2017 IEEE International Conference on Robotics and Automation (ICRA)
Singapore, May 29 - June 3, 2017

High-Precision Trajectory Tracking in Changing Environments Through $\mathcal{L}_1$ Adaptive Feedback and Iterative Learning

Karime Pereida, Rikky R. P. R. Duivenvoorden, and Angela P. Schoellig

***Abstract*— As robots and other automated systems are introduced to unknown and dynamic environments, robust and adaptive control strategies are required to cope with disturbances, unmodeled dynamics and parametric uncertainties. In this paper, we propose and provide theoretical proofs of a combined $\mathcal{L}_1$ adaptive feedback and iterative learning control (ILC) framework to improve trajectory tracking of a system subject to unknown and changing disturbances. The $\mathcal{L}_1$ adaptive controller forces the system to behave in a repeatable, predefined way, even in the presence of unknown and changing disturbances; however, this does not imply that perfect trajectory tracking is achieved. ILC improves the tracking performance based on experience from previous executions. The performance of ILC is limited by the robustness and repeatability of the underlying system, which, in this approach, is handled by the $\mathcal{L}_1$ adaptive controller. In particular, we are able to generalize learned trajectories across different system configurations because the $\mathcal{L}_1$ adaptive controller handles the underlying changes in the system. We demonstrate the improved trajectory tracking performance and generalization capabilities of the combined method compared to pure ILC in experiments with a quadrotor subject to unknown, dynamic disturbances. This is the first work to show $\mathcal{L}_1$ adaptive control combined with ILC in experiment.**

I. INTRODUCTION

Robots and automated systems are being increasingly deployed in unknown and dynamic environments. Operating in these environments requires sophisticated control methods that can guarantee high overall performance even in the presence of model uncertainties, unknown disturbances and changing dynamics. Examples of robotic applications in these increasingly challenging environments include autonomous driving, assistive robotics and unmanned aerial vehicle (UAV) applications such as airborne package delivery. In the latter example, UAVs are required to deliver packages with different mass properties (mass, center of gravity and inertia), which influence the dynamic behavior of the UAV. Designing a controller to achieve high performance for each package is not feasible and small changes in the conditions may result in a dramatic decrease in controller performance and potential instability (see [1], [2] and [3]).

The goal of this work is to design a controller such that the system shows a repeatable and reliable behavior (that is, achieves, for the same reference input, the same output) even in the presence of unknown disturbances and changing dynamics, and improves its performance over time. In this paper, we focus on improving the trajectory tracking performance over task iterations, and propose and provide theoretical proofs of a combined $\mathcal{L}_1$ adaptive feedback and iterative learning control (ILC) framework (see Fig. 1). The $\mathcal{L}_1$ adaptive controller forces the system to behave in a repeatable, predefined way, even if it is subject to model uncertainties and unknown disturbances. As a result, we obtain a repeatable system; however, perfect trajectory tracking is not achieved. To learn from previous iterations and gradually improve the trajectory tracking performance of the overall system, we implement ILC. Experimental results on a quadrotor show that the proposed approach achieves high tracking performance despite dynamic disturbances. Moreover, we show that learned trajectories can be generalized across different system configurations because the $\mathcal{L}_1$ controller handles any (dynamic) disturbances that affect the system.

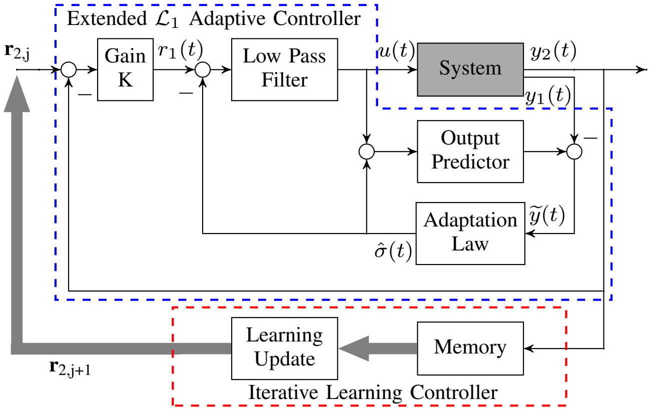

Fig. 1. Proposed framework to achieve high performance control in changing environments. The extended $\mathcal{L}_1$ adaptive controller forces the system to behave in a predefined, repeatable way. The iterative learning controller improves the tracking performance in each iteration j based on experience from previous executions.

$\mathcal{L}_1$ adaptive control and ILC have previously been combined to improve trajectory tracking performance (see [4], [5], and [6]). In previous work, the control input to the system ($u(t)$ in Fig. 1) was constructed by combining both $\mathcal{L}_1$ and ILC inputs in a parallel architecture. In contrast, the serial architecture proposed in this paper places the $\mathcal{L}_1$ adaptive control as an underlying controller, while the ILC acts as a high-level adaptation scheme that mainly compensates for systematic tracking errors. This serial architecture allows us to decouple the task of making the system behave in a predefined way even in the presence of disturbances, from

The authors are with the Dynamic Systems Lab (www.dynsyslab.org) at the University of Toronto Institute for Aerospace Studies (UTIAS), Canada. Email: {karime.pereida, rikky.duivenvoorden}@robotics.utias.utoronto.ca, schoellig@utias.utoronto.ca.

This research was supported in part by NSERC grant RGPIN-2014-04634, the Connaught New Researcher Award, and the Mexican National Council of Science and Technology (abbreviated CONACYT).

978-1-5090-4633-1/17/$31.00 ©2017 IEEE

the task of improving the tracking performance. Furthermore, the results presented in [4], [5], and [6] are restricted to simulations while the proposed approach is the first work to show the $\mathcal{L}_1$-ILC architecture in experiment.

