Automatica 56 (2015) 44–52

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Brief paper

Extremum seeking of dynamical systems via gradient descent and stochastic approximation methods☆

Sei Zhen Khong[a,1], Ying Tan[b], Chris Manzie[c], Dragan Nešić[b]

[a] Department of Automatic Control, Lund University, SE-221 00 Lund, Sweden

[b] Department of Electrical and Electronic Engineering, The University of Melbourne, Parkville, VIC 3010, Australia

[c] Department of Mechanical Engineering, The University of Melbourne, Parkville, VIC 3010, Australia

ARTICLE INFO

Article history:
Received 15 August 2013
Received in revised form 13 October 2014
Accepted 8 March 2015
Available online 11 April 2015

Keywords:
Extremum seeking
Infinite-dimensional nonlinear systems
Sampled-data control
Gradient descent method
Stochastic approximation methods

ABSTRACT

This paper examines the use of gradient based methods for extremum seeking control of possibly infinite-dimensional dynamic nonlinear systems with general attractors within a periodic sampled-data framework. First, discrete-time gradient descent method is considered and semi-global practical asymptotic stability with respect to an ultimate bound is shown. Next, under the more complicated setting where the sampled measurements of the plant's output are corrupted by an additive noise, three basic stochastic approximation methods are analysed; namely finite-difference, random directions, and simultaneous perturbation. Semi-global convergence to an optimum with probability one is established. A tuning parameter within the sampled-data framework is the period of the synchronised sampler and hold device, which is also the waiting time during which the system dynamics settle to within a controllable neighbourhood of the steady-state input–output behaviour.

1. Introduction

Extremum seeking locates via online computations an optimal operating regime of the steady-state input–output map of a dynamical system without explicit knowledge of a model (Ariyur & Krstić, 2003; Zhang & Ordóñez, 2011). Two categories of extremum seeking controllers can be found in the literature. The first of which is continuous-time controllers which exploit dither/excitation signals to probe the local behaviour of the system to be optimised and continuously transition the system input to one that results in an optimum. See Ariyur and Krstić (2003), Krstić and Wang (2000) and Tan, Nešić, and Mareels (2006) for such methods that utilise periodic dithers and Liu and Krstić (2012), Manzie and Krstić (2009) for stochastic dithers. The convergence proofs of the former rely on averaging and singular perturbation techniques (Khalil, 2002; Teel, Moreau, & Nešić, 2003), while the latter on stochastic averaging (Liu & Krstić, 2012). On the contrary, discrete-time extremum seeking controllers based on nonlinear programming methods are examined in Teel and Popović (2001) within a sampled-data framework. The convergence proof therein is established using Lyapunov arguments.

An alternative and more direct proof for convergence to an extremum in a sampled-data framework is given in Khong, Nešić, Tan, and Manzie (2013) using trajectory-based techniques. In the same paper, the sampled-data framework of extremum seeking is further examined to accommodate global nonconvex optimisation methods, such as those described in Strongin and Sergeyev (2000). These results demonstrate that a wide range of optimisation algorithms in the literature can be applied to extremum seeking of dynamic plants. Making use of the results in Khong, Nešić, Tan et al. (2013), deterministic gradient descent based extremum seeking control is reviewed in this paper. Furthermore, stochastic gradient descent (a.k.a. stochastic approximation) methods are accommodated for extremum seeking in a way that is robust against measurement errors.

Stochastic approximation methods (Kushner & Clark, 1978; Kushner & Yin, 2003; Spall, 2003) are a family of well-studied iterative gradient-based optimisation algorithms that find applications in a broad range of areas, such as adaptive control and neural networks (Bertsekas & Tsitsiklis, 1996). In contrast to the standard

☆ This work was supported by the Swedish Research Council through the LCCC Linnaeus centre and the Australian Research Council (DP120101144). The material in this paper was presented at the 9th Asian Control Conference, June 23–26, 2013, Istanbul, Turkey. This paper was recommended for publication in revised form by Associate Editor Raul Ordóñez under the direction of Editor Miroslav Krstic.

E-mail addresses: seizhen@control.lth.se (S.Z. Khong), yingt@unimelb.edu.au (Y. Tan), manziec@unimelb.edu.au (C. Manzie), dnesic@unimelb.edu.au (D. Nešić).

