Iterative Learning Control of Iteration-Varying Systems via Robust Update Laws with Experimental Implementation

Berk Altın[a,*], Jeroen Willems[d], Tom Oomen[b], Kira Barton[c]

[a] Department of Computer Engineering, University of California, Santa Cruz, CA 95064, USA
[b] Department of Mechanical Engineering, Eindhoven University of Technology, 5600 MB Eindhoven, The Netherlands
[c] Department of Mechanical Engineering, University of Michigan, Ann Arbor, MI 48109, USA
[d] Nobleo Technology B.V., 5627 BV Eindhoven, The Netherlands

ARTICLE INFO

ABSTRACT

Iterative learning control (ILC) is an efficient way of improving the tracking performance of repetitive systems. While ILC can offer significant improvement to the transient response of complex dynamical systems, the fundamental assumption of iteration invariance of the process limits potential applications. Utilizing abstract Banach spaces as our problem setting, we develop a general approach that is applicable to the various frameworks encountered in ILC. Our main result is that robust invariant update laws lead to stable behavior in ILC systems, where iteration-varying systems converge to bounded neighborhoods of their nominal counterparts when uncertainties are bounded. Furthermore, if the uncertainties are convergent along the iteration axis, convergence to the nominal case can be guaranteed.

1. Introduction

Iterative learning control (ILC) has been recognized as an efficient way of improving the tracking performance of repetitive systems since the early 1980s (Arimoto, Kawamura, and Miyazaki, 1984). ILC can offer significant improvement to the transient response of complex dynamical systems with a high level of uncertainty through relatively simple algorithms (Bristow, Tharayil, and Alleyne, 2006; Moore, 1993). The fundamental assumption that enables the success of these algorithms has been iteration invariance of the: 1) plant dynamics, 2) exogenous disturbances, 3) initial conditions, and 4) reference signals. This assumption greatly simplifies the ILC problem and enables the control engineer to design an asymptotically stable recurrence relation in the iteration domain by employing a contraction mapping. Even though the assumption is unrealistic, similar to feedback control of linear time-invariant (LTI) systems, it yields good results in practice provided that the variation of the process (dynamics, exogenous disturbances, initial conditions etc.) from trial to trial is small.

1.1. The feedback analogy

The restrictive nature of the invariance assumption is perhaps best understood via an analogy to feedback control, since a common

interpretation of ILC is that of a feedback controller in the iteration domain, as per the following discussion: Let $\overline{P}: U \to Y$ be a bounded linear operator, where U is the space of admissible inputs and Y is the space of outputs. Assuming that $\overline{P}$ is known and there are no exogenous signals apart from u_k affecting the output, the classical ILC problem can be stated as that of finding a controller C that maps the input history $u_0, u_1, \ldots, u_{k-1} \in U$ to the current input u_k, such that the output $y_k = \overline{P} u_k$ converges to a desired reference r in the image of $\overline{P}$ as $k \to \infty$. In most cases, C is designed to consider only the previous iteration, thus giving rise to the name first-order ILC. The internal model principle then dictates that the controller (update law) C includes integral action to guarantee perfect tracking in the limit, so $C(u_k) = u_{k-1} + L(r - \overline{P}u_{k-1})$, as can be seen in Fig. 1, which guarantees $y_k \to r$ even in the case where the output is corrupted by a constant vector $d \in Y$ such that $y_k = \overline{P}u_k + d$. Essentially, the ILC problem is that of designing a "time"-invariant feedback controller for a constant static plant to track step references (Moore, 1993), under the assumption of constant disturbance signals.

The objective of this paper is to generalize the ILC problem by relaxing the invariance assumption, which restricts the feedback analogy to setpoint tracking, and fails to capture the generality associated with the feedback paradigm. In practice, initial conditions and disturbances are *always subject to variations*, while references

* Corresponding author.
   E-mail addresses: berkaltin@ucsc.edu (B. Altın), jeroenmaartenwillems@gmail.com (J. Willems), T.A.Oomen@tue.nl (T. Oomen), bartonkl@umich.edu (K. Barton).

Fig. 1. Feedback control in the iteration domain interpretation of ILC: The integral action of the control law ensures that the output y_k converges to r for constant $d_k = d$ for the static plant $\overline{P}$, provided the feedback loop is stable. Here, w^{-1} represents the trial-delay operator, and d_k is a term that represents disturbances and the effect of initial conditions.

and plants can commonly appear as outputs of higher-order internal models (HOIMs)[1] in the context of robotic manipulators doing different tasks, or freeway traffic models (Hou, Yan, Xu, and Li, 2012).

1.2. Literature review

Linear feedback control encompasses a wide array of problems and their accompanying solutions, such as stabilization, robustness, optimality, sensitivity reduction, fundamental limitations, and design trade-offs. Since the 1990s, there has been an increased effort in the ILC community to generalize the classical problem in these directions. These include the synthesis of 1) robust ILC algorithms (Norrlöf, 2004; Ahn, Moore, and Chen, 2007b; van de Wijdeven, Donkers, and Bosgra, 2009; Bristow, 2010; Moon, Doh, and Chung, 1998; De Roover and Bosgra, 2000; Altın and Barton, 2014), 2) norm-optimal ILC algorithms with quadratic cost functions, 3) adaptive ILC (AILC) methodologies (French, Munde, Rogers, and Owens, 1999; Tayebi, 2006; Tian and Yu, 2003; Wang, Su, and Hong, 2004), along with the study of performance guidelines and design trade-offs (Ahn et al., 2007b; Moore and Lashhab, 2010; Pipeleers and Moore, 2012). See also Bristow et al. (2006), Ahn, Chen, and Moore (2007a), Xu (2011) and the references therein.

