Monotonic Convergence of Iterative Learning Control Systems with Variable Pass Length

Thomas Seel[a*], Thomas Schauer[a] and Jörg Raisch[a,b]

[a]*Control Systems Group, Technische Universität Berlin, Germany;* [b]*Systems and Control Theory Group, Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany*

(*v5.0 released February 2015*)

A growing number of researchers consider Iterative Learning Control (ILC) a promising tool for numerous control problems in biomedical application systems. We will briefly discuss why classical ILC theory is technically too restrictive for some of these applications. Subsequently, we will extend the classical ILC design in the lifted systems framework to the class of repetitive trajectory tracking tasks with variable pass length. We will analyze the closed-loop dynamics for two standard learning laws, and we will discuss in which sense the tracking error can be reduced by which controller design strategies. Necessary and sufficient conditions for monotonic convergence will be derived. We then summarize all results in a set of practical controller design guidelines. Finally, a simulation study is presented, which demonstrates the usefulness of these guidelines and illustrates the special dynamics that occur in variable pass length learning.

Keywords: Iterative learning control; non-uniform trial duration; lifted system framework; monotonic convergence; biomedical applications;

1.   Introduction and Motivation

Several control problems in various application domains require the controller to repeat the same trajectory tracking task from the same initial condition multiple times. In such control systems, the tracking performance can be improved by learning from previous executions by means of Iterative Learning Control (ILC). Thanks to numerous contributions by many researchers, a rich body of methods exists for various implementations and frameworks of iterative learning control (see for example Moore et al. (1993), Rogers et al. (2007), Galkowski et al. (2003), Owens and Hätönen (2005), Norrlöf and Gunnarsson (2002), Paszke et al. (2010), and many more).

Depending on the analysis framework, the input and output trajectories can be described in the frequency domain or by lifted vectors. In many standard ILC application domains, especially in robotics, the analysis of the system dynamics as well as the controller design are carried out in the frequency domain, rather than in the lifted-systems framework. This yields some advantages[1], but for frequency domain methods to be applicable, the trial duration is assumed to be magnitudes larger than both the sampling time and the characteristic time constant of the system to be controlled. This requirement does not hold in many ILC applications, especially in those from the field of biomedical engineering. Therefore, it is advisable to use the lifted systems framework when aiming to derive new methods for such applications.

A brief review of recent contributions reveals that the potential application fields for ILC in

---

*Corresponding author. Email: seel@control.tu-berlin.de

[1]For example, the frequency-domain representations of the system dynamics and the controller dynamics commute, which is in general not the case for their lifted-system representations.

biomedical engineering are numerous, including diabetes treatment (e.g. Wang et al. (2010)), rehabilitation robotics (e.g. Duschau-Wicke et al. (2010); Liu et al. (2012)), and blood pressure measurement (e.g. Seel et al. (2013a)). However, in these applications, often even the basic assumptions of iterative learning control do not hold. One of these assumptions is that the dynamics (including the initial condition) do not change from pass to pass, i.e. the same input causes the same output when applied again in any of the following trials. This very restrictive assumption has been relaxed by several author's contributions, where ILC methods yielding robustness of convergence properties with respect to iteration-variant dynamics and initial conditions were proposed, see for example Donkers et al. (2008); Moon et al. (1998); Park et al. (2008). However, a more fundamental assumption of ILC is that each trial shall be of the same duration (i.e. the same pass length). In order to demonstrate that this assumption is often violated, two examples from the context of functional electrical stimulation (FES) are briefly discussed in the following:

Functional Electrical Stimulation for Upper Limb Movement.   A functional arm movement shall be produced via electrical stimulation of the respective muscles through skin electrodes, as considered for example by Freeman et al. (2010), Soska et al. (2012) and Klauer et al. (2010). If the desired motion is complex, then very precise and well-timed stimulation patterns need to be applied. However, the system dynamics are very nonlinear and uncertain, since they depend on many details such as the exact electrode position. Therefore, it can be hard to set stimulation parameters that generate (anything close to) the desired motion in the first trial. For successive learning of the corresponding stimulation profile, the use of a robust ILC algorithm based on a linear model seems promising. However, at least for the first passes, it is expected that the actually achieved trajectory leaves the neighborhood of the reference trajectory very early. Even if the ILC algorithm is very robust, the data gathered in a completely different section of the state space is hardly useful. Therefore, and for reasons of safety, a trial should be terminated whenever output and reference begin to differ too much. Nevertheless, between the beginning and the (early) termination of each trial, some valuable data is gathered that might be used for learning if there was a way to deal with varying pass lengths.
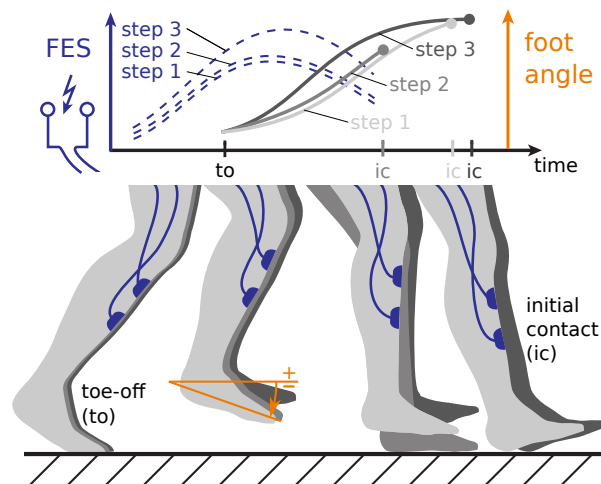

Figure 1.   Iterative Learning Control of Functional Electrical Stimulation (FES) for supporting the foot motion during gait. The stimulation intensity trajectory is adjusted from stride to stride to influence the foot-to-ground angle trajectory. Hip and knee joint are moved by the patients themselves. Depending on their strength and abilities, the steps are often cut short by suddenly putting the foot down. Thus, swing phase duration varies from step to step.

Functional Electrical Stimulation for Gait Assistance.   Many stroke patients suffer from the drop foot syndrome, a limited ability to dorsiflex the foot, i.e. to lift the foot towards the tibia.

Patients walking at constant speed on a treadmill can be supported, for example, by controlling the ankle joint angle via electrical stimulation of the tibialis anterior muscle. This is illustrated in Figure 1 and further described in Seel et al. (2013b). Due to the repetitive nature of gait, ILC is a promising method for improving the stimulation profile for the swing phase from step to step and to maintain the desired foot motion despite slow parameter variance. However, the foot is only in swing phase (i.e. above the ground) as long as the hip joint and the knee joint are sufficiently flexed. Thus, depending on the patient's strength and abilities, the steps are often cut short by suddenly putting the foot down. This means that the swing phase duration varies largely even when patients walk at constant speed. In order to apply ILC methods, we must find a way to learn from trials of varying duration.

Both application examples demonstrate that, in many situations, passes might be terminated early, either by events that depend on the states of the system or on the controller performance, or by randomly occurring events. A number of further examples for this phenomenon are found in the aforementioned literature on other biomedical application systems. Due to a lack of suitable methods, the fact that the trial duration varies is typically ignored completely, or a heuristic approach is used with no guarantee of convergence being maintained (cf. for example Nahrstaedt et al. (2008a), Longman and Mombaur (2006)).

Before we discuss methods that may be used to address the challenge of variable pass length, we briefly isolate the problem we are trying to solve from similar problems that may occur in the same context and applications. In particular, we shall note that there exist applications in which the entire reference duration varies (see for example Moore et al. (2007) or Kurniawan et al. (2013)). More precisely, the reference trajectory that shall be tracked in one trial is scaled to a different time length in another (shorter or longer) trial. This is illustrated by an example from the previously discussed application context. If FES is used to (successfully) perform the same arm motion (for example drawing a large circle or a point-to-point reaching motion) first within five seconds and then within three seconds, then the time scale of the entire tracking problem changes. We are not considering such problems in the following, although they might be an interesting subject for further research. Instead, we focus (as described above) on control problems with a fixed reference time duration but early trial termination.

1.1 *Previous Work on ILC Systems with Variable Pass Length*

Currently, there are very few results available that consider convergence of ILC systems in which the pass length (or trial duration) is not constant. However, a few similar problems are addressed in the literature.

Moore (2000) and Moore and Mathews (1998) consider the problem of controlling a gas-metal arc welding process, in which the time interval between detachments of mass droplets from the end of a consumable electrode is considered a trial. They propose a discontinuous, switching-based iterative learning control scheme that forces the system to exhibit a uniform trial length. However, they assume that the system can be forced to restart a trial by a known input, and they show convergence only for a simple input-update law with a diagonal learning gain.

Li et al. (2014), Li et al. (2015a), and Li et al. (2015b) address iterative learning control design problems for systems with randomly varying trial lengths. A stochastic matrix and an iteration-average operator are introduced to present a unified expression of the ILC scheme. While Li et al. (2014) and Li et al. (2015a) consider discrete-time linear systems, nonlinear affine and non-affine systems are considered in Li et al. (2015b). However, the authors assume the trial length to be a stochastic variable with a known (uniform) probability distribution, and they only show convergence of the expectation of the tracking error to a small value after a large number of trials.

Ahn et al. (2008) consider an ILC scheme that is implemented via a networked control system in which data dropout occurs. They present a Kalman filtering approach and show that it is possible

to design a learning gain such that the system eventually converges to a desired trajectory as long as there is no complete data dropout. With some substantial extensions, this method might be used for variable pass length systems. The authors assume, however, that the data dropout probabilities are known, and they do not establish monotonic convergence of the tracking error in any sense.

In fact, for ILC systems with variable pass lengths, the practically highly relevant property of monotonic convergence is, to the best of our knowledge, for the first time considered by Seel et al. (2011). Under the sole assumption of a (potentially very small) lower bound and a (potentially very large) upper bound of the trial duration, it is shown that monotonic convergence of the tracking error can be achieved in different senses by applying a slightly modified version of a standard input-update law. Guth et al. (2013) apply these results to trajectory tracking on a lab-scale gantry crane. In both contributions, however, the convergence analysis is restricted to simple learning laws in which the change of the input trajectory equals the tracking error multiplied by a learning gain matrix. Therefore, the methods proposed therein do not apply to the commonly used iterative learning controllers that include a Q-filter.