[1] Tel.: +46 462220362; fax: +46 46138118.

http://dx.doi.org/10.1016/j.automatica.2015.03.018
0005-1098/

optimisation algorithms such as the steepest descent or Newton methods (Boyd & Vandenberghe, 2004) which exploit direct gradient information, stochastic approximation methods operate based on *approximation* to the gradient constructed from noisy measurements of the objective/cost function. For the former, knowledge of the underlying system input–output relationships are often needed to calculate the gradient using for example, the chain rule. This is not necessary for stochastic approximation, making it well-suited for non-model based extremum seeking control.

This paper adapts within a periodic sampled-data framework three discrete-time multivariate stochastic approximation algorithms for extremum seeking control of dynamical systems which can be of infinite dimension and contain general attractors. Namely, Kiefer–Wolfowitz–Blum's Finite Difference Stochastic Approximation (FDSA) (Blum, 1954; Kiefer & Wolfowitz, 1952), Random Directions Stochastic Approximation (RDSA) (Kushner & Clark, 1978), and Simultaneous Perturbation Stochastic Approximation (SPSA) (Spall, 1992, 2003). It is shown that there exists a sufficiently long sampling period under which semi-global convergence with probability one to an extremum of the steady-state input–output relation can be achieved. This stands in comparison with the gradient descent method based extremum seeking control under ideal noise-free sample measurements, for which semi-global practical ultimately bounded *asymptotic* stability can be established. Note that the existence of Lyapunov functions satisfying the conditions in Teel and Popović (2001) is not known for the stochastic approximation methods, and hence the convergence results therein do not directly generalise to these methods.

A related work (Nusawardhana & Żak, 2004) considers an extremum seeking method based on the SPSA within a different setup (i.e. not sampled-data and has continuous plant output measurements). There, the steady-state input–output objective function is assumed to evaluate to a constant after some waiting time with respect to a constant input, and the output measurements are corrupted by noise. By contrast, this paper exploits the fact that the state trajectory of an asymptotically stable dynamical system converges to a neighbourhood of its steady-state value after the system's input is held constant for a pre-selected waiting time. Furthermore, the sampled output value is assumed to be corrupted by measurement noise. The SPSA method has also been applied to optimisation of variable cam timing engine operation in Popović, Janković, Magner, and Teel (2006), alongside several other optimisation algorithms. Azuma, Sakar, and Pappas (2012) adapts the SPSA method for *extreme source seeking* of randomly switching *static* distribution fields using a nonholonomic mobile robot. On a different note, Stanković and Stipanović (2010) considers a related problem of extremum seeking of *static* functions under noisy measurements using a discrete-time controller with sinusoidal dither signals. These works differ from the setting of the paper, where stochastic approximation methods based extremum seeking of the steady-state input–output maps of *dynamical* systems is analysed within a sampled-data framework.

The paper has the following structure. First, the next section states the properties of the nonlinear dynamical systems to which gradient descent and stochastic optimisation methods are applied. Section 3 depicts the sampled-data framework in which extremum seeking control is analysed. Subsequently, Section 4 examines the gradient descent method for extremum seeking. Stochastic optimisation methods are considered in Section 5. Illustrative simulation examples are provided in Section 6, followed by some concluding remarks in Section 7.

2. Dynamical systems

The class of nonlinear, possibly infinite-dimensional, systems with general attractors considered in this paper is introduced in this section. A function $\gamma : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is of class-$\mathcal{K}$ (denoted $\gamma \in \mathcal{K}$) if it is continuous, strictly increasing, and $\gamma(0) = 0$. If γ is also unbounded, then $\gamma \in \mathcal{K}_\infty$. A continuous function $\beta : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is of class-$\mathcal{KL}$ if for each fixed t, $\beta(\cdot, t) \in \mathcal{K}$ and for each fixed s, $\beta(s, \cdot)$ is decreasing to zero (Khalil, 2002). The Euclidean norm is denoted $\|\cdot\|_2$.

