Asian Journal of Control, Vol. 14, No. 3, pp. 613–623, May 2012
Published online 19 June 2011 in Wiley Online Library (wileyonlinelibrary.com) DOI: 10.1002/asjc.398

LIMITED BENEFIT OF JOINT ESTIMATION IN MULTI-AGENT ITERATIVE LEARNING

Angela P. Schoellig, Javier Alonso-Mora, and Raffaello D'Andrea

ABSTRACT

This paper studies iterative learning in a multi-agent framework, wherein a group of agents simultaneously and repeatedly perform the same task. Assuming similarity between the agents, we investigate whether exchanging information between the agents improves an individual's learning performance. That is, does an individual agent benefit from the experience of the other agents? We consider the multi-agent iterative learning problem as a two-step process of: first, estimating the repetitive disturbance of each agent; and second, correcting for it. We present a comparison of an agent's disturbance estimate in the case of (I) independent estimation, where each agent has access only to its own measurement, and (II) joint estimation, where information of all agents is globally accessible. When the agents are identical and noise comes from measurement only, joint estimation yields a noticeable improvement in performance. However, when process noise is encountered or when the agents have an individual disturbance component, the benefit of joint estimation is negligible.

Key Words: Multi-agent learning, iterative learning control, information exchange, estimation, Kalman filter.

I. INTRODUCTION

Exploiting previous experience when repeatedly executing the same task is a logical way to improve future performance in the presence of repetitive, unmodeled disturbances. Iterative learning control (ILC), as first proposed in [1], achieves precise tracking behavior by effectively incorporating past control information (such as applied input signals and measured outputs) when calculating the feedforward control action used in the next iteration [2, 3]. One way of viewing ILC is as a two-step process of estimation and control: first identifying the unknown repetitive disturbance and later compensating for it [4–9]. LQG-type solutions have been proposed in [10–13], which estimate the tracking error and, based on this result, calculate a new input trajectory by minimizing a quadratic cost function.

While ILC has proven to be successful in a variety of industrial applications (including chemical process control, rotary systems and robotics), we have yet to identify if – and how – ILC schemes can be generalized when facing homogeneous groups of agents or assemblies of similar units (for example, robot arms in an industrial environment, or a fleet of mobile robots in a warehouse [14, 15]). In other words, how can we cope with uncertainties in a multi-agent framework? Is there a benefit of exchanging information between the agents? What kind of information sharing makes sense? Cooperative iterative learning schemes were previously proposed in [16]. Recently, ILC was applied to multi-agent systems, cf. [17], with the goal of achieving formation control. While it has been established that the joint performance of all agents is fundamental to the formation problem, this

Manuscript received May 25, 2009; revised May 25, 2010; accepted February 12, 2011.

A. P. Schoellig and R. D'Andrea are with the Institute for Dynamic Systems and Control (IDSC), and J. Alonso-Mora is with the Institute of Robotics and Intelligent Systems (IRIS), ETH Zurich, 8092 Zurich, Switzerland (e-mail: aschoellig@ethz.ch, jalonso@ethz.ch, rdandrea@ethz.ch).

paper focuses on the potential for individual agents to improve their performance when conducting a task alongside a group of similar agents conducting the same task. Preliminary results were first published in [18]. Analogous questions were previously studied in the context of reinforcement learning (see [19]).

The results of our research show that the passing of information between agents has limited benefit for a large class of problems. This conclusion is based upon a comparison of independent learning with a cooperative scheme, where information about all agents is globally accessible to every other agent. Similarity between the agents is assured by assuming that they have the same nominal dynamics and share a common iteration-independent disturbance; however, they differ in an additional individual disturbance component that is also constant across iterations. We introduce iteration-dependent noise terms that account for measurement and process noise, and obtain results for two limit cases: (i) pure process noise, and (ii) pure measurement noise. The benefits of information sharing are negligible in (i). For (ii), we observe a greater improvement in performance when a high similarity between the agents is guaranteed. From this point, we are able to deduce the properties of the general mixed-noise case. In short, individual agents in an ILC framework do not, in most cases, benefit significantly from information sharing when simultaneously learning the same task.

The paper is organized as follows: Section II formalizes the multi-agent iterative learning problem and reduces it to a comparison of independent versus joint estimation. Section III compares both scenarios and presents the core result of the paper in terms of an upper bound on the performance improvement due to joint estimation. Several numerical examples are presented in Section IV visualizing the derived analytical results. The work is summarized in Section V, whereas proofs are partly presented in the Appendix.