From this review of previous contributions we conclude that there exist no methods that guarantee monotonic convergence in iterative learning control systems with pass lenghts that vary arbitrarily within certain bounds. In the present contribution, a number of novel methods are derived to fill this gap. We will extend the concept of monotonic convergence as well as the standard ILC design in the lifted-system framework to the class of repetitive trajectory tracking tasks with variable pass length. The methods we propose for this class of systems are summarized in a set of controller design guidelines, and their usefulness for practical applications is discussed in a simulation study.

The remainder of this contribution is organized as follows. In Section 2 we define the control problem and derive the closed-loop dynamics. With regard to controller design, we consider ILC systems with a single learning gain matrix in Section 3, whereas the extended version with lowpass filter in the update loop is considered in Section 4. For both cases, we will analyze the closed loop dynamics, and we will discuss in which sense the tracking error can be reduced by which controller design strategies. In particular, the maximum-pass-length (MPL) error will be introduced as a useful concept for convergence analysis of variable-pass-length systems. Necessary and sufficient conditions for monotonic convergence of different vector norms of this error will be derived. In Section 5, we will then summarize all results in a set of practical controller design guidelines. Finally, the special dynamics that occur in variable-pass-length learning will be illustrated by a simulation study in Section 6.

2.   Iterative Learning Control for Variable-Pass-Length Systems

In this section, a class of iterative learning control systems with variable pass length is introduced and its closed-loop dynamics are derived. To simplify notation, we consider a dynamic system with scalar input $u(t) \in \mathbb{R}$ and scalar output $y(t) \in \mathbb{R}$. It is assumed, that the pass length n can vary from pass to pass and that pass termination is sudden. This means that, for each pass $j \in \mathbb{N}_0$, the pass length n_j is not known before the pass ends. However, a possibly very small lower bound $\underline{n}$ and a possibly very large upper bound $\overline{n}$ shall be given. Consequently, a full-length desired output $\mathbf{r} \in \mathbb{R}^{\overline{n}}$ must be defined, and (for each pass j) a full-length input $\bar{\mathbf{u}}_j \in \mathbb{R}^{\overline{n}}$ with

$$\bar{\mathbf{u}}_j := [u(t_{0,j}), u(t_{0,j} + t_{\mathrm{s}}), \ldots, u(t_{0,j} + (\overline{n} - 1)t_{\mathrm{s}})]^T, \tag{1}$$

must be prepared, although only the first n_j samples of that trajectory will actually be applied to the system. In contrast, the output trajectory that is measured during the same trial is denoted by $\mathbf{y}_j \in \mathbb{R}^{n_j}$ and contains only $n_j \leq \overline{n}$ sample values.

In order to cope with the varying length of the output vector and for the sake of a more compact

notation, we introduce the *remove-last-$(\overline{n}-n_j)$-elements* operator $\langle \cdot \rangle_{n_j} : \mathbb{R}^{\overline{n}} \to \mathbb{R}^{n_j}$ and the *append-$(\overline{n}-n_j)$-zero-elements* operator $\langle \cdot \rangle_{\overline{n}} : \mathbb{R}^{n_j} \to \mathbb{R}^{\overline{n}}$. Clearly, these correspond to premultiplication by the matrices $[\mathbf{I}_{n_j}, 0_{n_j \times (\overline{n}-n_j)}]$ and $[\mathbf{I}_{n_j}, 0_{n_j \times (\overline{n}-n_j)}]^T$, respectively, where $\mathbf{I}_{n_j}$ denotes the n_j-dimensional identity matrix and $0_{n_j \times (\overline{n}-n_j)}$ denotes a zero matrix with n_j rows and $(\overline{n}-n_j)$ columns. Thus, for example, the desired output signal for pass length n_j becomes $\langle \mathbf{r} \rangle_{n_j}$, and the tracking error $\mathbf{e}_j \in \mathbb{R}^{n_j}$ measured in pass j can be denoted by

$$\mathbf{e}_j := \langle \mathbf{r} \rangle_{n_j} - \mathbf{y}_j . \tag{2}$$

For the last $(\overline{n}-n_j)$ samples of the reference trajectory, there exists no measurement information, since the trial actually ends after n_j sample periods. This means that the complete input trajectory $\overline{\mathbf{u}}_j$ must be updated using the incomplete error information $\mathbf{e}_j$. This can be achieved, for example, by using the following modified version of a classical ILC learning law:

$$\overline{\mathbf{u}}_{j+1} = \overline{\mathbf{u}}_j + \mathbf{L} \langle \mathbf{e}_j \rangle_{\overline{n}}, \tag{3}$$

where $\mathbf{L} \in \mathbb{R}^{\overline{n} \times \overline{n}}$ is an adjustable learning gain matrix. As in classical ILC systems, an initial input $\overline{\mathbf{u}}_0$ must be chosen based on the (limited) prior knowledge of the system dynamics. Every following input is then calculated using (3). Although many of the following arguments would also hold for a learning gain matrix $\mathbf{L}$ that depends on the trial index, we will consider only the standard case of an iteration-invariant learning gain in the sequel.

Rationale of the Variable-Pass-Length Learning Law. Recall from the standard ILC literature that, whenever certain sections of the measured output trajectory perfectly track the corresponding sections of the reference trajectory, then a well-designed learning controller will try to modify the input trajectory in such a way that these sections of the output trajectory are influenced as little as possible. By setting the last $\overline{n} - n_j$ samples of the error information vector to zero, we likewise demand a learning step that affects the last $\overline{n} - n_j$ samples of the output as little as possible. This is a very reasonable strategy, since there is no measurement information about the values that these samples would have under the current input trajectory.

In the following, we will introduce a linear system model and analyze the resulting closed-loop dynamics in the presence of variable pass length.

2.1 Linear System Dynamics

Assume that the input-output dynamics of the system to be controlled can be approximated by some linear, discrete-time process with Markov parameters $p_i, i = 1, 2, \ldots$. The output $\overline{\mathbf{y}}_j \in \mathbb{R}^{\overline{n}}$ of the system can then be described by

$$\overline{\mathbf{y}}_j := \begin{bmatrix} \mathbf{y}_j \\ \hat{\mathbf{y}}_j \end{bmatrix} := \mathbf{P}\overline{\mathbf{u}}_j + \mathbf{v} , \quad \mathbf{P} = \begin{bmatrix} p_1 & 0 & \cdots & 0 \\ p_2 & p_1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ p_{\bar{n}} & p_{\bar{n}-1} & \cdots & p_1 \end{bmatrix} , \tag{4}$$

where $\mathbf{P} \in \mathbb{R}^{\overline{n} \times \overline{n}}$ is the lifted-system matrix of the process and $\mathbf{v} \in \mathbb{R}^{\overline{n}}$ is a bounded, iteration-invariant disturbance.

Assume that the system dynamics $\mathbf{P}, \mathbf{v}$ are only imprecisely known, and recall that the trials end after n_j samples. The last $\overline{n} - n_j$ samples of $\overline{\mathbf{y}}_j$ are denoted by $\hat{\mathbf{y}}_j$ and represent the unknown output that would (hypothetically) be measured if the trial was of maximum length $n_j = \overline{n}$. This

portion of the complete output $\bar{\mathbf{y}}_j$ is not known to the controller when computing $\bar{\mathbf{u}}_{j+1}$. In fact, the only certain information the controller obtains in each iteration is the measured output $\mathbf{y}_j$.

2.2   Derivation of Closed-Loop Dynamics

Although we assume that the true system dynamics are only imprecisely known, we can nevertheless derive equations that describe how the complete output $\bar{\mathbf{y}}_j$ and the tracking error evolve for a particular $\mathbf{P}$ under consideration. By the latter, we will develop a method that enables us to verify, for a (large number of) given plant dynamics, whether a certain controller yields desired convergence properties.

In order to derive closed-loop dynamics, we define the maximum-pass-length (MPL) error $\bar{\mathbf{e}}_j \in \mathbb{R}^{\bar{n}}$ analogously to the previous definition of the maximum-pass-length output:

$$\bar{\mathbf{e}}_j := \begin{bmatrix} \mathbf{e}_j \\ \hat{\mathbf{e}}_j \end{bmatrix} := \mathbf{r} - \bar{\mathbf{y}}_j. \tag{5}$$

Therein, $\mathbf{e}_j \in \mathbb{R}^{n_j}$ is the actually measured control deviation, and $\hat{\mathbf{e}}_j \in \mathbb{R}^{(\bar{n}-n_j)}$ is the hypothetical error that would occur on the last $\bar{n} - n_j$ sample instants if the trial was of full length.

The previous definitions allow us to rewrite the update law (3) in the following form:

$$\bar{\mathbf{u}}_{j+1} \,=\, \bar{\mathbf{u}}_j \,+\, \mathbf{L} \, \langle\, \langle \mathbf{r} \rangle_{n_j} - \mathbf{y}_j \,\rangle_{\bar{n}} \,=\, \bar{\mathbf{u}}_j \,+\, \mathbf{L}\, \mathbf{H}_{n_j}\, \bar{\mathbf{e}}_j \,, \tag{6}$$

in which the last $\bar{n} - n_j$ samples of the MPL error are canceled by multiplication with the block-diagonal matrix[2] $\mathbf{H}_{n_j} := \mathrm{blockdiag}\{\mathbf{I}_{n_j}, \mathbf{0}_{(\bar{n}-n_j)\times(\bar{n}-n_j)}\}$.