$\mathcal{L}_1$ adaptive control is based on the model reference adaptive control (MRAC) architecture with the addition of a low-pass filter that decouples robustness from adaptation [7]. This allows arbitrarily high adaptation gains to be chosen for fast adaptation. This algorithm has been successfully implemented on UAVs to augment a baseline controller for improved disturbance rejection. Attitude control based on $\mathcal{L}_1$ adaptive control was shown in [8], where three algorithms were successfully implemented and tested on a quadrotor, hexacopter and octocopter, respectively. In [9], $\mathcal{L}_1$ adaptive control is implemented for a quadrotor in translational velocity output feedback control, and shows the ability of the controller to compensate for artificial reduction in the speed of a single motor. In this work, we also use $\mathcal{L}_1$ adaptive output feedback on translational velocity, as it guarantees robustness bounds, and has a-priori known steady-state and transient performance.

Iterative learning control efficiently uses information from previous trials to improve tracking performance within a small number of iterations by updating the feedforward input signal. ILC has successfully been applied to a variety of trajectory tracking scenarios such as motion control of industrial robot arms [10] and ground vehicles [11], manufacturing of integrated circuits [12], swinging up a pendulum [13], and quadrotor control [14]. For a survey on ILC, the reader is referred to [15]. In this paper, we use optimization-based ILC in conjunction with a model error estimator [16].

The remainder of this paper is organized as follows: We define the problem in Section II. Section III details the proposed approach and proves key features such as the transient behavior of the adaptive control. Section IV shows our experimental results, including examples with changing system dynamics. We compare our approach to one with a standard underlying feedback controller. Conclusions are provided in Section V.

II. PROBLEM STATEMENT

The goal of this work is to achieve high-precision tracking despite changing system dynamics and uncertain environment conditions. The system optimizes its performance, for a given desired trajectory, over multiple executions of the task. We aim to design an algorithm that does not require to re-learning if the system dynamics continue to change.

For simplicity of presentation, we assume the uncertain and changing system dynamics ('System' block in Fig. 1) can be described by a single-input single-output (SISO) system (this approach can be extended to multi-input multi-output (MIMO) systems as described in Section IV) identical to [7] for output feedback:

$$y_1(s) = A(s)(u(s) + d_{\mathcal{L}_1}(s)) , \qquad y_2(s) = \tfrac{1}{s} y_1(s) , \quad (1)$$

where $y_1(s)$ and $y_2(s)$ are the Laplace transforms of the translational velocity $y_1(t)$, and position $y_2(t)$, respectively, $A(s)$ is a strictly-proper *unknown* transfer function that can be stabilized by a proportional-integral controller, $u(s)$ is the Laplace transform of the input signal, and $d_{\mathcal{L}_1}(s)$ is the Laplace transform of the disturbance signal defined as $d_{\mathcal{L}_1}(t) \triangleq f(t, y_1(t))$, where $f : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ is an *unknown* map subject to the following assumption:

Assumption 1 (Global Lipschitz continuity)**.** *There exist constants $L > 0$ and $L_0 > 0$, such that the following inequalities hold uniformly in t:*

$$|f(t, v) - f(t, w)| \le L|v - w| , \text{ and} \quad (2)$$

$$|f(t, w)| \le L|w| + L_0 \quad \forall v, w \in \mathbb{R} . \quad (3)$$

The system is tasked to track a desired postition trajectory $y_2^*(t)$, which is defined over a finite-time interval and is assumed to be feasible with respect to the true dynamics of the $\mathcal{L}_1$-controlled system (Fig. 1, blue dashed box). This signal is discretized. We introduce the lifted representation, see [10], for the desired trajectory $\mathbf{y_2^*} = (y_2^*(1), \ldots, y_2^*(N))$, and the output of the plant $\mathbf{y_2} = (y_2(1), \ldots, y_2(N))$, where $N < \infty$ is the number of discrete samples. The tracking performance criterion J is defined as:

$$J \triangleq \min_{\mathbf{e}} \mathbf{e}^T \mathbf{Q} \mathbf{e}$$

where $\mathbf{e} = \mathbf{y}_2 - \mathbf{y_2^*}$ is the tracking error and $\mathbf{Q}$ is a positive definite matrix. The goal is to improve the tracking performance iteratively; that is, from execution to execution.

III. METHODOLOGY

We consider two main subsystems: the extended $\mathcal{L}_1$ adaptive controller (blue dashed box in Fig. 1) and the ILC (red dashed box in Fig. 1). The extended $\mathcal{L}_1$ adaptive controller is presented in Section III-A including proofs of its transient behavior. Section III-B introduces the ILC.

A. $\mathcal{L}_1$ Adaptive Control

In the proposed framework, the aim of the $\mathcal{L}_1$ adaptive controller is to make the system behave in a repeatable, predefined way, even when unknown, changing disturbances affect the system. In this subsection, we describe the extended $\mathcal{L}_1$ adaptive controller and provide proofs of the transient behavior.

In this work, the typical $\mathcal{L}_1$ adaptive output feedback controller for SISO systems [7] is nested within a proportional controller (see Fig. 1). This extended architecture is identical to [9]. The outer-loop proportional controller enables the system to remain within certain position boundaries. Given the proposed extended $\mathcal{L}_1$ adaptive control, we must show that the system performs provably close to a given reference model under the uncertainty defined in Section II. This is done by finding bounds for the transient behavior. The proof is inspired by [7], but is extended to include the proportional controller ('Gain K' in Fig. 1).

