor, it is everywhere continuous in $\hat{\theta}$. The following theorem is the fourth main result of this paper.

Theorem 3.4 *Suppose that Assumption 3.1 holds. Then, system (14) renders UGFXS the set \mathcal{A}_0 with $T^* = \frac{\pi}{2a\gamma\rho k}$, and every solution has an unbounded time-domain.*


Remark 3.2 *Note that the fixed-time T^* is independent of the initial estimate $\hat{\theta}(0)$, but dependent on the level of richness of the data of the regressor, i.e., dependent on $\gamma > 0$.*

3.4. Discussion

The results of Theorems 3.1-3.4 establish new convergence bounds for CL algorithms that explicitly show the dependence on the richness of the data, i.e., the constant γ . In particular, while the standard CL dynamics of Chowdhary and Johnson (2010) achieve exponential convergence with rate of convergence proportional to the level or richness of the matrix P (cf. Lemma 2), under suitable tuning of the restarting parameters the data-driven hybrid dynamics introduced in this paper can achieve rates of convergence proportional to the squared root of the level of richness of the matrix P , see Poveda and Li (2019), Teel et al. (2019) and Poveda and Li (2021). This acceleration property is induced by the addition of momentum to the dynamics, and the design of the flow set and the jump set. Similarly, for the non-smooth DD-ODES (13) and (14), our results establish finite and fixed-time convergence bounds that are similar to those obtained in Poveda and Krstic (2021) for adaptive model-free optimization, but which are new in the context of CL, with an explicit characterization of the convergence time in terms of the level of richness of the data. Finally, note that for applications where γ is small the bound (9) establishes a desirable “semi-acceleration” property for estimation problems that lead to convex optimization problems that are not necessarily strongly convex.

4. Numerical Experiments: Li-ion battery parameter estimation

To illustrate our algorithms in a practical context, we apply the methods to solve an estimation problem related to the characterization of the impedance parameters of an equivalent circuit model of a Lithium-Ion (Li-Ion) battery. These types of batteries have strict requirements to achieve a reliable life-time operation, see (Reddy and Linden, 2011, Part 4) and (Plett, 2016, Volume II, Chapter 1). In this setting, it is customary to design real-time algorithms that control the charge and estimate internal states based on a dynamic model of the battery. Under limited computational capability, circuit based models are typically used in battery management systems. A first order equivalent circuit model of a Li-ion battery is described below Feng et al. (2015); Liaw et al. (2004): $\dot{z} = \frac{I}{C_0}$, $\dot{i}_1 = \frac{I - i_1}{R_1 C_1}$, $V = \Phi(z) + IR_0 + i_1 R_1$, where $C_0[As]$ denotes the battery capacity, $R_0[\Omega], R_1[\Omega], C_1[F]$ denote the battery impedance parameters, z denotes the state of charge (SOC) of the battery, i.e. the relative capacity, $I[A]$ denotes the input current, $i_1[A]$ is the current through the parallel resistor R_1 , and Φ is mapping from state of charge to open circuit voltage. The impedance parameters R_0, R_1, C_1 are typically functions of the state of charge Hannan

Figure 1: Battery model simulation with nominal value of R_0

et al. (2017). For ease of presentation, we make some simplifying assumptions regarding the estimation problem. We note that the methods described here can be extended to the case without such assumptions, whereas standard methods like Recursive least squares are not suitable due to the form of the parameter dependency on the battery state of charge. The capacity C_0 evolves on a slower time scale, and it is typically estimated using rest voltage measurements Subbaraman et al. (2019), which decouples it from the battery impedance parameter estimation which can change over a single charge/discharge cycle due to variations of state of charge, temperature, etc. If C_0 , the initial SOC, and the OCV-SOC relation are known at least over a cycle, the battery SOC can be calculated accurately using Coulomb counting. Thus, we assume that the battery capacity and the initial SOC are known, and the estimation is restricted to battery impedance parameters, for which the functional dependency of SOC is restricted to R_0 , whereas R_1 and C_1 are constant. Thus, we write R_0 as $R_0(z) = \sum_{i=0}^N \alpha_i z^i$, and we define the parameter vector $w = [\alpha_0, \dots, \alpha_N, R_1]$. Then, the system's model can be rewritten as $(V - \Phi(z)) = [I, \dots, z^N I, i_1]w$. The initial SOC is 1.0, $C_0 = 3.4\text{Ah}$, and $R_1 C_1 = 100$. A sample drive-cycle behavior and the nominal value of R_0 as a function of the SOC are illustrated in Figure 1.