By combining (4), (5), and (6), and under the assumption of a regular system matrix $\mathbf{P}$, the following MPL error dynamics are derived:

$$\bar{\mathbf{e}}_{j+1} \,=\, \left(\mathbf{I}_{\bar{n}} - \mathbf{P}\mathbf{L}\mathbf{H}_{n_j} \right) \bar{\mathbf{e}}_j. \tag{7}$$

In order to gain more insight into the interrelation of the measured error $\mathbf{e}_j$ and the hypothetical error $\hat{\mathbf{e}}_j$, we introduce the following trial-dependent partition:

$$\begin{bmatrix} \mathbf{G}_{1,j} & \mathbf{G}_{2,j} \\ \mathbf{G}_{3,j} & \mathbf{G}_{4,j} \end{bmatrix} := \mathbf{I}_{\bar{n}} - \mathbf{P}\mathbf{L}, \tag{8}$$

$$\mathbf{G}_{1,j} \in \mathbb{R}^{n_j \times n_j}, \; \mathbf{G}_{2,j} \in \mathbb{R}^{n_j \times (\bar{n}-n_j)}, \; \mathbf{G}_{3,j} \in \mathbb{R}^{(\bar{n}-n_j)\times n_j}, \; \mathbf{G}_{4,j} \in \mathbb{R}^{(\bar{n}-n_j)\times(\bar{n}-n_j)}. \tag{9}$$

We then use (5) to rewrite the closed-loop dynamics as follows:

$$\begin{bmatrix} \mathbf{e}_{j+1} \\ \hat{\mathbf{e}}_{j+1} \end{bmatrix} = \begin{bmatrix} \mathbf{G}_{1,j} & \mathbf{0}_{n_j \times (\bar{n}-n_j)} \\ \mathbf{G}_{3,j} & \mathbf{I}_{\bar{n}-n_j} \end{bmatrix} \begin{bmatrix} \mathbf{e}_j \\ \hat{\mathbf{e}}_j \end{bmatrix} = \left(\begin{bmatrix} \mathbf{G}_{1,j} & \mathbf{G}_{2,j} \\ \mathbf{G}_{3,j} & \mathbf{G}_{4,j} \end{bmatrix} \begin{bmatrix} \mathbf{e}_j \\ \mathbf{0}_{(\bar{n}-n_j)\times 1} \end{bmatrix} + \begin{bmatrix} \mathbf{0}_{n_j \times 1} \\ \hat{\mathbf{e}}_j \end{bmatrix} \right). \tag{10}$$

Note that, in general, $\mathbf{e}_{j+1}$ and $\mathbf{e}_j$ do not have the same dimensions.

In the following sections, we will discuss in which sense the tracking error may reduce in variable-pass-length systems, we will investigate how the learning gain matrix $\mathbf{L}$ must be chosen to guarantee such a reduction, and we will discuss some fundamental effects that appear in learning from trials with variable length.

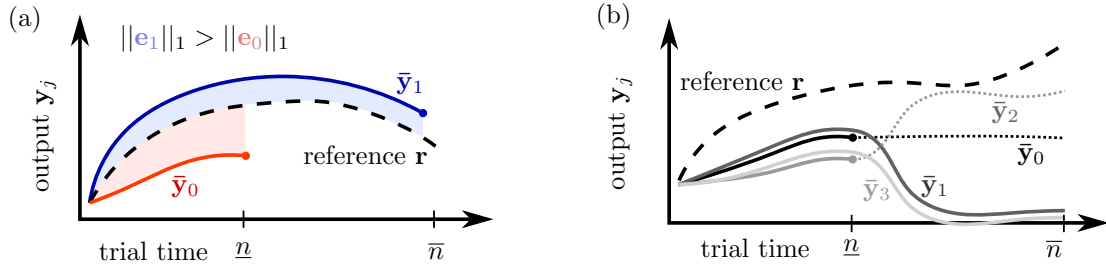[2]By definition, $\mathbf{H}_{n_j} a = \langle \langle a \rangle_{n_j} \rangle_{\bar{n}} \; \forall a \in \mathbb{R}^{\bar{n}}$

Figure 2. *Left*: In variable pass length ILC, the norm of the measured error $\mathbf{e}_j$ is not a useful performance measure, since it depends largely on the trial duration. The longer trial 1 might yield a larger error norm than the short trial 0, despite the former exhibiting better reference tracking than the latter. *Right*: Hidden error growth in ILC systems with variable pass length. Every second trial is cut short, dotted lines indicate the hypothetical output $\hat{\mathbf{y}}_j$. In each trial, the tracking error is reduced in the sense of (11). Nevertheless, the tracking gets worse on the entire time interval.

3. Learning Dynamics for the Standard Input-Update Law

In classical ILC systems, controller design and implementation aims at reducing the p-norm[3] $||\mathbf{e}_j||_p$ of the measured error from one trial to the next, for some $p \in \{1, 2, \infty\}$. In variable-pass-length ILC, the p-norm of the measured tracking error $\mathbf{e}_j$ will vary as n_j varies, even if the input trajectory is not updated at all. A long trial might yield a larger error norm than a short trial, despite the former yielding better reference tracking than the latter, as illustrated in Figure 2a.

Hence, we need a different concept to describe whether the reference tracking is improved by an update step in the presence of variable pass length. Since, in each trial j, the controller receives the measurement values of the first n_j samples, it is reasonable to demand that the reference tracking should be improved at least on those first n_j samples. Since every vector p-norm and its induced[4] matrix norm are submultiplicative[5] and since the induced p-norm of a matrix cannot be smaller than the induced p-norm of any of its sub-matrices, we can derive the following condition for this property from (10):

$$\frac{||\langle \bar{\mathbf{e}}_{j+1} \rangle_{n_j}||_p}{||\mathbf{e}_j||_p} < 1 \; \forall j \in \mathbb{N}_0, n_j \in \{\underline{n}, ..., \overline{n}\}, \mathbf{r}, \mathbf{v}, \bar{\mathbf{u}}_0 \in \mathbb{R}^{\overline{n}} \quad \Leftrightarrow \quad \gamma_p := ||\mathbf{I}_{\overline{n}} - \mathbf{PL}||_p < 1. \qquad (11)$$

This statement introduces the convergence indicator γ_p and represent a both simple and useful result with strong analogy to the classical case of constant pass length. However, the demanded property is still not sufficient to guarantee that the tracking error remains bounded in the presence of variable pass length. Figure 2b illustrates this remarkable fact. While the example dynamics chosen therein may seem unrealistic or extreme, they nevertheless demonstrate that, in each trial, the performance on the first n_j samples might be improved at the cost of (severely) worsening the performance on the last $\overline{n} - n_j$ samples. Hence, a more rigorous converge concept is needed that also accounts for hidden growth of the hypothetical tracking error $\hat{\mathbf{e}}_j$.

3.1 *Monotonic Convergence of the Maximum-Pass-Length Error*

The convergence properties of $\bar{\mathbf{e}}_j$ well describe whether the controller performance actually improves. Therefore, we aim at designing controllers that ensure reduction of a p-norm $||\bar{\mathbf{e}}_j||_p$ of the MPL error, for example its sum-of-absolute-values norm ($p = 1$), its Euclidean norm ($p = 2$), or its maximum-absolute-value norm ($p = \infty$). Obviously, if $||\bar{\mathbf{e}}_j||_p$ is monotonically decreasing in j,

---

[3]$||\mathbf{x}||_p := (\sum_{i=1}^{n} |\mathbf{x}(i)|^p)^{(1/p)} \; \forall \mathbf{x} \in \mathbb{R}^n, p \geq 1$, with the limit case $||\mathbf{x}||_\infty = \max_i |\mathbf{x}(i)| \; \forall \mathbf{x} \in \mathbb{R}^n$

[4]i.e. $||\mathbf{A}||_p := \max\{||\mathbf{Ax}||_p, ||\mathbf{x}||_p = 1\}$

[5]i.e. $||\mathbf{ABx}||_p \leq ||\mathbf{A}||_p ||\mathbf{B}||_p ||\mathbf{x}||_p \; \forall \mathbf{A}, \mathbf{B}, \mathbf{x}$

then also $||\mathbf{e}_j||_p$ is bounded by trial-wise decreasing upper bounds, since

$$||\mathbf{e}_j||_p \leq ||\bar{\mathbf{e}}_j||_p \, \forall j \in \mathbb{N}_0 \, \forall p \in \{1, 2, \infty\}. \tag{12}$$

Therefore, the concept of the MPL error is very useful for convergence analysis. In order to have monotonic convergence (MC) for an arbitrary sequence of pass lengths $\{n_j\}$ and an arbitrary disturbance signal $\mathbf{v}$, the following condition must hold:

Lemma 1 (Monotonic convergence of the MPL error p-norm (Seel et al., 2011)): *For a repetitive process* (4) *with regular system matrix* $\mathbf{P}$ *and input-update law* (6),

$$||\bar{\mathbf{e}}_{j+1}||_p \leq ||\bar{\mathbf{e}}_j||_p \quad \forall j \in \mathbb{N}_0, n_j \in \{\underline{n}, ..., \overline{n}\}, \mathbf{r}, \mathbf{v}, \bar{\mathbf{u}}_0 \in \mathbb{R}^{\overline{n}}$$
$$\Leftrightarrow \quad ||\mathbf{I}_{\overline{n}} - \mathbf{PLH}_n||_p \leq 1 \qquad \forall \, n \in \{\underline{n}, ..., \overline{n}\} \tag{13}$$

for some vector norm $|| \cdot ||_p$ *and its induced matrix norm.*

Proof. Necessity and sufficiency follow from the definition of an induced matrix norm, i.e. $||\mathbf{A}||_p := \max\{||\mathbf{Ax}||_p, ||\mathbf{x}||_p = 1\}$, from the MPL error closed-loop dynamics (7), and from the fact that (in each trial j) the pass length may take any value $n_j \in \{\underline{n}, ..., \overline{n}\}$. □

In practice, checking the condition of Lemma 1 would entail the need to calculate the norm of $(\overline{n} - \underline{n} + 1)$ matrices of dimension $\overline{n} \times \overline{n}$, which is a potentially large number of large-scale matrices when high sampling rates are used. Fortunately, simplified conditions can be found for the 1-norm $||\mathbf{e}_j||_1$ by exploiting the special structure of the closed-loop error dynamics matrices $(\mathbf{I}_{\overline{n}} - \mathbf{PLH}_n), n \in \{\underline{n}, ..., \overline{n}\}$.

The following theorem exploits this special structure and provides a necessary and sufficient condition for monotonic convergence of the 1-norm of the MPL error.

