Automatica 66 (2016) 238–245

Contents lists available at ScienceDirect

Automatica

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

Brief paper

Iterative learning control based on extremum seeking[☆]

Sei Zhen Khong[a,1], Dragan Nešić[b], Miroslav Krstić[c]

[a] Institute for Mathematics and its Applications, University of Minnesota, Minneapolis, MN 55455, USA
[b] Department of Electrical and Electronic Engineering, The University of Melbourne, Parkville, VIC 3010, Australia
[c] Department of Mechanical and Aerospace Engineering, University of California, San Diego, CA 92093-0411, USA

ARTICLE INFO

Article history:
Received 30 April 2015
Received in revised form
29 October 2015
Accepted 11 November 2015

Keywords:
Extremum seeking
Iterative learning control
Non-model based approach
Local and global optimisation

ABSTRACT

This paper proposes a non-model based approach to iterative learning control (ILC) via extremum seeking. Single-input–single-output discrete-time nonlinear systems are considered, where the objective is to recursively construct an input such that the corresponding system output tracks a prescribed reference trajectory as closely as possible on finite horizon. The problem is formulated in terms of extremum seeking control, which is amenable to a range of local and global optimisation methods. Contrary to the existing ILC literature, the formulation allows the initial condition of each iteration to be incorporated as an optimisation variable to improve tracking. Sufficient conditions for convergence to the reference trajectory are provided. The main feature of this approach is that it does not rely on knowledge about the system's model to perform iterative learning control, in contrast to most results in the literature.

1. Introduction

Iterative learning control (ILC) is a learning based method for tracking a prescribed trajectory. It carries out the same task multiple times with respect to recursively updated control inputs while improving the tracking performance by learning from previous executions (Moore, 1993; Moore, Dahleh, & Bhattacharyya, 1992; Xu & Tan, 2003). ILC is known to achieve good performance in the presence of repeating disturbances and certain model uncertainty due to its iteratively learning feature. Practically, ILC has been applied to a wide range of engineering applications, including robotics (Messner, Horowitz, Kao, & Boals, 1991), induction motors (Saab, 2004), rolling mills (Garimella & Srinivasan, 1998), stroke rehabilitation (Freeman, Rogers, Burridge, Hughes, & Meadmore, 2015; Freeman, Rogers, Hughes, Burridge, & Meadmore, 2012), and aluminium extruders (Pandit & Buchheit, 1999); see Ahn, Chen, and Moore (2007) for a classification of the ILC literature. It is also useful within the context of motion planning (Srinivasan & Ruina, 2006). Bristow, Tharayil, and Alleyne (2006) contains an excellent survey of a particular ILC algorithm by Moore (1993), where several topics in analysis (e.g. performance, transients, robustness) and design methods (e.g. plant inversion, quadratically optimal) are covered.

This paper proposes an *extremum-seeking* based framework within which to perform iterative learning control of discrete-time single-input–single-output time-varying nonlinear systems on finite horizon. It is noted here that multi-input–multi-output systems are addressable with the same approach. A key feature of extremum seeking is its ability to locate an optimum with respect to some measure without assuming knowledge about the underlying *models* governing the dynamics of the nonlinear systems (Ariyur & Krstić, 2003; Zhang & Ordóñez, 2011). Such knowledge may be unavailable due to the difficulty associated with modelling of complicated nonlinear systems. Extremum seeking has found applications in a wide array of problems, including biochemical reactors (Guay, Dochain, & Perrier, 2003; Wang, Krstić, & Bastin, 1999), gas-turbine combustors (Moase, Manzie, & Brear, 2010), power electronics (Scheinker, Bland, Krstić, & Audia, 2014), multi-agent source seeking (Khong, Tan, Manzie, & Nešić, 2014), and finite-horizon optimal control (Frihauf, Krstić, & Başar, 2013). Within the context of ILC, extremum seeking has been applied to pulse shaping in a double-pass laser amplifier (Ren, Frihauf, Rafac, & Krstić, 2012).

In this paper, we propose a unifying framework in which to apply optimisation-based extremum seeking algorithms to ILC in the spirit of Khong, Nešić, Tan, and Manzie (2013) and

☆ This work was supported by the Swedish Research Council through the LCCC Linnaeus Centre and the Australian Research Council. The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Warren E. Dixon under the direction of Editor Andrew R. Teel.