1) Problem Formulation: The objective of the extended $\mathcal{L}_1$ adaptive output feedback controller is to design a control input $u(t)$ such that $y_2(t)$ tracks a bounded piecewise continuous reference input $r_2(t)$. To achieve this, one method is for the output of the $\mathcal{L}_1$ adaptive controller nested within the proportional feedback loop $y_1(t)$ to track $r_1(t)$ according to a first-order reference system:

$$M(s) = \tfrac{m}{s+m}\,, \quad m > 0\,. \tag{4}$$

2) Definitions and $\mathcal{L}_1$-Norm Condition: The system in (1) can be rewritten in terms of the reference system (4):

$$y_1(s) = M(s)(u(s) + \sigma(s))\,, \tag{5}$$

where uncertainties in $A(s)$ and $d_{\mathcal{L}_1}(s)$ are combined into σ:

$$\sigma(s) \triangleq \frac{(A(s) - M(s))u(s) + A(s)d_{\mathcal{L}_1}(s)}{M(s)}\,. \tag{6}$$

We consider a strictly-proper low-pass filter $C(s)$ (see Fig. 1) with $C(0) = 1$, and a proportional gain $K \in \mathbb{R}^+$, such that:

$$H(s) \triangleq \frac{A(s)M(s)}{C(s)A(s) + (1 - C(s))M(s)} \quad \text{is stable,} \tag{7}$$

$$F(s) \triangleq \frac{1}{s + H(s)C(s)K} \quad \text{is stable,} \tag{8}$$

and the following $\mathcal{L}_1$-norm condition is satisfied:

$$\|G(s)\|_{\mathcal{L}_1} L < 1\,, \text{where } G(s) \triangleq H(s)(1 - C(s))F(s) \tag{9}$$

and L is the Lipschitz constant defined in Assumption 1.

The $\mathcal{L}_1$-norm condition is used to prove bounded-input bounded-output (BIBO) stability of a reference model that will describe the repeatable behavior of the underlying $\mathcal{L}_1$ controlled system. The solution of the $\mathcal{L}_1$-norm condition in (9) exists under the following assumptions:

Assumption 2 (Stability of $H(s)$)**.** *$H(s)$ is assumed to be stable for appropriately chosen low-pass filter $C(s)$ and first-order reference eigenvalue $-m < 0$.*

As indicated in [7], this assumption holds in cases where $A(s)$ can be stabilized by a proportional-integral controller.

Assumption 3 (Stability of $F(s)$)**.** *$F(s)$ is assumed to be stable for appropriately chosen proportional gain K.*

A sufficient condition for this assumption to be valid is if $A(s)$ is minimum phase stable, which holds if there is a controller within the system $A(s)$ that is stabilizing a plant without any unstable zeros. In the case of velocity control of a quadrotor, this assumption is valid. Less conservative conditions that guarantee the stability of $F(s)$ exist, but are not necessary for the application in this paper.

3) Extended $\mathcal{L}_1$ Adaptive Control Architecture: The SISO extended $\mathcal{L}_1$ adaptive controller architecture is shown in Fig. 1. With the exception of the proportional feedback loop, this architecture (from r_1 to y_1) is identical to [7]. The integrator from y_1 to y_2 allows the outer-loop to control the position, while the $\mathcal{L}_1$ adaptive feedback controls the velocity. The equations describing the implementation of the extended $\mathcal{L}_1$ output feedback architecture are presented below in (10), (11), (12), and (13).

Output Predictor: The following output predictor is used within the $\mathcal{L}_1$ adaptive output feedback architecture:

$$\dot{\hat{y}}_1(t) = -m\hat{y}_1(t) + m(u(t) + \hat{\sigma}(t))\,, \qquad \hat{y}_1(0) = 0\,,$$

where $\hat{\sigma}(t)$ is the adaptive estimate of $\sigma(t)$. In the Laplace domain, this is equivalent to:

$$\hat{y}_1(s) = M(s)(u(s) + \hat{\sigma}(s))\,. \tag{10}$$

Adaptation Law: The adaptive estimate $\hat{\sigma}(t)$ is updated according to the following update law:

$$\dot{\hat{\sigma}}(t) = \Gamma \text{Proj}(\hat{\sigma}(t), -mP\tilde{y}(t))\,, \qquad \hat{\sigma}(0) = 0\,, \tag{11}$$

where $\tilde{y}(t) \triangleq \hat{y}_1(t) - y_1(t)$, and $P > 0$ solves the algebraic Lyapunov equation $mP + Pm = 2mP = -Z$ for $Z > 0$. The variable $\Gamma \in \mathbb{R}^+$ is the adaptation rate subject to the lower bound as specified in [7]. Typically in $\mathcal{L}_1$ adaptive control, Γ is set very large. Experiments with this controller were carried out with an adaptation rate of $\Gamma = 1000$. The projection operator defined in [7] ensures that the estimation of σ is guaranteed to remain within a specified convex set.

Control Law: The control input signal is the difference between the $\mathcal{L}_1$ desired trajectory signal r_1 and the adaptive estimate $\hat{\sigma}$ after passing through the low-pass filter $C(s)$:

$$u(s) = C(s)(r_1(s) - \hat{\sigma}(s))\,. \tag{12}$$

This means that only the low frequencies of the uncertainties within $A(s)$ and $d_{\mathcal{L}_1}(s)$, which the system is capable of counteracting, are compensated for. The high frequency portion is attenuated by the low-pass filter.