Theorem 1 (Monotonic convergence of the MPL error 1-norm (Seel et al., 2011)): *Given a repetitive process* (4) *with regular* $\mathbf{P}$ *and input-update law* (6), *the 1-norm of the MPL error* $\bar{\mathbf{e}}_j$ *is monotonically decreasing, i.e.*

$$||\bar{\mathbf{e}}_{j+1}||_1 \leq ||\bar{\mathbf{e}}_j||_1 \, \forall j \in \mathbb{N}_0, n_j \in \{\underline{n}, ..., \overline{n}\}, \mathbf{r}, \mathbf{v}, \bar{\mathbf{u}}_0 \in \mathbb{R}^{\overline{n}}, \tag{14}$$

if and only if

$$\gamma_1 = ||\mathbf{I}_{\overline{n}} - \mathbf{PL}||_1 \leq 1 . \tag{15}$$

Proof. Necessity follows from Lemma 1, since (15) repeats a special case ($n = \overline{n}$, $p = 1$) of the condition therein. Sufficiency is proved by the following argument, which holds for every $n \in \{\underline{n}, ..., \overline{n}\}$:

$$(\mathbf{I}_{\overline{n}} - \mathbf{PLH}_n) = (\mathbf{I}_{\overline{n}} - \mathbf{PL})\mathbf{H}_n + (\mathbf{I}_{\overline{n}} - \mathbf{H}_n). \tag{16}$$

Note that the last $\overline{n} - n$ columns of $(\mathbf{I}_{\overline{n}} - \mathbf{PL})\mathbf{H}_n$ are zero and that $(\mathbf{I}_{\overline{n}} - \mathbf{H}_n) = \mathrm{blockdiag}\{\mathbf{0}_{n \times n}, \mathbf{I}_{(\overline{n}-n)}\}$. Recall that the induced matrix norm $|| \cdot ||_1$ is the maximum absolute column sum of the matrix. Therefore, we obtain

$$||\mathbf{I}_{\overline{n}} - \mathbf{PLH}_n||_1 = \max(||(\mathbf{I}_{\overline{n}} - \mathbf{PL})\mathbf{H}_n||_1, 1) \tag{17}$$
$$\leq \max(||\mathbf{I}_{\overline{n}} - \mathbf{PL}||_1 ||\mathbf{H}_n||_1, 1) \leq \max(\gamma_1, 1). \tag{18}$$

And with Lemma 1, monotonic convergence follows. □

Of course, *strict* monotonic convergence is more desirable than just monotonic convergence. But even for $\gamma_1 < 1$, it turns out that $||\mathbf{I}_{\overline{n}} - \mathbf{PLH}_n||_1 = 1 \ \forall n < \overline{n}$, since the last $\overline{n} - n$ columns each have an absolute-value sum of one. To clarify this issue, the following corollary is given:

Corollary 1 (Almost-strict monotonic convergence of the MPL error 1-norm): *If the matrix* $\mathbf{L}$ *satisfies* $\gamma_1 = ||\mathbf{I}_{\overline{n}} - \mathbf{PL}||_1 < 1$, *then* $||\bar{\mathbf{e}}_{j+1}||_1 < ||\bar{\mathbf{e}}_j||_1$ *for arbitrary pass length* $n_j \in \{\underline{n}, ..., \overline{n}\}$, *unless the first* n_j *entries of* $\bar{\mathbf{e}}_j$ *are zero. This follows from* (10) *and from the special property* $||\bar{\mathbf{e}}_j||_1 = ||\mathbf{e}_j||_1 + ||\hat{\mathbf{e}}_j||_1 \forall j$ *of the 1-norm, which allows the following conclusion:*

$$||\bar{\mathbf{e}}_{j+1}||_1 \leq ||(\mathbf{I}_{\overline{n}} - \mathbf{PL})||_1 ||\mathbf{e}_j||_1 + ||\hat{\mathbf{e}}_j||_1 < ||\bar{\mathbf{e}}_j||_1 \,. \tag{19}$$

In practice, this simply means that the learning algorithm is not guaranteed to improve the last $\overline{n} - n_j$ input samples until long enough passes occur. But it always learns from measured errors on the first n_j samples and (unless they are all zero) improves the overall controller performance by reducing the 1-norm of the MPL error, even for short pass lengths.

3.1.1 Convergence of the MPL error in different p-norms

The simplified condition that Theorem 1 provides for monotonic convergence of the MPL error 1-norm uses a single scalar convergence indicator, which is quickly calculated from the given input-output dynamics $\mathbf{P}$ and the iterative learning controller $\mathbf{L}$. Beyond this, Lemma 1 provides a condition for monotonic convergence of the MPL error in any vector p-norm. However, as mentioned before, it requires the calculation of the matrix norm of a (potentially) large number of (potentially) large-scale matrices. Therefore, we shall briefly comment on how the simplified condition of Theorem 1 can be used to assess convergence in any other p-norm of the MPL error. To this end, note that any vector p-norm is less than or equal to the 1-norm of the same vector, i.e. $||\mathbf{a}||_p \leq ||\mathbf{a}||_1 \ \forall \mathbf{a} \in \mathbb{R}^n, p \in \{1, 2, ..., \infty\}$. Therefore, note the following properties.

P1: If the MPL error converges to zero in the 1-norm, then so it does in any p-norm.
P2: If its 1-norm converges to a small value ε then any p-norm of the MPL error converges to or even falls below ε.
P3: If almost-strict MC is established for the MPL error 1-norm, then any p-norm is bounded from above by the monotonically nonincreasing bound $||\bar{\mathbf{e}}_j||_1$, i.e. $||\bar{\mathbf{e}}_j||_p \leq ||\bar{\mathbf{e}}_j||_1 \ \forall j$.

Therefore, it is in general sufficient to design a learning gain matrix $\mathbf{L}$ that satisfies (15).

4. Learning Dynamics for Input-Update Laws with Q-Filter

In the previous section, a modified version of the standard learning law $\mathbf{u}_{j+1} = \mathbf{u}_j + \mathbf{Le}_j$ has been employed. This is a reasonable choice for many applications. Often, however, one wants to restrict the learning process to a low-frequency range in which model uncertainties are at least limited. In order to avoid that high-frequency components accumulate in the input trajectory via repetitive application of the update law, the aforementioned standard learning law is often extended by a lowpass filter in standard (i.e. constant-pass-length) ILC. Thus, in analogy, we propose the following for the variable-pass-length case:

$$\bar{\mathbf{u}}_{j+1} = \mathbf{Q}\left(\bar{\mathbf{u}}_j + \mathbf{L}\,\mathbf{H}_{n_j}\,\bar{\mathbf{e}}_j\right) \,, \tag{20}$$

where $\mathbf{Q} \in \mathbb{R}^{\overline{n} \times \overline{n}}$ is a symmetric matrix with Toeplitz structure containing the Markov parameters of a low pass filter with cutoff frequency denoted by $f_{\mathbf{Q}}$. Multiplying a lifted signal by $\mathbf{Q}$ corresponds to applying a zero-phase lowpass filter to the signal, see for example Elci et al. (2002).

Using this modified learning law leads to more complex closed-loop dynamics. By combining (4), (5), and (20), and under the assumption of a regular system matrix $\mathbf{P}$, we obtain the following maximum-pass-length error dynamics:

$$\begin{bmatrix} \mathbf{e}_{j+1} \\ \hat{\mathbf{e}}_{j+1} \end{bmatrix} = \mathbf{PQP}^{-1}(\mathbf{I}_{\bar{n}} - \mathbf{PLH}_{n_j}) \begin{bmatrix} \mathbf{e}_j \\ \hat{\mathbf{e}}_j \end{bmatrix} + (\mathbf{I}_{\bar{n}} - \mathbf{PQP}^{-1})(\mathbf{r} - \mathbf{v}). \tag{21}$$

As before, note that, in general, the measured error $\mathbf{e}_{j+1}$ of trial $j + 1$ on the left-hand side and the measured error $\mathbf{e}_j$ of trial j on the right-hand side do not have the same dimensions. For convenience, introduce the following trial-dependent partition:

$$\begin{bmatrix} \mathbf{K}_{1,j} & \mathbf{K}_{2,j} \\ \mathbf{K}_{3,j} & \mathbf{K}_{4,j} \end{bmatrix} := \mathbf{PQP}^{-1}, \tag{22}$$

$$\mathbf{K}_{1,j} \in \mathbb{R}^{n_j \times n_j}, \ \mathbf{K}_{2,j} \in \mathbb{R}^{n_j \times (\bar{n}-n_j)}, \ \mathbf{K}_{3,j} \in \mathbb{R}^{(\bar{n}-n_j) \times n_j}, \ \mathbf{K}_{4,j} \in \mathbb{R}^{(\bar{n}-n_j) \times (\bar{n}-n_j)}. \tag{23}$$

This, together with (8), allows us to rewrite (21) in the following form:

$$\begin{bmatrix} \mathbf{e}_{j+1} \\ \hat{\mathbf{e}}_{j+1} \end{bmatrix} = \begin{bmatrix} \mathbf{K}_{1,j} & \mathbf{K}_{2,j} \\ \mathbf{K}_{3,j} & \mathbf{K}_{4,j} \end{bmatrix} \left(\begin{bmatrix} \mathbf{G}_{1,j} & \mathbf{G}_{2,j} \\ \mathbf{G}_{3,j} & \mathbf{G}_{4,j} \end{bmatrix} \begin{bmatrix} \mathbf{e}_j \\ \mathbf{0}_{(\bar{n}-n_j) \times 1} \end{bmatrix} + \begin{bmatrix} \mathbf{0}_{n_j \times 1} \\ \hat{\mathbf{e}}_j \end{bmatrix} \right) + \begin{bmatrix} \mathbf{I}_{n_j} - \mathbf{K}_{1,j} & -\mathbf{K}_{2,j} \\ -\mathbf{K}_{3,j} & \mathbf{I}_{n_j} - \mathbf{K}_{4,j} \end{bmatrix} (\mathbf{r} - \mathbf{v}). \tag{24}$$

These dynamics exhibit a different structure than the previously considered dynamics (7). Lemma 1 and Theorem 1 cannot be applied or trivially extended to the present case.