Simulation Results: We select the regressor vector as $\phi = [I, Iz, \dots, Iz^N, i_1]^\top$, with $N = 7$, where I and i_1 are the input-current and the current through resistor R_1 , respectively. To achieve good numerical performance, in addition to being compliant with Assumption 3.1, we pay special attention to the convergence rates of the different dynamics and how they are closely related to the magnitude of the minimum eigenvalue of the matrix P . With this in mind, we aim to select data that is both SR, and which at the same time maximizes said eigenvalue. To do this, we follow the procedure presented in Chowdhary and Johnson (2011), and we find a subset of 50 data points from the battery sample drive-cycle behavior. We highlight that the resulting $\lambda_{\min}(P) \approx 10^{-6}$, and even though this value can increase slightly when larger data-sets are considered, the size is kept small in order to facilitate faster computation, specially for the FTCL and FXCL algorithms which involve calculating the norm of B . With the selected data-set we implement the different methods presented in Section 3 and compare them with the standard concurrent learning dynamics (4). The results are presented in Figure 2. All the proposed algorithms outperform the standard CL approach for the number of iterations considered.

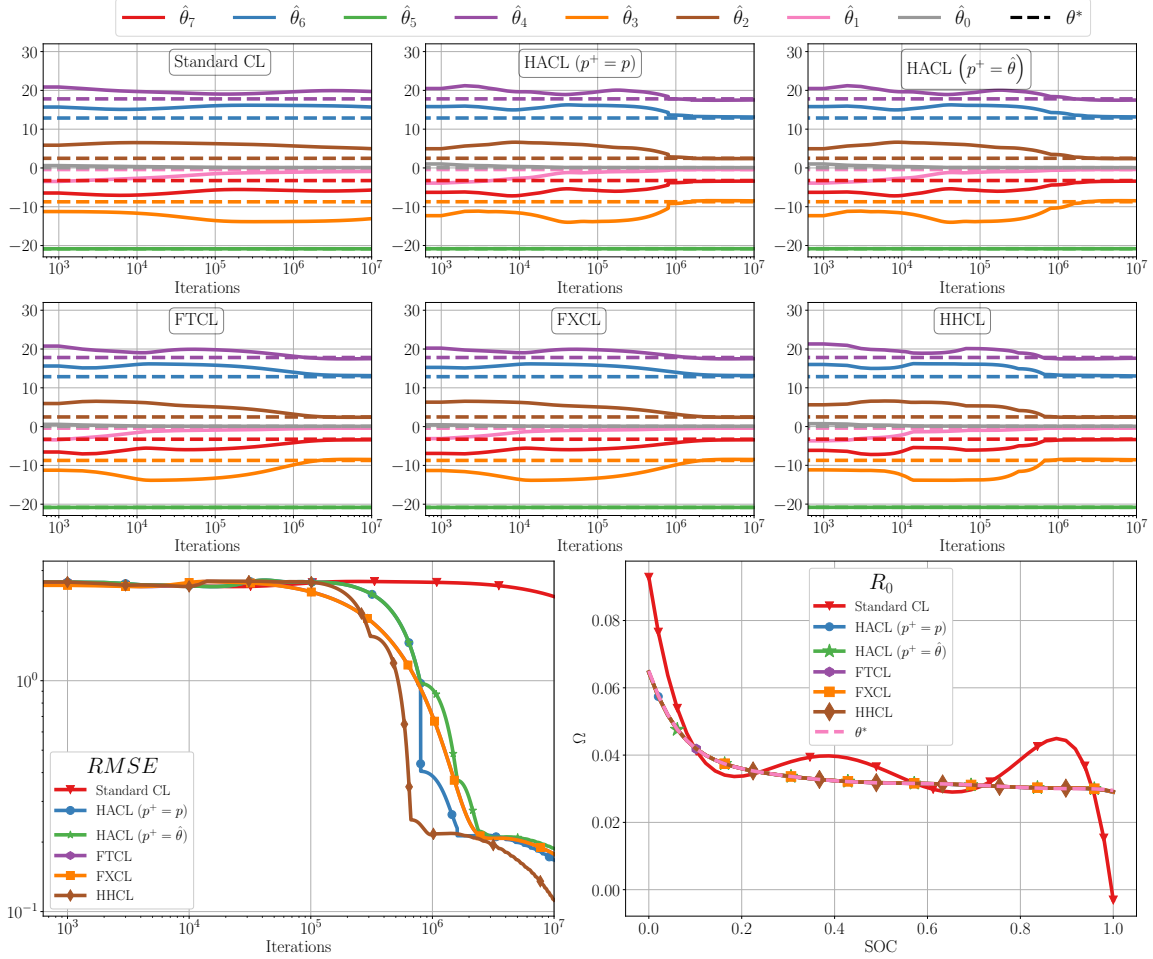


Figure 2: Estimation of parameter R_0 in a Li-Ion battery using the proposed algorithms.

5. Conclusions

In this paper, we introduced a new class of concurrent learning algorithms with acceleration and finite/fixed-time convergence properties. The algorithms are suitable for identification and parameter estimation problems that arise in the context of adaptive control, model-free optimization, and reinforcement learning. The proposed algorithms are modeled as non-smooth ODEs or hybrid dynamical systems, for which suitable stability, convergence, and robustness properties can be established via Lyapunov-based tools and invariance principles. We illustrated the advantages of the methods via numerical examples in a Lithium-Ion battery estimation problem where standard concurrent learning dynamics can generate prohibitively slow rates of convergence. Future research directions will consider data-driven accelerated adaptive controllers for dynamical systems, as well as extensions to *multi-agent* problems and algorithms based on *multi-time scale* stochastic hybrid dynamics.