Closed-Loop Feedback: The following equation describes the closed-loop feedback acting on the input to the $\mathcal{L}_1$ adaptive output feedback controller r_1 based on the output of the system y_1. As discussed above: $y_2(s) \triangleq \frac{1}{s}y_1(s)$, and the negative feedback is defined as follows:

$$r_1(s) = K(r_2(s) - y_2(s))\,, \tag{13}$$

where the objective is for y_2 to track r_2.

4) Transient and Steady-State Performance: The extended $\mathcal{L}_1$ adaptive controller is required to perform repeatably and consistently. This is done by guaranteeing that the difference between the output of a known BIBO stable reference system and the output of the actual system is uniformly bounded. Intuitively, the reference system describes the desired behavior of the actual system.

The proof starts off by presenting a BIBO stable closed-loop reference system. This reference system is then compared to the actual extended $\mathcal{L}_1$ adaptive output feedback controller.

Lemma 1. *Let $C(s)$, $M(s)$ and K satisfy the $\mathcal{L}_1$-norm condition in* (9)*. Then the following closed-loop reference system:*

$$\begin{aligned} y_{2,\mathit{ref}}(s) &= F(s)H(s)\big(C(s)Kr_2(s) + (1 - C(s))d_{\mathit{ref}}(s)\big) \\ d_{\mathit{ref}}(t) &\triangleq f(t, y_{2,\mathit{ref}}(t)) \end{aligned} \tag{14}$$

is BIBO stable.

Proof. Since $r_2(t)$ is bounded and $H(s)$, $C(s)$ and $F(s)$ are strictly-proper stable transfer functions, taking the norm of the reference system and making use of Assumption 1 yields the following bound:

$$\begin{aligned}\|y_{2,\mathrm{ref}_\tau}\|_{\mathcal{L}_\infty} \le K\|H(s)C(s)F(s)\|_{\mathcal{L}_1}\|r_2\|_{\mathcal{L}_\infty} \\ + \|G(s)\|_{\mathcal{L}_1}(L\|y_{2,\mathrm{ref}_\tau}\|_{\mathcal{L}_\infty} + L_0)\,,\end{aligned} \quad (15)$$

where $\|y_{2,\mathrm{ref}_\tau}\|_{\mathcal{L}_\infty}$ is the truncated $\mathcal{L}_\infty$-norm of the signal $y_{2,\mathrm{ref}}(t)$ up to $t = \tau$. Let ρ_r be defined as follows:

$$\rho_r \triangleq \frac{K\|H(s)C(s)F(s)\|_{\mathcal{L}_1}\|r_2\|_{\mathcal{L}_\infty} + \|G(s)\|_{\mathcal{L}_1}L_0}{1 - \|G(s)\|_{\mathcal{L}_1}L}\,. \quad (16)$$

From the $\mathcal{L}_1$-norm condition in (9) and the definition of ρ_r in (16):

$$\|y_{2,\mathrm{ref}_\tau}\|_{\mathcal{L}_\infty} \le \rho_r\,. \quad (17)$$

This result holds uniformly, so $\|y_{2,\mathrm{ref}}\|_{\mathcal{L}_\infty}$ is bounded. Hence, the closed-loop reference system in (14) is BIBO stable. □

Theorem 1. *Consider the system in* (1)*, with a control input from the extended $\mathcal{L}_1$ output feedback adaptive controller defined in* (10)*,* (11)*,* (12)*, and* (13)*. Suppose $C(s)$, $M(s)$ and K satisfy the $\mathcal{L}_1$-norm condition in* (9)*. Then the following bounds hold:*

$$\|\tilde{y}\|_{\mathcal{L}_\infty} \le \gamma_0\,, \quad (18)$$

$$\|y_{2,\mathit{ref}} - y_2\|_{\mathcal{L}_\infty} \le \gamma_1\,, \quad (19)$$

where $\tilde{y}(t) \triangleq \hat{y}_1(t) - y_1(t)$, $\gamma_0 \propto \sqrt{\frac{1}{\Gamma}}$ is defined in [7], and

$$\gamma_1 \triangleq \frac{\left\|\dfrac{F(s)H(s)C(s)}{M(s)}\right\|_{\mathcal{L}_1}}{1 - \|G(s)\|_{\mathcal{L}_1}L}\gamma_0\,. \quad (20)$$

Proof. See Appendix. □

The bounds given in (18) and (19) show that the difference between the output predictor and the system output $y_1(t)$ and the difference between the reference system and the system output $y_2(t)$ are uniformly bounded with bounds inversely proportional to the square root of the adaptation gain Γ. This means that for high adaptation gains, the actual system approaches the behavior of the reference system (14). Hence, the system achieves repeatable and consistent performance, which is required for ILC.

B. Iterative Learning Control

We use ILC to improve the tracking performance of the underlying, repeatable system. The algorithm updates the feedforward signal $r_2(t)$ based on data gathered during previous iterations. The ILC implementation in this work is based on [16]. In this subsection, we give a brief summary of the optimization-based ILC used in this work and highlight the differences to the approach in [16], where a more detailed description is found.

We consider a repeatable system as seen by the ILC, which includes both the plant and the extended $\mathcal{L}_1$ adaptive controller (blue dashed box and shadowed box in Fig. 1), and whose key dynamics can be represented by the following model:

$$\dot{x}(t) = g(x(t), r_2(t))\,, \quad y_2(t) = h(x(t))\,, \quad (21)$$

where g and h are nonlinear function, $r_2(t) \in \mathbb{R}$ is the control input to the system, $x(t) \in \mathbb{R}^{n_x}$ is the state and $y_2(t) \in \mathbb{R}$ is the output. To satisfy the typical ILC assumption of identical initial conditions, despite unknown disturbances, experiments start when the system state is in close vicinity of the desired initial state. This is possible as the $\mathcal{L}_1$ adaptive controller compensates for the effect of unknown disturbances.