To assess the right-hand side of (24), recall that any p-norm is submultiplicative and that the p-norm of a matrix cannot be smaller than the p-norm of any of its sub-matrices. It is helpful to introduce the following scalar convergence indicator $\tilde{\gamma}_p$, residual indicator $\tilde{\epsilon}_p$, and crosstalk indicator $\tilde{\kappa}_p$:

$$\tilde{\gamma}_p := ||\mathbf{PQP}^{-1}(\mathbf{I}_{\bar{n}} - \mathbf{PL})||_p \quad \text{with} \quad \tilde{\gamma}_p \geq \left\| \mathbf{PQP}^{-1} \begin{bmatrix} \mathbf{G}_{1,j} \\ \mathbf{G}_{3,j} \end{bmatrix} \right\|_p, \tag{25}$$

$$\tilde{\epsilon}_p := ||\mathbf{I}_{\bar{n}} - \mathbf{PQP}^{-1}||_p, \tag{26}$$

$$\tilde{\kappa}_p := ||\mathbf{PQP}^{-1}||_p \quad \text{with} \quad \tilde{\kappa}_p \geq \left\| \begin{bmatrix} \mathbf{K}_{2,j} \\ \mathbf{K}_{4,j} \end{bmatrix} \right\|_p. \tag{27}$$

Note that these indicators do *neither* depend on the current pass length, *nor* on the reference or disturbance. They are constants that can easily be calculated for any given system $\mathbf{P}$ and iterative learning controller $\mathbf{L}, \mathbf{Q}$. Therefore, we will try to formulate convergence criteria and controller design guidelines in terms of these indicators.

The above definitions allow us to state the following result:

Theorem 2 (Monotonic decrease of the MPL error norm above a threshold): *Consider a repetitive process* (4) *with regular* $\mathbf{P}$ *and input-update law* (20). *Assume that the learning gain matrix* $\mathbf{L}$ *is chosen such that* $\tilde{\gamma}_p < 1$. *Then the p-norm of the MPL error decreases from trial to trial, i.e.* $||\bar{\mathbf{e}}_{j+1}||_p < ||\bar{\mathbf{e}}_j||_p$, *if the norm of the measured tracking error is sufficiently large, namely*

$$||\mathbf{e}_j||_p > \frac{\tilde{\epsilon}_p ||\mathbf{r} - \mathbf{v}||_p + \tilde{\kappa}_p ||\hat{\mathbf{e}}_j||_p}{1 - \tilde{\gamma}_p}. \tag{28}$$

Proof. From the closed-loop error dynamics (24) and the triangle inequality, it follows that

$$\frac{||\bar{\mathbf{e}}_{j+1}||_p}{||\bar{\mathbf{e}}_j||_p} \leq \frac{||\bar{\mathbf{e}}_{j+1}||_p}{||\mathbf{e}_j||_p} \leq \frac{\left|\left|\mathbf{PQP}^{-1}\begin{bmatrix}\mathbf{G}_{1,j}\\\mathbf{G}_{3,j}\end{bmatrix}\mathbf{e}_j\right|\right|_p + \left|\left|\begin{bmatrix}\mathbf{K}_{2,j}\\\mathbf{K}_{4,j}\end{bmatrix}\hat{\mathbf{e}}_j\right|\right|_p + \tilde{\epsilon}_p||\mathbf{r}-\mathbf{v}||_p}{||\mathbf{e}_j||_p}\,. \tag{29}$$

Using the inequalities (25), (27) and $||\mathbf{e}_j||_p > \tilde{\epsilon}_p||\mathbf{r}-\mathbf{v}||_p + \tilde{\kappa}_p||\hat{\mathbf{e}}_j||_p + \tilde{\gamma}_p||\mathbf{e}_j||_p$ from (28), we immediately see that the right-hand side of (29) is less than one. $\square$

This means that the MPL error norm $||\bar{\mathbf{e}}_j||_p$ decreases (at least) until the norm $||\mathbf{e}_j||_p$ of the measured tracking error falls below a certain threshold, which is small if the convergence indicator $\tilde{\gamma}_p$, the residual indicator $\tilde{\epsilon}_p$, and the hypothetical error norm $||\hat{\mathbf{e}}_j||_p$ are small. The usefulness of such statements largely depends on the conservativeness of the threshold (28). Therefore, we will aim at reducing this conservativeness in the following.

4.1   *Frequency-Weighted Residual Indicator*

Assume that a nontrivial Q-filter is used, which is neither $\mathbf{Q} \approx \mathbf{0}_{\bar{n}\times\bar{n}}$ nor $\mathbf{Q} \approx \mathbf{I}_{\bar{n}}$, i.e. the cutoff frequency $f_{\mathbf{Q}}$ is neither close to zero nor close to the Nyquist frequency. Then, for all regular lifted system matrices $\mathbf{P}$, we find that $\tilde{\kappa}_p = ||\mathbf{PQP}^{-1}||_p$ can hardly be less than one. This is not surprising, since multiplication by a lowpass Q-filter leaves low-frequency lifted signal almost unchanged. At the same time, premultiplying a high-frequency lifted signal by $\mathbf{Q}$ yields a lifted signal with very small norm. Therefore, the residual indicator $\tilde{\epsilon}_p = ||(\mathbf{I}_{\bar{n}} - \mathbf{PQP}^{-1})||_p$ can hardly be smaller than one.

Since this seems to be a very restrictive result, we briefly interpret it in more detail by inspecting the system dynamics (4) and the learning law (20), which leads to the following conclusion. If there are high-frequency portions in the reference $\mathbf{r}$ or disturbance $\mathbf{v}$, then these lead to high-frequency portions in the tracking error that will not be removed by the learning controller unless $f_{\mathbf{Q}}$ is larger than those frequencies. This is in accordance with our previous statement that the Q-filter restricts the learning process to a low-frequency range.

In that sense, however, approximating $||(\mathbf{I}_{\bar{n}} - \mathbf{PQP}^{-1})(\mathbf{r}-\mathbf{v})||_p$ by $\tilde{\epsilon}_p||\mathbf{r}-\mathbf{v}||_p$ is highly conservative. In the very realistic case that both $\mathbf{r}$ and $\mathbf{v}$ contain only frequencies (far) below the chosen $f_{\mathbf{Q}}$, we may expect the error norm to be reduced to a much smaller value than (28) suggests.

In order to obtain a less conservative result, we define the frequency-weighted matrix norm $||\cdot||_{\leq f_0, p} : \mathbb{R}^{\bar{n}\times\bar{n}} \to \mathbb{R}$ as follows:

$$||\mathbf{A}||_{\leq f_0, p} := ||\mathbf{A}\mathbf{Q}_{f_0}||_p \,\forall\mathbf{A} \in \mathbb{R}^{\bar{n}\times\bar{n}}\,, \tag{30}$$

where $\mathbf{Q}_{f_0} \in \mathbb{R}^{\bar{n}\times\bar{n}}$ is the regular lifted system matrix of a lowpass filter with some cutoff frequency f_0 (between 0 and the Nyquist frequency) and unity p-norm $||\mathbf{Q}_{f_0}||_p = 1$. Recall that the induced matrix p-norm is the largest gain by which the p-norm of a vector increases when multiplied with the matrix. Following this definition, the frequency-weighted matrix norm $||\cdot||_{\leq f_0, p}$ yields the largest gain by which the p-norm of a low-frequency vector increases when multiplied with the matrix[6]. Furthermore, note that, by definition, the proposed matrix norm $||\cdot||_{\leq f_0, p}$ exhibits the fundamental properties of a matrix norm, i.e. it is definite, homogeneous, and subadditive.

Using this frequency-weighted matrix norm, we now define the frequency-weighted residual indicator

$$\tilde{\epsilon}_{\leq f_0, p} := ||(\mathbf{I}_{\bar{n}} - \mathbf{PQP}^{-1})||_{\leq f_0, p}\,, \tag{31}$$

---

[6]more precisely: $||\mathbf{A}||_{\leq f_0, p} = \max\{||\mathbf{Ax}||_p, \mathbf{x} = \mathbf{Q}_{f_0}\,\hat{\mathbf{x}}, ||\hat{\mathbf{x}}||_p = 1\}$.

which is found to be clearly below one for Q-filters with $f_{\mathbf{Q}} > f_0$. This allows us to reduce the conservativeness of (28) in the sense that $\tilde{\epsilon}_p||\mathbf{r} - \mathbf{v}||_p$ can be replaced by $\tilde{\epsilon}_{\leq f_0,p}||\mathbf{r} - \mathbf{v}||_p$ if the reference and disturbance signals contain only frequencies below a known frequency f_0, i.e. if multiplication by the lowpass filter from the definition (30) yields

$$\mathbf{Q}_{f_0}(\mathbf{r} - \mathbf{v}) = (\mathbf{r} - \mathbf{v}). \tag{32}$$

4.2   *Monotonic Decrease of the Maximum-Pass-Length Error 1-Norm*

A second source of conservativeness of Theorem 2 lies in the fact that $\tilde{\kappa}_p$ is typically close to one, as discussed above. However, if we decide to use the 1-norm for convergence analysis, we find the following condition for monotonic convergence, which is clearly less conservative.

Theorem 3 (Monotonic decrease of MPL error 1-norm above a threshold): *Consider a repetitive process* (4) *with regular* $\mathbf{P}$ *and input-update law* (20), *and assume that the reference and disturbance signals contain only frequencies below a known frequency* f_0, *i.e.*

$$\mathbf{Q}_{f_0}(\mathbf{r} - \mathbf{v}) = (\mathbf{r} - \mathbf{v}), \tag{33}$$

where $\mathbf{Q}_{f_0}$, *with* $||\mathbf{Q}_{f_0}||_1 = 1$, *is a regular lifted matrix of a lowpass filter with cutoff frequency* f_0. *Assume furthermore that the learning gain matrix* $\mathbf{L}$ *is chosen such that* $\tilde{\gamma}_1 < 1$. *Then the 1-norm of the maximum-pass-length error is reduced from trial* j *to trial* $j + 1$, *i.e.* $||\bar{\mathbf{e}}_{j+1}||_1 < ||\bar{\mathbf{e}}_j||_1$, *if its measured portion satisfies*

$$||\mathbf{e}_j||_1 > \frac{\tilde{\epsilon}_{\leq f_0,1}||\mathbf{r} - \mathbf{v}||_1 + (\tilde{\kappa}_1 - 1)||\hat{\mathbf{e}}_j||_1}{1 - \tilde{\gamma}_1}. \tag{34}$$