E-mail addresses: seizhen@control.lth.se (S.Z. Khong), dnesic@unimelb.edu.au (D. Nešić), krstic@ucsd.edu (M. Krstić).

[1] Tel.: +46 46 2220362; fax: +46 46 138118.

http://dx.doi.org/10.1016/j.automatica.2015.12.019
0005-1098/

Nešić, Tan, Moase, and Manzie (2010). In particular, the proposed framework is shown to be amenable to a broad range of local and global optimisation methods (Khong, Nešić, Manzie, & Tan, 2013; Pintér, 1996; Teel & Popović, 2001). This allows complexity of implementation and convergence speed of the algorithms to be taken into account in the control design stage. For instance, if large variations in the control input is undesirable but convergence to local optima is tolerable, local optimisation methods may be selected. Furthermore, Newton-based methods can be employed if a quadratic convergence rate is solicited. In the proposed framework, the cost function is defined as the distance between the system output and the reference trajectory. For local optimisation methods, ultimately bounded asymptotic stability of local minima is demonstrated. In the case of global optimisation, it is shown that the proposed ILC converges to a global minimum.

Several optimisation-based ILC methods can be found in the literature, but the vast majority of them rely on knowledge on the models. For instance, the updating control laws as well as convergence of the ILC methods in Gunnarsson and Norrlöf (2001) and Owens and Hätönen (2005) depend on the precise knowledge of the nominal model. Owens, Hatonen, and Daley (2009) proposes a robust monotone gradient-based scheme for ILC of linear time-invariant (LTI) systems, where the multiplicative modelling uncertainty is assumed to be bounded. The robustness analysis therein reminisces that performed in Bristow et al. (2006). Schoellig, Mueller, and D'Andrea (2012) considers the case where an LTI model is subject to noisy disturbances and proposes a combined model-based Kalman filter and convex optimisation approach to ILC. Mishra, Topcu, and Tomizuka (2011) proposes a primal barrier method to ILC of LTI systems contingent on the availability of knowledge about the gradient and Hessian of the quadratic cost function, which in turn is dependent on the models.

While the standard ILC literature considers learning controllers for systems that perform the same operation repetitively under the same initial conditions, we depart from such a setting and incorporate the initial conditions as parts of the optimisation variables, so that they may vary from one iteration to the next for improved tracking. Indeed, the former is subsumed by the latter by setting the initial conditions to be constant across all iterations. The formulation in this paper differs from that of repetitive control (Longman, 2000) and repetitive learning control (Sun, Ge, & Mareels, 2006), where the initial conditions of the current iteration are set to be the final conditions of the previous trial. It is also noteworthy that the proposed extremum seeking based ILC, which updates the control input signal, differs from iterative feedback tuning (Hjalmarsson, Gevers, Gunnarsson, & Lequin, 1998), where non-model based optimisation methods are exploited to iteratively tune controller's parameters in order to achieve tracking of an output trajectory given a fixed reference input.

The paper has the following structure. A formal definition of ILC and the class of nonlinear systems considered in this paper are stated in the next section. In Section 3, ILC is formulated in terms of an extremum seeking problem. Subsequently, local and global optimisation based extremum seeking approaches are discussed in Sections 4 and 5 respectively. Section 6 contains simulation examples illustrating the main results. Finally, some concluding remarks are provided in Section 7.

2. Iterative learning control

The problem of iterative learning control (ILC) is formulated in this section. The special case where the plant is linear time-invariant (LTI) and a commonly used ILC method are reviewed.