The desired output trajectory $y_2^*(t)$ is assumed to be feasible based on the nominal model (21), where $(r_2^*(t),\ x^*(t),\ y_2^*(t))$ satisfy (21). We assume that the system stays relatively close to the reference trajectory; hence, we only consider small deviations from the above nominal trajectories, $\tilde{r}_2(t)$, $\tilde{x}(t)$ and $\tilde{y}_2(t)$. The system is linearized about the nominal trajectories to obtain a time-varying, linear state-space model, which approximates the system dynamics along the reference trajectory. The system is discretized and rewritten in the lifted representation as in [16]. We define $\bar{\mathbf{y}}_2 = (\tilde{y}_2(1), \ldots, \tilde{y}_2(N)) \in \mathbb{R}^N$ and analogously we define $\bar{\mathbf{r}}_2$. The lifted representation for the extended system is written as:

$$\bar{\mathbf{y}}_{2,j} = \mathbf{F}_{\mathrm{ILC}}\bar{\mathbf{r}}_{2,j} + \mathbf{d}_j\,, \quad (22)$$

where the subscript j denotes the iteration number, $\mathbf{F}_{\mathrm{ILC}}$ is a constant matrix derived from the nominal model and $\mathbf{d}$ represents a repetitive disturbance that is initially unknown.

Using the approach presented in [14] and [16], an iteration-domain Kalman filter for the system (22) is used to compute the estimate $\widehat{\mathbf{d}}_{j|j}$ based on measurements from iterations $1, \ldots, j$.

An optimization-based update step computes the next reference sequence $\bar{\mathbf{r}}_{2,j+1}$ that compensates for the identified disturbance $\widehat{\mathbf{d}}_{j|j}$ and estimated output error $\widehat{\mathbf{y}}_{j+1|j}$, where $\widehat{\mathbf{y}}_{j+1|j} = \mathbf{F}_{\mathrm{ILC}}\bar{\mathbf{r}}_{2,j} + \widehat{\mathbf{d}}_{j|j}$. In the input update step, the following quadratic cost function is minimized:

$$\min_{\bar{\mathbf{r}}_{2,j+1}} \left(\widehat{\mathbf{y}}_{j+1|j}^T\mathbf{Q}\widehat{\mathbf{y}}_{j+1|j} + \bar{\mathbf{r}}_{2,j+1}^T\mathbf{S}\bar{\mathbf{r}}_{2,j+1} + \ddot{\bar{\mathbf{r}}}_{2,j+1}^T\mathbf{R}\ddot{\bar{\mathbf{r}}}_{2,j+1}\right) \quad (23)$$

subject to

$$\ddot{\bar{r}}_{2,j+1} \le \mathbf{a}_{\max}\,,$$

where $\mathbf{a}_{\max}$ is a constraint based on the maximum acceleration achievable by the physical system. The sequence $\ddot{\bar{\mathbf{r}}}_{2,j+1}$ represents the discrete approximation of the second derivative of the input reference. The constant matrices $\mathbf{Q}$, $\mathbf{R}$, $\mathbf{S}$ are symmetric positive definite matrices that weight different components of the cost function. The cost function tries to minimize the tracking error of the system (weighted by $\mathbf{Q}$), the control effort required (weighted by $\mathbf{S}$) and the rate of change of the reference signal derivative (weighted by $\mathbf{R}$). We use the IBM CPLEX optimizer to solve the above optimization problem. The cost function used in this work is different from the cost function in [16] as it includes both the input and its second derivative to improve the performance of the given task.

In previous work (see [17] and [18]) the convergence for optimization-based ILC with Kalman filter, such as the one used in this paper, was proven. However, the cost function in [17] and [18] differs from the cost function in this paper. Instead of including $\bar{\mathbf{r}}_{2,j+1}$ as in (23), the cost function in [17] and [18] only includes the reference input change from iteration to iteration $\Delta\bar{\mathbf{r}}_{2,j+1} = \bar{\mathbf{r}}_{2,j+1} - \bar{\mathbf{r}}_{2,j}$. Future work will extend the proof of [17], [18] to our setup (23).

IV. EXPERIMENTAL RESULTS

The proposed framework combining $\mathcal{L}_1$ adaptive control and ILC ($\mathcal{L}_1$-ILC) is used to minimize the trajectory tracking error of a quadrotor flying a three-dimensional trajectory under different dynamic disturbances. The SISO architecture derived in the previous section is extended to the MIMO quadrotor system by implementing (3×3) diagonal transfer function matrices for the low-pass filter and first-order output predictor. The signals $r_1(t)$, $r_2(t)$, $y_1(t)$, and $y_2(t)$ are the desired translational velocity, desired position, quadrotor translational velocity and quadrotor position, respectively. This implementation is identical to [9], which ensures that the quadrotor remains within the boundaries of the indoor flying space. Each element of the three-dimensional signals and each diagonal element of the transfer function matrices correspond to the x, y and z inertial directions, respectively.

The experiments were performed using the commercial quadrotor platform AR.Drone 2.0 from Parrot. An overhead motion capture camera system is used to obtain position information. To test the performance of the proposed approach under unknown, changing disturbances, we change the dynamic behavior of the quadrotor by adding a *mass disturbance*. To create the mass disturbance a 50 g mass is suspended 55 cm below the back-left leg, 17 cm from the geometric center of the frame, creating a pendulum.