Proof. The proof follows the lines of the proof of Theorem 2, but it exploits the fact that the 1-norm of a vector is the sum of the absolute values of its entries. Therefore, we find

$$\frac{||\bar{\mathbf{e}}_{j+1}||_1}{||\bar{\mathbf{e}}_j||_1} = \frac{||\bar{\mathbf{e}}_{j+1}||_1}{||\mathbf{e}_j||_1 + ||\hat{\mathbf{e}}_j||_1} \leq \frac{\tilde{\gamma}_1||\mathbf{e}_j||_1 + \tilde{\epsilon}_{\leq f_0,1}||\mathbf{r} - \mathbf{v}||_1 + \tilde{\kappa}_1||\hat{\mathbf{e}}_j||_1}{||\mathbf{e}_j||_1 + ||\hat{\mathbf{e}}_j||_1}, \tag{35}$$

which is less than one, since (34) yields $||\mathbf{e}_j||_1 > \tilde{\gamma}_1||\mathbf{e}_j||_1 + \tilde{\epsilon}_{\leq f_0,1}||\mathbf{r} - \mathbf{v}||_1 + (\tilde{\kappa}_1 - 1)||\hat{\mathbf{e}}_j||_1$. $\square$

   Theorem 2 and Theorem 3 provide practically useful conditions for convergence analysis. In particular, the conditions (28) and (34) yield upper bounds of the residual errors that may occur when a Q-filter is used. It is known from standard ILC theory that these residual errors depend on $\mathbf{r} - \mathbf{v}$. In variable-pass-length systems, however, they also depend on the hypothetical tracking error $\hat{\mathbf{e}}_j$ that would occur on the last $\bar{n} - n_j$ samples. Assume for example that $\tilde{\epsilon}_{\leq f_0,1}||\mathbf{r} - \mathbf{v}||_1$ is negligibly small, that $\tilde{\kappa}_1 = 1.1$ and that $\tilde{\gamma}_1 = 0.5$. Then the 1-norm of the MPL error is guaranteed to decrease as long as $||\mathbf{e}_j||_1 > 0.2\,||\hat{\mathbf{e}}_j||_1$. This means that the tracking is improved at least until the measured portion of the error becomes smaller than a fifth of the hidden portion.

4.3   *Multiple-Input Multiple-Output Systems*

To simplify the preceding derivations and arguments, we have assumed that the input and the output of the system to be controlled are both scalar. At this point, we briefly comment on the value of the above results for multiple-input multiple-output systems with $u(t), y(t), r(t), v(t) \in \mathbb{R}^m$.

Obviously, for such systems $\bar{\mathbf{u}}_j, \bar{\mathbf{y}}_j, \bar{\mathbf{e}}_j, \mathbf{r}, \mathbf{v} \in \mathbb{R}^{m\bar{n}}$ and $\mathbf{y}_j, \mathbf{e}_j \in \mathbb{R}^{mn_j}$ with

$$\bar{\mathbf{u}}_j := \left[u(t_{0,j})^T, u(t_{0,j} + t_\mathrm{s})^T, \ldots, u(t_{0,j} + (\bar{n} - 1)t_\mathrm{s})^T\right]^T, \quad \mathbf{P}, \mathbf{L}, \mathbf{Q} \in \mathbb{R}^{m\bar{n} \times m\bar{n}},$$

$$\langle \cdot \rangle_{n_j} : \mathbb{R}^{m\bar{n}} \to \mathbb{R}^{mn_j}, \quad \langle \cdot \rangle_{\bar{n}} : \mathbb{R}^{mn_j} \to \mathbb{R}^{m\bar{n}}, \quad \mathbf{H}_{n_j} := \mathrm{blockdiag}\{\mathbf{I}_{mn_j}, \mathbf{0}_{m(\bar{n}-n_j)\times m(\bar{n}-n_j)}\}.$$

With these slight modifications and with $\mathbf{I}_{mn_j}$ instead of $\mathbf{I}_{n_j}$, we find that the learning law, the definition of the maximum-pass-length error, and the closed-loop dynamics are equivalent to (3), (5), (7), respectively. Likewise, we proceed and redefine the dimensions of the partition

$$\mathbf{G}_{1,j} \in \mathbb{R}^{mn_j \times mn_j}, \ \mathbf{G}_{2,j} \in \mathbb{R}^{mn_j \times m(\bar{n}-n_j)}, \ \mathbf{G}_{3,j} \in \mathbb{R}^{m(\bar{n}-n_j)\times mn_j}, \ \mathbf{G}_{4,j} \in \mathbb{R}^{m(\bar{n}-n_j)\times m(\bar{n}-n_j)},$$

and equivalently for $\mathbf{K}_{i,j}, i = 1, 2, 3, 4$. This finally reveals that all of the results of Sections 3 and 4 also hold analogously for the considered case of multiple-input multiple-output systems.

In the following, we summarize the above findings in guidelines for controller design.

13

5.   Controller Design Guidelines

As Theorem 2 suggests, it is advantageous to choose a learning gain matrix that leads to a small convergence indicator $\tilde{\gamma}_p$. As also discussed above, the choice $\mathbf{L} = \mathbf{P}^{-1}$ leads to the fastest convergence. This, however, assumes that the input-output dynamics of the system to be controlled are precisely known. If model uncertainties are significant, then a more careful choice is advisable. As discussed in the classical ILC literature (see for example Bristow et al. (2006) and references therein), a small diagonal learning gain will yield $||\mathbf{I}_{m\overline{n}} - \mathbf{P}\mathbf{L}||_p$ just below one for a large class of system dynamics, which results in slow convergence with a potentially large residual error but good robustness with respect to model uncertainties.

When faced with the question of how to choose the cutoff frequency $f_{\mathbf{Q}}$ of the Q-filter, one should note the following relationships. For a given f_0, the indicator $\tilde{\epsilon}_{\leq f_0, p}$ becomes small if $f_{\mathbf{Q}}$ is raised above f_0. However, as discussed before, a small $f_{\mathbf{Q}}$ improves robustness with respect to model uncertainties by preventing high-frequency portions of the measured output from influencing the learning process. As in classical control theory, one must balance between controller performance (bandwidth) on the one hand and robustness and measurement noise rejection on the other hand.

From the preceding discussion we conclude the following controller design guidelines for ILC systems of the form (4), (20):

G1. An upper bound f_0 for the largest frequency that occurs in the reference signal $\mathbf{r}$ or in the disturbance signal $\mathbf{v}$ should be known.
G2. The model $\mathbf{P}$ should approximate the input-output dynamics of the system (at least with small or moderate uncertainties) for frequencies between 0 and f_0.
G3. The cutoff frequency $f_{\mathbf{Q}}$ of the Q-filter $\mathbf{Q}$ should be smaller than the smallest frequency at which large model uncertainties occur.
G4. At the same time, $f_{\mathbf{Q}}$ should be chosen such that a small residual indicator $\tilde{\epsilon}_{\leq f_0, p}$ is achieved for the given system dynamics $\mathbf{P}$. This typically implies $f_{\mathbf{Q}} \geq f_0$.
G5. The learning gain matrix $\mathbf{L}$ should be chosen such that a small convergence indicator $\tilde{\gamma}_p$ is achieved for some induced matrix norm $|| \cdot ||_p$.

Following these guidelines will prevent hidden error growth and will assure that the tracking error is quickly reduced (at least) until its norm falls below a small threshold.

6.   Simulation Study of an Example System

The goal of this section is to illustrate some variable-pass-length learning phenomena and to demonstrate the usefulness of the criteria derived above. For this purpose, controller design and simulations are carried out for a simple example application system that is motivated by Nahrstaedt et al. (2008a). There, functional electrical stimulation (FES) of the tibialis anterior muscle is used to induce a predefined foot movement during the swing phase of gait (i.e. between toe-off and initial contact, when the foot moves forward without ground contact).

The stimulation pulse width of the FES current pulses (in tenths of milliseconds) and the ankle joint angle (in degrees) are the input variable $u(t)$ and the output variable $y(t)$, respectively. Then the following linear difference equation model roughly approximates the input-output dynamics of the system at a sampling interval of $t_s = 0.02\,\text{s}$ (cf. Nahrstaedt et al. (2008b)):

$$y(t) - 0.8097002\, y(t - t_s) - 0.0777289\, y(t - 2t_s) = 0.6634\, u(t - 2t_s), \tag{36}$$

The ankle joint angle $y(t)$ shall be controlled to follow a predefined reference trajectory in each trial, i.e. in the swing phase of each stride. Based on experimental data (Seel et al., 2013b), the duration of swing phase in the gait of stroke patients is estimated to vary between $0.6\,\text{s}$ and $1.0\,\text{s}$,

which corresponds to $\underline{n} = 30$ and $\overline{n} = 50$. Therefore, the following full-length lifted signal vectors are defined:

$$\bar{\mathbf{u}}_j = [u_j(t_0),\ u_j(t_0 + t_\mathrm{s}),\dots,\ u_j(t_0 + 49t_\mathrm{s})]^T \in \mathbb{R}^{\overline{n}}, \tag{37}$$

$$\bar{\mathbf{y}}_j = [y_j(t_0 + 2t_\mathrm{s}),\ y_j(t_0 + 3t_\mathrm{s}),\dots,\ y_j(t_0 + 51t_\mathrm{s})]^T \in \mathbb{R}^{\overline{n}}, \tag{38}$$

$$\mathbf{r} = [r(t_0 + 2t_\mathrm{s}),\ r(t_0 + 3t_\mathrm{s}),\dots,\ r(t_0 + 51t_\mathrm{s})]^T \in \mathbb{R}^{\overline{n}}, \tag{39}$$

$$\mathbf{v} = [v(t_0 + 2t_\mathrm{s}),\ v(t_0 + 3t_\mathrm{s}),\dots,\ v(t_0 + 51t_\mathrm{s})]^T \in \mathbb{R}^{\overline{n}}, \tag{40}$$

where the input-output time shift of two samples is chosen to compensate the relative degree of the system dynamics (36). The Markov parameters of (36) are determined and, from those, the (regular, lower triangular, Toeplitz) lifted system matrix $\mathbf{P}$ is calculated.

Based on experimental data, reasonable values for the sequence of pass lengths $\{n_j\}$, for the desired output $\mathbf{r}$ and for the zero-input response signal $\mathbf{v}$ are found. The initial input $\bar{\mathbf{u}}_0$ is chosen to be a constant pulse width of $0.2\,\mathrm{ms}$.