Let $\mathcal{X}$ be a Banach space whose norm is denoted $\|\cdot\|$. Given any subset $\mathcal{Y}$ of $\mathcal{X}$ and a point $x \in \mathcal{X}$, define the distance of x from $\mathcal{Y}$ as $\|x\|_{\mathcal{Y}} := \inf_{a \in \mathcal{Y}} \|x - a\|$. Also let

$$\mathcal{U}_\epsilon(\mathcal{Y}) := \{x \in \mathcal{X} \mid \|x\|_{\mathcal{Y}} < \epsilon\}.$$

Definition 1. Let the state of a time-invariant dynamical system be represented by $x : \mathbb{R}_{\geq 0} \to \mathcal{X}$, where $\mathcal{X}$ is a Banach space with norm $\|\cdot\|$. The input to and output of the system are denoted, respectively, by $u : \mathbb{R}_{\geq 0} \to \mathbb{R}^m$ and $y : \mathbb{R}_{\geq 0} \to \mathbb{R}$. Given any $u \in \Omega \subset \mathbb{R}^m$ and $x_0 \in \mathcal{X}$, let $x(\cdot, x_0, u)$ be the state of the dynamical system starting at x_0 with input u.

Parts of the following assumption are based on Teel and Popović (2001, Assumption 1). Remarks on each of the assumptions follow.

Assumption 2. Given a system described in Definition 1 and an open bounded set $\Omega \subset \mathbb{R}^m$, the following hold:

(i) There exists a function $\mathcal{A}$ mapping from $\mathbb{R}^m$ to subsets of $\mathcal{X}$ such that for each constant $u \in \Omega$, $\mathcal{A}(u)$ is a nonempty closed set and a *global attractor* (Ruelle, 1989) which satisfies:
  (a) Given any $x_0 \in \mathcal{X}$ and $\epsilon > 0$, there exists a sufficiently large $t > 0$ such that $x(t, x_0, u) \in \mathcal{U}_\epsilon(\mathcal{A}(u))$;
  (b) If $x(t_0, x_0, u) \in \mathcal{A}(u)$, then $x(t, x_0, u) \in \mathcal{A}(u)$ for all $t \geq t_0$;
  (c) There exists no proper subset of $\mathcal{A}(u)$ having the first two properties above.

  Furthermore,
  $$\sup_{u \in \Omega} \sup_{x \in \mathcal{A}(u)} \|x\| < \infty.$$
(ii) Given any $\Delta > 0$, there exists a class-$\mathcal{KL}$ function β such that
  $$\|x(t, x_0, u)\|_{\mathcal{A}(u)} \leq \beta(\|x_0\|_{\mathcal{A}(u)}, t)$$
  for all $t \geq 0$, $u \in \Omega$, and $\|x_0\|_{\mathcal{A}(u)} \leq \Delta$.
(iii) There exists a locally Lipschitz function $h : \mathcal{X} \to \mathbb{R}$ such that the system output
  $$y(t) = h(x(t, x_0, u)) \quad \forall t \geq 0$$
  for any constant input $u \in \Omega$ and $x_0 \in \mathcal{X}$. Moreover, $h(x_a) = h(x_b)$ for every $x_a, x_b \in \mathcal{A}(u)$. Since $\mathcal{A}(u)$ is a global attractor and h is locally Lipschitz, for any $u \in \mathbb{R}^m$ and $x_0 \in \mathcal{X}$,
  $$\begin{aligned} Q(u) &:= \lim_{t\to\infty} h(x(t, x_0, u)) \\ &= h\left(\lim_{t\to\infty} x(t, x_0, u)\right) \\ &= h(x_l), \quad \text{for some } x_l \in \mathcal{A}(u) \end{aligned}$$
  is a well-defined steady-state input–output map.
(iv) Q is thrice continuously differentiable and has bounded derivatives on Ω.
(v) The Jacobian $\nabla Q = 0$ in a nonempty, compact set $\mathcal{C} \subset \mathbb{R}^m$, i.e. Q achieves its minimum on $\mathcal{C}$.

Remark 3. Property (i) of Assumption 2 states that for each constant input to the system, there exists a corresponding set to which the state of the system converges. Property (ii) stipulates that the state converges asymptotically stably. Property (iii) guarantees the existence of a corresponding output and hence an input–output map Q in steady state. The last two conditions are properties of Q which are assumed for convergence of the approximate gradient optimisation methods used in this paper, and are consistent with corresponding assumptions in e.g. Spall (2003).