We compare the performance of the proposed $\mathcal{L}_1$-ILC approach with that of a pure ILC with an underlying, non-adaptive proportional-derivative controller (PD-ILC). To quantify the controller performance, the error in the system is defined as:

$$e = \frac{\sum_{i=1}^{N} \sqrt{(e_x(i))^2 + (e_y(i))^2 + (e_z(i))^2}}{N} \qquad (24)$$

where $e_x(i) = r^*_{2,x}(i) - y_{2,x}(i)$, $e_y(i) = r^*_{2,y}(i) - y_{2,y}(i)$ and $e_x(i) = r^*_{2,z}(i) - y_{2,z}(i)$ are the deviations from the desired trajectory in each axis. We consider three scenarios to compare the performance of the control frameworks:learning convergence and generalizability, repeatability, and performance under changing conditions. In all three scenarios the $\mathcal{L}_1$-ILC approach outperforms the PD-ILC approach.

A. Learning Convergence and Generalizability

The quadrotor learns to track a desired trajectory using each of the two frameworks: PD-ILC and $\mathcal{L}_1$-ILC. The errors of this initial learning process (iteration 1-10) are depicted in Fig. 2a. The proposed $\mathcal{L}_1$-ILC shows lower errors consistently and converges faster.

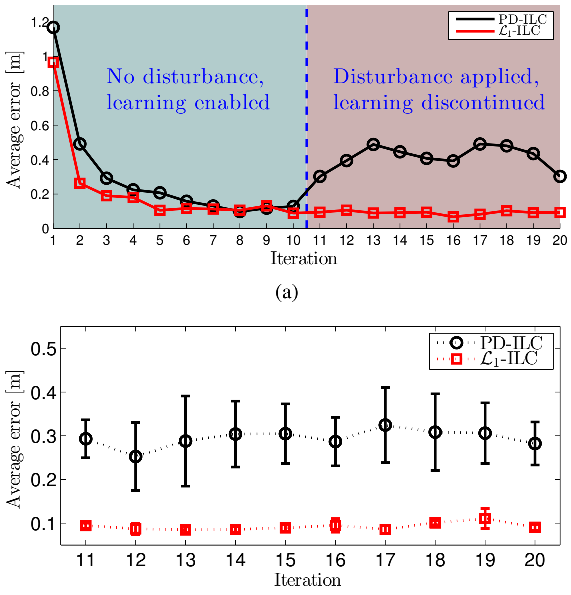

(a)

(b)

Fig. 2. (a) The $\mathcal{L}_1$-ILC approach shows a faster learning convergence initially. At iteration 11 a disturbance is applied and learning is disabled: the $\mathcal{L}_1$-ILC error is not affected while the PD-ILC error increases significantly. (b) The mean of the error across five 10-iteration sets shows the repeatability of the learned trajectory after a mass disturbance is applied to the system. The PD-ILC approach displays a significantly larger error and standard deviation compared to the $\mathcal{L}_1$-ILC approach.

After this initial learning process a mass disturbance is applied to the system and the learning is discontinued. The learned trajectory at iteration ten is repeated for ten more iterations with both the $\mathcal{L}_1$-ILC and PD-ILC framework, see Fig. 2a. The PD-ILC framework shows a 323% increase after the mass disturbance is applied. The $\mathcal{L}_1$-ILC approach shows no noticeable increase in the error because the $\mathcal{L}_1$ adaptive controller achieves repeatable behavior, despite the disturbances applied to the system.

B. Repeatability

To assess the repeatability of the overall control after a mass disturbance has been applied to the system, we discontinued learning and performed five experiments with ten iterations each for both control frameworks. Fig. 2b shows the average error of the five sets at each iteration along with their standard deviation. The system is more repeatable with the $\mathcal{L}_1$-ILC framework as the error and standard deviation are much smaller than with the PD-ILC framework.

C. Performance under Changing Conditions

The ability of the system to continue to learn after a disturbance has been applied is also explored. The errors while the system is learning without disturbance (first ten iterations) and with a mass disturbance (last ten iterations) are shown in Fig. 3a. The error increases significantly in the PD-ILC framework after the disturbance is applied. This error rapidly decreases as the system continues to learn;

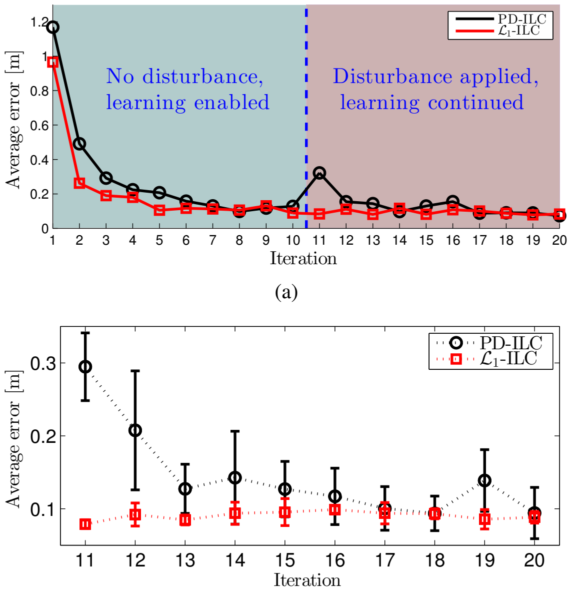

(a)

(b)

Fig. 3. (a) Learning behavior after a mass disturbance is applied to the system at the end of iteration ten. The error of the PD-ILC framework after the disturbance increases dramatically; while the error of the $\mathcal{L}_1$-ILC framework is virtually unchanged. (b) Average error across five sets of ten iterations of learning after a mass disturbance is applied. The PD-ILC approach displays a significantly larger error and standard deviation than that of the $\mathcal{L}_1$-ILC approach.

however, for some applications, this behavior may not be acceptable. The error in the $\mathcal{L}_1$-ILC framework does not change even after the mass disturbance has been applied.