A two-parameter learning gain matrix $\mathbf{L}_\mathrm{opt}$ is employed with $(\mathbf{L}_\mathrm{opt})_{i,i} = l_1 \,\forall i \in [1,\overline{n}]$ on the main diagonal, and $(\mathbf{L}_\mathrm{opt})_{i,i-1} = l_2 \,\forall i \in [2,\overline{n}]$ just below the main diagonal. If no Q-filter was used, this would correspond to the following first-order controller dynamics:

$$\begin{aligned}
u_{j+1}(t) = u_j(t) &+ l_1\left(r(t + 2t_\mathrm{s}) - y_j(t + 2t_\mathrm{s})\right) \\
&+ l_2\left(r(t + 1t_\mathrm{s}) - y_j(t + 1t_\mathrm{s})\right), \quad t = t_0,\dots,t_0 + (n_j - 1)t_\mathrm{s}.
\end{aligned} \tag{41}$$

The parameters of $\mathbf{L}_\mathrm{opt}$ are chosen by numerically minimizing the norm $||\mathbf{I}_{\overline{n}} - \mathbf{P}\mathbf{L}_\mathrm{opt}||_1$, as suggested by Theorem 1, over wide ranges of both l_1 and l_2. For $l_1 = 1.5$, $l_2 = -1.35$, an approximate plant inversion with a small convergence indicator of $\gamma_1 < 0.13$ is achieved. Therefore, when employing the learning law (6), *almost-strict*[7] monotonic convergence of $||\bar{\mathbf{e}}_j||_1$ (for all $j, n_j, \mathbf{v}, \mathbf{u}_0$) follows from Corollary 1 without further computational effort. By calculation of $\max_{n \in \{\underline{n},\dots,\overline{n}\}} ||\mathbf{I}_{\overline{n}} - \mathbf{P}\mathbf{L}_\mathrm{opt}\mathbf{H}_n||_2 \approx 1.004$, it is furthermore found that the Euclidean norm $||\bar{\mathbf{e}}_j||_2$ of the maximum-pass-length error is not only bounded by $||\bar{\mathbf{e}}_j||_1$ from above, but also bounded to rise (if at all) in no pass by more than 0.4%.

A zero-phase lowpass filter (2^nd-order Butterworth) is used as a Q-filter. Since $\mathbf{r}$ and $\mathbf{v}$ are known to contain only frequencies below $f_0 = 2\,\mathrm{Hz}$, the Q-filter cutoff frequency is chosen to be $f_\mathbf{Q} = 5\mathrm{Hz}$, such that $\tilde{\epsilon}_{\leq f_0,2} = 0.214$. Calculation of the convergence indicator defined in Section 4 yields $\tilde{\gamma}_2 = 0.092$. According to Theorem 2, this implies that the Euclidean norm of the MPL error decreases, (at least) as long as the measured error norm is larger than a small threshold.

In practice, much less reliable model knowledge might be available. Therefore, and for the sake of comparison, a second learning gain $\mathbf{L}_\mathrm{diag}$ is designed by simply choosing a small diagonal learning gain $\mathbf{L}_\mathrm{diag} = 0.12\,\mathbf{I}_{\overline{n}}$, which yields a convergence indicator of $\tilde{\gamma}_2 = 0.954$. Obviously, this also guarantees monotonic decrease of $||\bar{\mathbf{e}}_j||_2$ for large measured error norms, but a much slower decrease and a potentially larger residual error norm should be expected.

The lifted system model $\mathbf{P},\mathbf{v}$ is simulated with both of the designed learning controllers $\mathbf{L}_\mathrm{diag}, \mathbf{Q}$ and $\mathbf{L}_\mathrm{opt}, \mathbf{Q}$. Input and output trajectories of the first seven passes (trials) (with pass lengths of $\{n_j\}_{j=0}^6 = \{34, 40, 30, 46, 50, 50, 37\}$) are presented in Figure 3. The reduction of the MPL error 1-norm and 2-norm is illustrated in Figure 4a.

In the plots for $\mathbf{L}_\mathrm{opt}$, a core phenomenon of variable-pass-length learning becomes apparent. Whenever a pass is long enough to encounter samples that have not been reached before, the controller starts to learn this piece of trajectory almost from scratch, i.e. starting with the according piece of $\bar{\mathbf{y}}_0$ that (hypothetically) would have been measured in the first pass if it had been long enough. Due to approximate plant inversion, the very next output already resembles the reference

---

[7] in the sense of Corollary 1, i.e. $||\bar{\mathbf{e}}_{j+1}||_1 < ||\bar{\mathbf{e}}_j||_1$ for all $n_j \in \{\underline{n},\dots,\overline{n}\}$ unless $\mathbf{e}_j$ is the zero vector.
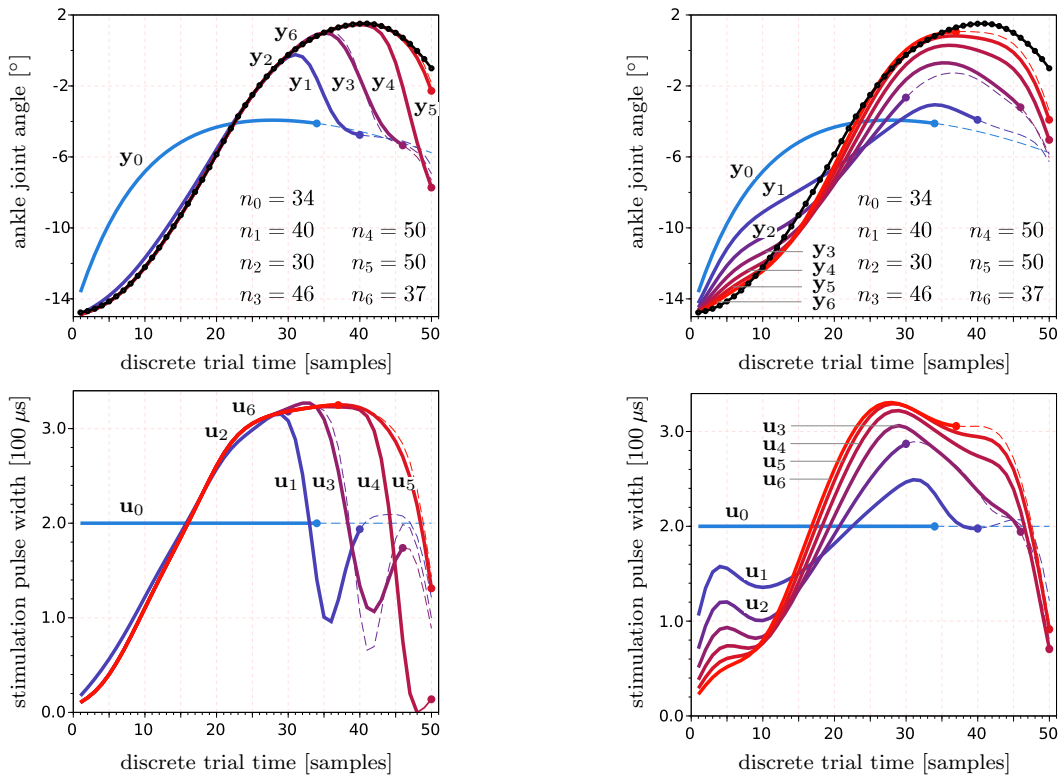
Figure 3. Variable-pass-length ILC. The end of each pass is marked with a dot (which also helps to identify the curves, since n_j is given). Lines are continued in dashed style to illustrate the (hypothetical) maximum-pass-length errors $\bar{\mathbf{e}}_j = \mathbf{r} - \bar{\mathbf{y}}_j$ with respect to the reference trajectory (black solid with dots). *Left*: Starting from a constant input $\bar{\mathbf{u}}_0$, the controller with optimized learning gain $\mathbf{L}_{\mathrm{opt}}$ quickly reduces the control deviation for all sample instants seen so far. *Right*: The controller with diagonal learning gain matrix learns slower and in a qualitatively different fashion.

(up to the longest pass length seen so far). The similarity of $\mathbf{y}_2$ and $\mathbf{y}_3$ shows that, when shorter passes occur, controller performance is not improved significantly any more.

In contrast, the diagonal controller learns slower and in a qualitatively different fashion. The measured errors result in input updates that clearly affect the output less precisely. After the first three learning steps (with pass lengths $n_0 = 34$, $n_1 = 40$, and $n_2 = 30$), the error on the first 30 samples is still significant. However, the samples $40 - 46$ of $\mathbf{y}_3$ are closer to the reference than the same samples of $\mathbf{y}_3$ in the left subplot. This extrapolating behavior of the less precise input update law seems advantageous. However, if either $\mathbf{v}$ or $\mathbf{r}$ would exhibit steeper slopes on the time interval $t \in [\underline{n}\,t_{\mathrm{s}}, ..., \overline{n}\,t_{\mathrm{s}}]$, this behavior would obviously lead to inferior performance.

In several further simulations, similarly quick reduction of the tracking error is observed for both learning gain matrices. Furthermore, simulations with modified controller parameters l_0, l_1, l_2 confirm that γ_p and $\tilde{\gamma}_p$ are indeed useful convergence indicators.

Influence of time-variant disturbances

A closer observation of experimental data reveals that the zero-input response $\mathbf{v}$ varies with time, which is most likely due to changes in the muscle tone of antagonistic muscles and due to small variations of the shank motion, both of which act as disturbances on the output of the system to be controlled. We model this effect by choosing an iteration-dependent signal vector $\mathbf{v}_j$ that varies unpredictably within an experimentally identified range, as illustrated by Figure 4b. Note that the signal changes dramatically within as few as three iterations.