References

- K. J. Astrom and Bjorn Wittenmark. *Adaptive Control*. Addison-Wesley Publishing Company, 1989.
- S. Boyd and S. S. Sastry. Necessary and sufficient conditions for parameter convergence in adaptive control. *Automatica*, 22(6):629–639, 1986.
- G. Chowdhary and E. Johnson. Concurrent learning for convergence in adaptive control without persistence of excitation. *49th IEEE Conference on Decision and Control*, pages 3674–3679, 2010.
- G. Chowdhary and E. Johnson. A singular value maximizing data recording algorithm for concurrent learning. In *Proceedings of the 2011 American Control Conference*, pages 3547–3552. IEEE, 2011.
- G. Chowdhary, T. Wu, M. Cutler, N. K. Ure, and J. P. How. Experimental results of concurrent learning adaptive controllers. *AIAA Guidance, Navigation and Control Conference*, pages 1–14, 2012a.
- G. Chowdhary, T. Yucelen, M. Muhlegg, and E. N. Johnson. Concurrent learning adaptive control of linear systems with exponentially convergent bounds. *International Journal of Adaptive Control and Signal Processing*, 27(4):280–301, 2012b.
- T. Feng, L. Yang, X. Zhao, H. Zhang, and J. Qiang. Online identification of lithium-ion battery parameters based on an improved equivalent-circuit model and its implementation on battery state-of-power prediction. *Journal of Power Sources*, 281:192–203, 2015.
- R. Goebel, R. G. Sanfelice, and A. R. Teel. *Hybrid Dynamical Systems: Modeling, Stability, and Robustness*. Princeton University Press, 2012.
- M. A Hannan, MS. H Lipu, A. Hussain, and Azah Mohamed. A review of lithium-ion battery state of charge estimation and management system in electric vehicle applications: Challenges and recommendations. *Renewable and Sustainable Energy Reviews*, 78:834–854, 2017.
- P. A. Ioannou and J. Sun. *Robust Adaptive Control*. Dover Publications Inc., Mineola, NY., 2012.
- M. U. Javed, J. I. Poveda, and X. Chen. Excitation conditions for uniform exponential stability of the cooperative gradient algorithm over weakly connected digraphs. *IEEE Control and Systems Letters*, DOI:10.1109/LCSYS.2021.3049153, pages 1–6, 2021.
- L. Pack Kaelbling, M. L. Littman, and A. W. Moore. Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4:237–285, 1996.
- R. Kamalapurkar, J. R. Klotz, and W. E. Dixon. Concurrent learning-based approximate feedback-nash equilibrium solution of n-player nonzero-sum differential games. *IEEE/CAA Journal of Automatica Sinica*, 1:239–247, 2014.