The learning behavior is further explored by obtaining a total of five 10-iteration sets of the learning systems after the mass disturbance is applied. The average of the error and the standard deviation across the five sets are shown in Fig. 3b. The average error at iteration eleven for the PD-ILC framework is significantly higher than for the $\mathcal{L}_1$-ILC framework. The standard deviation is notably higher for the PD-ILC approach for all iterations. The $\mathcal{L}_1$-ILC experiments show that the learned input trajectory can be re-used even if the system dynamics are changed.

V. CONCLUSIONS

In this paper, we introduced an $\mathcal{L}_1$-ILC framework for trajectory tracking. The $\mathcal{L}_1$ adaptive controller forces the system to remain close to a predefined nominal system behavior, even in the presence of unknown and changing disturbances. However, having a repeatable system does not imply achieving zero tracking error. We use ILC to learn from previous iterations and improve the tracking performance over time. We proved that the proposed framework is stable and achieves learning convergence. Experiments on quadrotors showed significant performance improvements of the proposed $\mathcal{L}_1$-ILC approach compared to a non-adaptive PD-ILC approach in terms of learning convergence, repeatability, and behavior under disturbances. The learned reference trajectories of the $\mathcal{L}_1$-ILC framework are re-usable even if the system dynamics are changed, because the $\mathcal{L}_1$ adaptive controller compensates for the unknown, changing disturbances. As far as the authors are aware, this is the first work to show such an $\mathcal{L}_1$-ILC framework in real-world experiments and on quadrotor vehicles, specifically.

APPENDIX

Below we sketch the proof of Theorem 1:

Proof. Theorem 4.1.1 in [7] proves the bound in (18) under the same assumptions as made in this paper. The bound in (19) remains to be shown. The following definitions will become useful:

$$H_0(s) \triangleq \frac{A(s)}{C(s)A(s)+(1-C(s))M(s)}\,,\ \text{and} \tag{25}$$

$$H_1(s) \triangleq \frac{(A(s)-M(s))C(s)}{C(s)A(s)+(1-C(s))M(s)}\,. \tag{26}$$

In [7], it is shown that both $H_0(s)$ and $H_1(s)$ are strictly-proper stable transfer functions. Furthermore, the following expressions using (25) and (26) can be verified:

$$M(s)H_0(s) = H(s)\,,\ \text{and} \tag{27}$$

$$M(s)\Big(C(s){+}H_1(s)(1-C(s))\Big) = H(s)C(s)\,. \tag{28}$$

Let $\tilde{\sigma}(t) \triangleq \hat{\sigma}(t) - \sigma(t)$ where $\hat{\sigma}$ is the adaptive estimate, and σ is defined in (6). The control law in (12) can be expressed as:

$$u(s) = C(s)r_1(s) - C(s)(\tilde{\sigma}(s)+\sigma(s))\,. \tag{29}$$

Substitution of (29) into (6) and making use of the definitions in (25) and (26) results in the following expression for $\sigma(s)$:

$$\sigma(s) = H_1(s)(r_1(s)-\tilde{\sigma}(s)) + H_0(s)d_{\mathcal{L}_1}(s)\,. \tag{30}$$

Substitution of (29) and (30) into the system (5) results in:

$$\begin{aligned} y_1(s) = {} & M(s)\Big(C(s)+H_1(s)(1-C(s))\Big)\big(r_1(s)-\tilde{\sigma}(s)\big) \\ & + M(s)H_0(s)(1-C(s))d_{\mathcal{L}_1}(s)\,. \end{aligned}$$

From (28) and (27), this expression simplifies to:

$$y_1(s) = H(s)C(s)\big(r_1(s)-\tilde{\sigma}(s)\big) + H(s)(1-C(s))d_{\mathcal{L}_1}(s)\,. \tag{31}$$

An expression for y_2 is obtained by substituting (31) and (13) into $y_2(s) = \frac{1}{s}y_1(s)$ and making use of the definition in (8):

$$\begin{aligned} y_2(s) = {} & F(s)H(s)\Big(C(s)Kr_2(s) + (1-C(s))d_{\mathcal{L}_1}(s)\Big) \\ & - F(s)H(s)C(s)\tilde{\sigma}(s)\,. \end{aligned} \tag{32}$$

Substitution of (10) and (5) into the definition of $\tilde{y}$ in the adaptation law results in the following expression for $\tilde{y}(s)$:

$$\tilde{y}(s) = M(s)\tilde{\sigma}(s)\,. \tag{33}$$

Recalling the reference system in (14) and using the expression for y_2 in (32), the error between reference and actual systems, $y_{2,\text{ref}} - y_2$ is:

$$y_{2,\text{ref}}(s) - y_2(s) = F(s)H(s)\big(1 - C(s)\big)\big(d_{\text{ref}}(s) - d_{\mathcal{L}_1}(s)\big) - \frac{F(s)H(s)C(s)}{M(s)} M(s)\tilde{\sigma}(s)\,.$$

Substituting the expression for $\tilde{y}(s)$ in (33) and the definition of G(s) in (9), we obtain:

$$y_{2,\text{ref}}(s) - y_2(s) = G(s)(d_{\text{ref}}(s) - d_{\mathcal{L}_1}(s)) - \frac{F(s)H(s)C(s)}{M(s)}\tilde{y}(s)\,.$$