Implicit in the vast majority of these earlier works is the invariance assumption in some form. To date, there has been relatively limited material attempting to relax these assumptions. Among these, initial condition invariance was by far the most discussed topic earlier in the literature, since perfect resetting can be hard to achieve for certain systems (Heinzinger, Fenwick, Paden, and Miyazaki, 1992). The central result of Heinzinger et al. (1992) shows that initial condition resetting errors and bounded disturbances affect the tracking error continuously, provided they are uniformly bounded in the iteration domain. The effects of varying disturbance signals have been studied in stochastic settings (Bristow, 2010; Norrlöf, 2004; Ahn et al., 2007b; Saab, 2006). Varying references are also increasingly studied in ILC theory; AILC is one of the avenues in which this objective is pursued (Xu and Xu, 2004; Xu, 2011), while some other works consider parametrizing the set of references by basis functions (Hoelzle, Alleyne, and Johnson, 2011; Bolder and Oomen, 2015; Bolder, Oomen, Koekebakker, and Steinbuch, 2014; van Zundert, Bolder, and Oomen, 2016) or library-based interpolations (Hoelzle and Barton, 2012). Lastly, iteration-varying plant models are actively studied in the case that they can be described by a HOIM (Yin, Xu, and Hou, 2010), with generalizations to iteration-varying references and signals considered in Zhu, Xu, Huang, and Hu (2015).

Despite all these efforts, the feedback interpretation of ILC still paints mostly an incomplete picture, and lacks the fundamental notions of *asymptotic and input-output stability*. In this sense, the introduction of the w-transform (z-transform in the iteration domain)

---

[1] That is, systems wherein the plant operator P_k at trial k is a function of $P_{k-1}, P_{k-2}, \ldots, P_{k-n}$ for some n. However, to the best of our knowledge, there have been no studies on whether HOIMs occur naturally in physical systems.

in Chen and Moore (2002) has been crucial in adopting a more holistic view of ILC as an input-output system, induced by feedback control in the iteration domain. The transform enables the integration of iteration-varying signals into the ILC problem and is a good step towards the establishment of input-output stability properties in ILC. However, it restricts the analysis to iteration-invariant plants and update laws. On the other hand, while Norrlöf and Gunnarsson (2002) presents a framework to investigate the stability of discrete-time iteration-varying systems, the analysis is restricted to iteration-invariant signals. Finally, a robust ILC framework for discrete-time systems in state-space form is analyzed recently in Meng and Moore (2016), Meng and Moore (2014), wherein the treatment is limited to classical D-type ILC algorithms. While the results of these two papers are theoretically important, the authors make no comments on how the learning gain matrices can be designed when the sole information on the uncertainty is boundedness.

Our aim in this paper is to construct *a general framework encapsulating a broad class of systems* in order to, 1) analyze stability properties of ILC in the presence of iteration-varying signals (including references) and plant operators, where the operators are assumed to belong to a bounded set and otherwise unknown, and 2) *connect our analysis to the robust ILC literature* by showing that robust updates lead to stable behavior in ILC. In addition, we will compare the performance of this uncertain iteration-varying system to its nominal invariant counterpart, discuss how nominal performance can be recovered, and verify the theory with simulation examples and experimental implementation.

1.3. Organization of the paper

The remainder of the manuscript is organized as follows: Section 2 introduces preliminaries and the ILC problem. Section 3 proves the basic boundedness result of the algorithm. In Section 4, asymptotic performance and design trade-offs are investigated. Section 5 describes the experimental setup, which also forms the basis for the simulation examples. Simulation examples are presented in Section 6, with the experimental results following in Section 7. Finally, concluding remarks are given in Section 8.

2. Background and problem statement

Consider the classical first-order ILC problem discussed in Section 1. We assume U and Y to be Banach spaces equipped with suitable norms. We base this assumption on the fact that Banach spaces are the natural settings of contraction mapping-based ILC, which relies on the Banach fixed-point theorem. Furthermore, $\mathcal{L}_p$ and l_p spaces, the natural framework for one-dimensional dynamic systems, are complete. The motivation for this assumption is to come up with a general framework that contains the variety of different settings in ILC, consistent with the vector space approach in Moore (1993).

The Banach space framework is discussed further in Appendix A. For simplicity, the reader can assume $\overline{P}$ to be an appropriate real lower triangular (causal) matrix describing a discrete-time linear system, or a stable transfer function $\overline{P}(s)$, without any loss of generality.

2.1. Notation and preliminaries

We take $\mathbb{N}$ to represent the set of nonnegative integers and $\mathbb{N}^+$ the set of positive integers. For normed vector spaces X and V, $B(X, V)$ is the space of all bounded linear operators from X to V. We use $\|.\|$ to denote vector and induced operator norms in the relevant spaces. For a family of operators indexed by a subset of $\mathbb{N}$, the product notation indicates the composition of the operators in increasing order; e.g. $\prod_{i=j}^{k} H_i \triangleq H_k H_{k-1} \ldots H_j$ for $j \leq k$ and $\prod_{i=j}^{k} H_i \triangleq I$ for $j > k$, where I is the identity. The uniform distribution over $[a, b]$ is denoted $\mathcal{U}(a, b)$.

For a rigorous study of the convergence and stability properties of