- H. K. Khalil. *Nonlinear Systems*. Prentice Hall, Upper Saddle River, NJ, 2002.
- B. Y. Liaw, G. Nagasubramanian, R. G Jungst, and D. H Doughty. Modeling of lithium ion cells—a simple equivalent-circuit model approach. *Solid state ionics*, 175(1-4):835–839, 2004.
- K. S. Narendra and A. Annaswamy. Persistent excitation in adaptive systems. *International Journal of Control*, 45(1):127–160, 1987.
- K. S. Narendra and A. M. Annaswamy. *Stable Adaptive Systems*. Courier Corporation, 2012.
- Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic course*. Kluwer Academic Publishers, Boston, MA., 2004.
- D. Ochoa, J. I. Poveda, C. Uribe, and N. Quijano. Optimal resource allocation with momentum. *IEEE Conference on Decision and Control*, pages 3954–3959, 2019.
- D. Ochoa, J. I. Poveda, C. Uribe, and N. Quijano. Robust optimization over networks using distributed restarting of accelerated dynamics. *IEEE Control Systems Letters*, DOI 10.1109/LCSYS.2020.3001632, pages 1–6, 2021.
- D. E. Ochoa, J. I Poveda, A. Subbaraman, G. S. Schmidt, and F. R. Pour-Safaei. Accelerated concurrent learning algorithms via data-drive hybrid dynamics and nonsmooth odes. *Technical Report, University of Colorado, Boulder*, 2020. <http://deot95.github.io/technicalreports/14dc2021>.
- O’Donoghue and E. J. Candes. Adaptive restart for accelerated gradient schemes. *Foundations of Computational Mathematics*, 15(3):715–732, 2013.
- G. L. Plett. *Battery Management Systems - Volume I - Battery Modeling & Volume II - Equivalent-Circuit Methods*. Artech House, 2016.
- J. I. Poveda and M. Krstic. Non-smooth extremum seeking control with user-prescribed fixed-time convergence. *IEEE Transactions on Automatic Control*, 10.1109/TAC.2021.3063700, 2021.
- J. I. Poveda and N. Li. Inducing uniform asymptotic stability in non-autonomous accelerated optimization dynamics via hybrid regularization. *58th IEEE Conference on Decision and Control*, pages 3000–3005, 2019.
- J. I. Poveda and N. Li. Robust hybrid zero-order optimization algorithms with acceleration via averaging in time. *Automatica*, 123, 2021.
- J. I. Poveda, M. Benosman, and K. Vamvoudakis. Data-Enabled Extremum Seeking: A Cooperative Concurrent Learning-Based Approach. *International Journal of Adaptive Control and Signal Processing*, DOI: 10.1002/acs.3189, pages 1–29, 2020.
- L. Praly. Convergence of the gradient algorithm for linear regression models in the continuous and discrete time cases. *Research Report, PSL Research University, Mines ParisTech*, 2017.

- T. B. Reddy and D. Linden, editors. *Linden's Handbook of Batteries*. McGraw Hill, 2011.
- W. Su, S. Boyd, and E. Candes. A differential equation for modeling Nesterov's accelerated gradient method: Theory and insights. *J. of Machine Learning Research*, 17(153):1–43, 2016.
- A. Subbaraman, N. Ravi, R. Klein, G. S. Schmidt, and C. Mayhew. Method and system for estimating battery open cell voltage, state of charge, and state of health during operation of the battery. *US Patent 10,312,699*, 2019.
- A. R. Teel, J. I. Poveda, and J. Le. First-order optimization algorithms with resets and hamiltonian flows. *In proc. of IEEE Conference on Decision and Control*, pages 5838–5843, 2019.
- A. Wibisono, A. C. Wilson, and M. I. Jordan. A variational perspective on accelerated methods in optimization. *Proceedings of the National Academy of Sciences*, 113(47):E7351–E7358, 2016.