Finally, since the $\mathcal{L}_1$-norm of G(s) exists, and $\frac{F(s)H(s)C(s)}{M(s)}$ is strictly proper and stable, the following bound can be derived by taking the truncated $\mathcal{L}_\infty$-norm and by making use of Assumption 1:

$$\begin{aligned} \big\|y_{2,\text{ref}_t} - y_{2_t}\big\|_{\mathcal{L}_\infty} &\leq \big\|G(s)\big\|_{\mathcal{L}_1} L \big\|y_{2,\text{ref}_t} - y_{2_t}\big\|_{\mathcal{L}_\infty} \\ &\quad + \left\|\frac{F(s)H(s)C(s)}{M(s)}\right\|_{\mathcal{L}_1} \big\|\tilde{y}_t\big\|_{\mathcal{L}_\infty} \\ &\leq \frac{\left\|\frac{F(s)H(s)C(s)}{M(s)}\right\|_{\mathcal{L}_1}}{1 - \big\|G(s)\big\|_{\mathcal{L}_1} L} \big\|\tilde{y}_t\big\|_{\mathcal{L}_\infty}\,, \end{aligned}$$

which holds uniformly. From the bound in (18) proven in [7], the following bound is derived:

$$\big\|y_{2,\text{ref}} - y_2\big\|_{\mathcal{L}_\infty} \leq \frac{\left\|\frac{F(s)H(s)C(s)}{M(s)}\right\|_{\mathcal{L}_1}}{1 - \big\|G(s)\big\|_{\mathcal{L}_1} L} \gamma_0 = \gamma_1\,,$$

proving the second bound in (19). □

REFERENCES

[1] R. Skelton, “Model error concepts in control design,” *International Journal of Control*, vol. 49, no. 5, pp. 1725–1753, 1989.

[2] M. Morari and J. H. Lee, “Model predictive control: past, present and future,” *Computers & Chemical Engineering*, vol. 23, no. 4, pp. 667–682, 1999.

[3] S. Skogestad and I. Postlethwaite, *Multivariable feedback control: analysis and design*. Wiley New York, 2007, vol. 2.

[4] K. Barton, S. Mishra, and E. Xargay, “Robust iterative learning control: $\mathcal{L}_1$ adaptive feedback control in an ILC framework,” in *Proc. of the American Control Conference (ACC)*, 2011, pp. 3663–3668.

[5] B. Altin and K. Barton, “$\mathcal{L}_1$ adaptive control in an iterative learning control framework: Stability, robustness and design trade-offs,” in *Proc. of the American Control Conference (ACC)*, 2013, pp. 6697–6702.

[6] B. Altın and K. Barton, “Robust iterative learning for high precision motion control through $\mathcal{L}_1$ adaptive feedback,” *Mechatronics*, vol. 24, no. 6, pp. 549–561, 2014.

[7] N. Hovakimyan and C. Cao, *Adaptive Control Theory: Guaranteed Robustness with Fast Adaptation*. Philadelphia, PA: Society for Industrial and Applied Mathematics, 2010.

[8] S. Mallikarjunan, B. Nesbit, E. Kharisov, E. Xargay, N. Hovakimyan, and C. Cao, “$\mathcal{L}_1$ adaptive controller for attitude control of multirotors,” in *Proc. of the AIAA Guidance, Navigation and Control Conference*, 2012, p. 4831.

[9] B. Michini and J. P. How, “$\mathcal{L}_1$ adaptive control for indoor autonomous vehicles: Design process and flight testing,” in *Proc. of the AIAA Guidance, Navigation and Control Conference*, 2009, p. 5754.

[10] S. Gunnarsson and M. Norrlöf, “On the design of ILC algorithms using optimization,” *Automatica*, vol. 37, no. 12, pp. 2011–2016, 2001.

[11] C. J. Ostafew, A. P. Schoellig, and T. D. Barfoot, “Visual teach and repeat, repeat, repeat: Iterative learning control to improve mobile robot path tracking in challenging outdoor environments,” in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2013, pp. 176–181.

[12] D. Yu, Y. Zhu, K. Yang, C. Hu, and M. Li, “A time-varying Q-filter design for iterative learning control with application to an ultra-precision dual-stage actuated wafer stage,” *Proc. of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering*, vol. 228, no. 9, pp. 658–667, 2014.

[13] A. P. Schoellig and R. D’Andrea, “Optimization-based iterative learning control for trajectory tracking,” in *Proc. of the European Control Conference (ECC)*, 2009, pp. 1505–1510.

[14] F. L. Mueller, A. P. Schoellig, and R. D’Andrea, “Iterative learning of feed-forward corrections for high-performance tracking,” in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012, pp. 3276–3281.

[15] D. Bristow, M. Tharayil, and A. G. Alleyne, “A survey of iterative learning control,” *IEEE Control Systems*, vol. 26, no. 3, pp. 96–114, 2006.

[16] A. P. Schoellig, F. L. Mueller, and R. D’Andrea, “Optimization-based iterative learning for precise quadrocopter trajectory tracking,” *Autonomous Robots*, vol. 33, no. 1-2, pp. 103–127, 2012.

[17] N. Degen and A. P. Schoellig, “Design of norm-optimal iterative learning controllers: The effect of an iteration-domain Kalman filter for disturbance estimation,” in *Proc. of the IEEE Conference on Decision and Control (CDC)*, 2014, pp. 3590–3596.

[18] J. H. Lee, K. S. Lee, and W. C. Kim, “Model-based iterative learning control with a quadratic criterion for time-varying linear systems,” *Automatica*, vol. 36, no. 5, pp. 641–657, 2000.