2.1. Nonlinear plants

Consider the following dynamical discrete-time time-varying nonlinear state-space system defined over a *finite* time interval/horizon $k = 0, 1, \ldots, T$:

$$\begin{aligned} x(k+1) &= f(x(k), u(k), k) \quad x(0) = \bar{x}; \\ y(k) &= h(x(k), u(k), k), \end{aligned} \tag{1}$$

where $f : \mathbb{R}^n \times \mathbb{R} \times \mathbb{T} \to \mathbb{R}^n$ and $h : \mathbb{R}^n \times \mathbb{R} \times \mathbb{T} \to \mathbb{R}$ are locally Lipschitz functions in each argument and $\mathbb{T} := \{0, 1, \ldots, T\}$. Repeated disturbances that are present on both the state-update differential and state-to-output algebraic equations are accounted for by f and h being functions of the time unit k. The corresponding input–output operator for system (1) is denoted by Σ, whereby $y = \Sigma(\bar{x}, u)$. Also, given a $z : \mathbb{T} \to \mathbb{R}$, define the ℓ^2 norm by

$$\|z\|_2 := \sqrt{\sum_{t=0}^{T} z(k)^2}.$$

With a slight abuse of notation, $\|v\|_2$ is also used to denote the Euclidean norm for the vector $v \in \mathbb{R}^n$. Only discrete-time plants are considered in this paper. It is a natural formulation because ILC uses information from previous trials which needs to be stored on suitable digital media. By the same token, the dynamics of the plant are assumed to evolve along a *finite* horizon $[0, T]$.

Denote by $r : \mathbb{T} \to \mathbb{R}$ the reference trajectory. The *control objective* is to construct a u^* and an $\bar{x}^*$ such that the corresponding system output $y^* = \Sigma(\bar{x}^*, u^*)$ tracks r as accurately as possible. In other words,

$$(\bar{x}^*, u^*) := \arg\min_{u \in \mathcal{U}; \bar{x} \in \Omega} \|r - \Sigma(\bar{x}, u)\|_2,$$

where $\mathcal{U}$ is an appropriate compact subset of $\{u : \mathbb{T} \to \mathbb{R}\}$ and Ω a compact subset of $\mathbb{R}^n$. In general, any ℓ^p-norm may be employable when defining the distance above. Note that a reference r may not be realisable by the system, i.e. there exist no $\bar{x}^*$ and u^* such that $\Sigma(\bar{x}^*, u^*) = r$. In this case, the achievable minimum of the optimisation problem above is nonzero. When Σ is an LTI operator, realisability of references may be studied using the notions of controllability and observability. Characterising this when Σ is nonlinear is a lot harder, and may require knowledge about the solutions to (1).

When f and h are known precisely, a brute-force optimisation over $\bar{x}, u(1), \ldots, u(T)$ can be used to generate a y^* such that the error $e(k) := r(k) - y^*(t)$ is minimised. Alternatively, should this prove to be an infeasible approach, by introducing an additional *iteration-time domain* j, several model-based ILC algorithms in the literature (Moore, 1993; Xu & Tan, 2003) can be used to iteratively design u_j based on previous trials' outputs $y_i = \Sigma(\bar{x}, u_i)$ for $i < j$ such that $u_j \to u^*$ in the ℓ^2-norm for a fixed $\bar{x}$ across all iterations. By tuning the parameters of the ILC algorithms appropriately, the desired transient properties, such as monotone convergence, may be achieved.

Control design in ILC can be specified in the following form. If $\bar{x}_j$ is the initial condition and u_j is the input applied to the plant at trial $j = 0, 1, 2, \ldots$ and $e_j := r - y_j = r - \Sigma(\bar{x}_j, u_j)$ is the resulting tracking error, the control design involves constructing an iteratively updated control law expressed as a functional relationship typified by the equation

$$\begin{aligned} \bar{x}_{j+1} &= g_1(e_j, \ldots, e_{j-s}, u_j, \ldots, u_{j-t}, \bar{x}_j, \ldots, \bar{x}_{j-t}) \\ u_{j+1} &= g_2(e_j, \ldots, e_{j-s}, u_j, \ldots, u_{j-t}, \bar{x}_j, \ldots, \bar{x}_{j-t}), \end{aligned}$$

where $s, t \leq j$. Ideally, the control law should have the property that $u_\infty := \lim_{j\to\infty} u_j = u^*$ and $\bar{x}_\infty := \lim_{j\to\infty} \bar{x}_j = \bar{x}^*$, or equivalently, $e_\infty := \lim_{j\to\infty} e_j = 0$. A looser requirement on this is that there exists some small $\epsilon > 0$ such that $\|u_\infty - u^*\|_2 < \epsilon$ and