Repeating the simulations described above with the time-variant $\mathbf{v}_j$ yields the results presented in Figure 4c. The MPL error 1-norm and 2-norm of the optimized learning gain $\mathbf{L}_{\mathrm{opt}}$ slightly increase
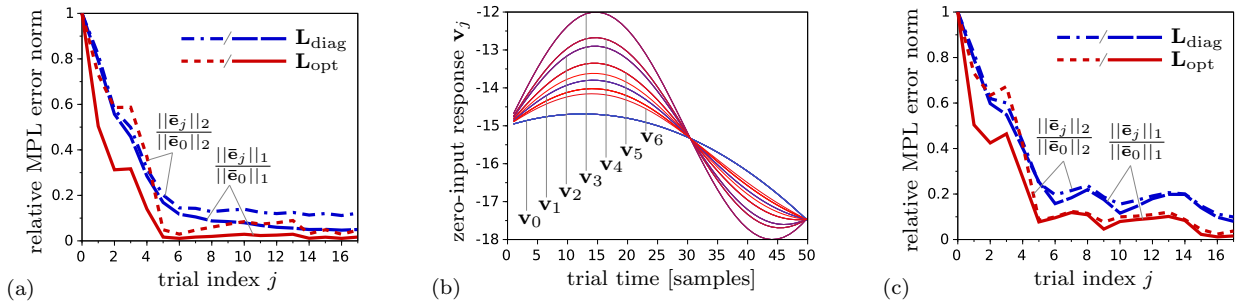
Figure 4. (a) Reduction of the MPL error 1-norm and 2-norm over the first 18 iterations when using an iteration-invariant zero-input response signal $\mathbf{v}_j$. (b) Iteration-variant zero-input response signal $\mathbf{v}_j$ as determined from experimental data. (c) Reduction of the MPL error 1-norm and 2-norm over the first 18 iterations when using the iteration-variant $\mathbf{v}_j$ from (b).

from trial 2 to trial 3. However, both norms quickly converge to a tenth of their original value. Rapid reduction of both norms is also observed for $\mathbf{L}_{\mathrm{diag}}$. After this first phase of convergence, the error norms fluctuate slightly more than in the case of a constant signal $\mathbf{v}$. However, we shall note that, despite the drastic variations in $\mathbf{v}_j$, the controllers manage to keep the MPL error norms close to the values of trials 5 and 6, which are only slightly above the remaining error norms in Figure 4a. This demonstrates that time-variant disturbances are compensated by the proposed iterative learning controllers, even if the disturbances change dramatically within few trials.

7.   Conclusions

A framework for the analysis of repetitive systems with variable pass length was developed. Iterative learning control schemes with and without lowpass filtering were proposed. The resulting closed-loop learning dynamics were derived and analyzed with respect to monotonous reduction of the tracking error. To cope with the variable dimension of the measured tracking error, the hypothetical maximum-pass-length error was introduced. This concept proved to be useful for convergence analysis in the presence of variable pass lengths.

For the standard input-update law without Q-filter, we obtained necessary and sufficient conditions that guarantee monotonic convergence of the MPL error norm. For the extended learning law with Q-filter, we derived conditions that assure a reduction of the tracking error at least until it falls below a potentially very small threshold. In this context, the issue of conservativeness in matrix-norm-based convergence criteria was addressed by proposing a frequency-weighted residual indicator to assess the non-zero residual error that is associated with Q-filter usage.

Based on all previous findings, we derived practical guidelines for the design of iterative learning controllers in the presence of variable pass length. In a simulation study, we demonstrated the usefulness of these guidelines and discussed typical phenomena of variable pass length learning.

Since we did not restrict $\mathbf{P}$ to be Toeplitz, all of the derived criteria can be employed for a very large class of systems including time-variant (but iteration-invariant) systems. Furthermore, all findings and guidelines hold as well if the pass length is constant and can therefore also be employed for convergence analysis and controller design in classical ILC systems.

Finally, it should be noted that in the motivational examples in Section 1, as in most other biomedical applications, the variability of the pass length is only one of the difficulties that complicate the employment of classical ILC theory. Issues like input saturation or iteration-variance (in the process dynamics, the initial conditions, the disturbance, or the reference signal) should be addressed individually as well as in the context of variable-pass-length systems.

17

References

H.-S. Ahn, K.L. Moore, YQ. Chen (2008). Discrete-time Intermittent Iterative Learning Controller with Independent Data Dropouts. *Proceedings of the 17th IFAC World Congress*, pp. 12442–47.

D.A. Bristow, M. Tharayil, A.G. Alleyne (2006). A Survey of Iterative Learning Control. *IEEE Control Systems Magazine*, 26(3):69–114.

H. Elci, R.W. Longman, M.Q. Phan, J.N. Juang, R. Ugoletti (2002). Simple Learning Control Made Practical by Zero-Phase Filtering: Applications to Robotics. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, 49(6):753–767.

T. Donkers, J. van de Wijdeven, O. Bosgra (2008). Robustness against model uncertainties of norm optimal iterative learning control. *American Control Conference 2008*, pp. 4561–4566.

A. Duschau-Wicke, A. Morger, H. Vallery, R. Riener. (2010). Adaptive Patient Support for Rehabilitation Robots. *at – Automatisierungstechnik*, 58(5):260–268.

C.T. Freeman, A.-M. Hughes, J.H. Burridge, P.H. Chappell, P.L. Lewin, E. Rogers (2010). Iterative learning control of FES applied to the upper extremity for rehabilitation. *Control Engineering Practice*, 17(3):368–381.

K. Galkowski, J. Lam, E. Rogers, S. Xu, B. Sulikowski, W. Paszke and D.H. Owens (2003). LMI based stability analysis and robust controller design for discrete linear repetitive processes. *International Journal of Robust and Nonlinear Control*, 13:1195–1211.

M. Guth, T. Seel, J. Raisch (2013). Iterative Learning Control with Variable Pass Length Applied to Trajectory Tracking on a Crane with Output Constraints. In *Proc. of the 52nd IEEE Conference on Decision and Control (CDC)*, pp. 6676–6681.

C. Klauer, T. Schauer, J. Raisch (2010). High Performance Motion Control by Neuro-Muscular Electrical Stimulation Applied to the Upper-Limb. *Proceedings of the 15th Annual IFESS Conference and 10th Vienna Int. Workshop on FES*, pp. 318–320.

E. Kurniawan, Z. Cao, Z. Man (2013). Design of Robust Repetitive Control with Time Varying Sampling Periods. *IEEE Transactions on Industrial Electronics*, 61(6):2834–41.

X.-F. Li; J.-X. Xu, D. Huang (2014). An Iterative Learning Control Approach for Linear Systems With Randomly Varying Trial Lengths. *IEEE Transactions on Automatic Control*, 59(7):1954–60.

X.-F. Li, J.-X. Xu (2015). Lifted system framework for learning control with different trial lengths. *International Journal of Automation and Computing*, 12(3):273–280.

X.-F. Li, J.-X. Xu, D. Huang (2015). Iterative learning control for nonlinear dynamic systems with randomly varying trial lengths. *International Journal of Adaptive Control and Signal Processing*, 29(11):1341–1353.

J. Liu, Y. Wang, H. Tong, R.P.S. Han (2012). Iterative Learning Control Based on Radial Basis Function Network for Exoskeleton Arm. *Advanced Materials Research*, Vols. 415–417, pp. 116–22.

R.W. Longman and K.D. Mombaur (2006). Investigating the Use of Iterative Learning Control and Repetitive Control to Implement Periodic Gaits. *Lecture notes in control and information sciences*, 340:189–218.

K.L. Moore, M. Johnson, M.J. Grimble (1993). Iterative Learning Control for Deterministic Systems. *Springer (Advances in Industrial Control)*.

J.-H. Moon, T.-Y. Doh, M.J. Chung (1998). A Robust Approach to Iterative Learning Control Design for Uncertain Systems. *Automatica*, 34(8):1001–1004.

K.L. Moore, A. Mathews (1998). Iterative Learning Control with Non-Standard Assumptions Applied to the Control of Gas-Metal ARC Welding (Book Section). *Iterative Learning Control*, Springer.

K.L. Moore (2000). A non-standard iterative learning control approach to tracking periodic signals in discrete-time non-linear systems. *International Journal of Control*, 73(10):955–967.

K.L. Moore, M. Ghosh, Y.Q. Chen (2007). Spatial-based iterative learning control for motion control applications. *Meccanica*, 42(2):167–175.

H. Nahrstaedt, T. Schauer, R. Shalaby, S. Hesse, J. Raisch (2008a). Automatic Control of a Drop-Foot Stimulator Based on Angle Measurement Using Bioimpedance. *Artificial Organs*, 32(8):649–654.

H. Nahrstaedt, T. Schauer, S. Hesse, J. Raisch (2008b). Iterative Learning Control of a Gait Neuroprosthesis. *at - Automatisierungstechnik*, 56(9):494–501.

M. Norrlöf and S. Gunnarsson (2002). Time and frequency domain convergence properties in iterative learning control. *International Journal of Control*, 75(14):1114–1126.

D.H. Owens and J. Hätönen (2005). Iterative learning control - An optimization paradigm. *Annual Reviews in Control*, 29:57–70.

K.H. Park, Z. Bien (2008). A generalized iterative learning controller against initial state error. *International Journal of Control*, 73(10):871–881.

W. Paszke, E. Rogers, K. Galkowski (2010). On the Design of ILC Schemes for Finite Frequency Range Tracking Specifications. *Proceedings of the IEEE Conference on Decision & Control*, pp. 6979–84

E. Rogers, K. Galkowski, D.H. Owens (2007). Control Systems Theory and Applications for Linear Repetitive Processes. *Springer (Lecture Notes in Control and Information Sciences, 349).*

T. Seel, T. Schauer, J. Raisch (2011). Iterative Learning Control for Variable Pass Length Systems. *Proceedings of the 18th IFAC World Congress*, pp. 4880–85.

T. Seel, S. Weber, K. Affeld, T. Schauer (2013a). Iterative Learning Cascade Control of Continuous Noninvasive Blood Pressure Measurement. *IEEE International Conference on Systems, Man, and Cybernetics*, pp. 2207–12.

T. Seel, M. Valtin, T. Schauer (2013b). Design and Control of an Adaptive Peroneal Stimulator with Inertial Sensor-based Gait Phase Detection. *Proceedings of the 18th International FES Society Conference.*

A. Soska, C.T. Freeman, E. Rogers (2012). ILC for FES-based Stroke Rehabilitation of Hand and Wrist. *IEEE Multi-Conference on System and Control*, 1267–72.

Y. Wang, E. Dassau, F.J. Doyle III (2010). Closed-loop control of artificial pancreatic-cell in type 1 diabetes mellitus using model predictive iterative learning control. *IEEE Transactions on Biomedical Engineering*, 57(2):211–219.