II. PROBLEM STATEMENT

2.1 Motivation

We begin by considering a group of N agents that simultaneously and repeatedly perform the same task. A common way of describing an agent's dynamics during a single run is the so-called lifted system representation [20–22]. For each agent $i \in \mathscr{I} = \{1, 2, \ldots, N\}$, the input-state relationship is modeled by a static matrix equation,

$$x^i = F^i u^i + d^i, \tag{1}$$

which maps a given discrete-time input signal $u^i = [u^i(0), u^i(1), \ldots, u^i(T)]^{\mathrm{T}} \in \mathbb{R}^{(T+1)n_u}$ to the corresponding lifted states $x^i \in \mathbb{R}^{(T+1)n_x}$. In this context, $(T+1)$ samples represent a single iteration and n_u and n_x denote the dimension of the input and state, respectively. The vectors x^i and u^i are defined as the deviation from the desired task trajectory and the corresponding nominal input (see for example [9]). The vector d^i represents an exogenous disturbance constant across iterations, which captures model errors along the trajectory as well as repeating disturbances and nonzero initial conditions [3, 23, 24]. We include a non-repetitive noise signal ξ^i_j in model (1) to account for process noise, which varies from trial to trial. Introducing the iteration index $j \in \{1, 2, \ldots\}$, the state in the jth trial is given by

$$x^i_j = F^i u^i_j + d^i + \xi^i_j, \tag{2}$$

where ξ^i_j is assumed to be zero-mean Gaussian white noise. The vector d^i is viewed as an agent-dependent, normally-distributed random signal.

The agents' output y^i_j is corrupted by measurement noise and similarly represented in the lifted domain,

$$y^i_j = G^i x^i_j + \mu^i_j, \tag{3}$$

where μ^i_j is zero-mean Gaussian white noise. Similar to x^i_j and u^i_j, the output y^i_j is interpreted as the deviation from the nominal output trajectory.

Note that (2) and (3) might be the result of linearizing the agent dynamics about a desired task trajectory. Refer to [9, 25] for a more detailed derivation.

In the above context, the goal of the iterative learning algorithm is to make the state x^i_j (that is, the deviation from the desired task trajectory) small or, more precisely, to reduce x^i_j with an increasing number of iterations j. The performance of each individual agent is gradually improved by taking into account all information on previous iterations when estimating the disturbance vector d^i. As the accuracy of the disturbance estimate increases, a more appropriate open-loop input is determined, thereby compensating for the deficiencies in the modeled dynamics represented by the matrix F^i. From x^i_j, conclusions can be drawn as to the performance of execution j.

We now consider a homogeneous fleet of agents with the same nominal dynamics:

$$\begin{aligned} F^i &= F \\ G^i &= I \quad \forall i \in \mathscr{I}, \end{aligned} \tag{4}$$

where I denotes the identity matrix. That is, the state is assumed to be measured directly. Differences between the agents are captured in the disturbance vector d^i, which is composed of a common part d^0 that is identical for all agents, and an individual part $d^{i,ind}$,

$$d^i = d^0 + d^{i,ind} \quad \forall i \in \mathcal{I}. \tag{5}$$

In this context, the question arises: Does an individual agent benefit from sharing information with its companions? To what degree can the disturbance estimate d^i be improved by taking into account the measurements of the other agents?

2.2 Simplified model

Our main objective and central problem is to identify the disturbance d^i for each agent i in the presence of both process and measurement noise. Based on the disturbance estimate, it is possible to find the correcting input u^i_j that best compensates for the repetitive disturbance using a problem-specific optimization criterion (see for example [9]). Importantly, the correcting input u^i_j applied in each iteration is known. Focusing on the estimation problem, we consider a condensed form of the above multi-agent system representation (2)-(3),

$$x^i_j = d^i + \xi^i_j \tag{6}$$

$$y^i_j = x^i_j + \mu^i_j, \tag{7}$$

which features the key noise and disturbance characteristics, but omits the known part Fu^i_j without loss of generality. Equations (6) and (7) are summarized by

$$y^i_j = d^i + \upsilon^i_j, \tag{8}$$

where $\upsilon^i_j = \xi^i_j + \mu^i_j$ captures both process and measurement noise.

Moreover, assuming independence of the single entries in the vectors d^i and υ^i_j and identical noise characteristics, the problem reduces to the scalar case,

$$y^i_j = d^0 + d^{i,ind} + \upsilon^i_j, \tag{9}$$

where all variables are scalar valued. The probability distributions are given by

$$\begin{aligned} d^0 &\sim \mathcal{N}(0, \alpha) \\ d^{i,ind} &\sim \mathcal{N}(0, \beta) \\ \upsilon^i_j &\sim \mathcal{N}(0, 1), \quad \alpha, \beta \geq 0, \end{aligned} \tag{10}$$

where all quantities, υ^i_j, $i \in \mathcal{I}$, $j \in \{1, 2, \ldots\}$, $d^{i,ind}$, $i \in \mathcal{I}$, and d^0, are assumed to be mutually independent. The notation $\mathcal{N}(0, \alpha)$ represents a normal distribution with mean 0 and variance α. Note that in (10), the variance of the individual disturbance $d^{i,ind}$ is assumed to be identical for all agents $i \in \mathcal{I}$. Without loss of generality, the variances are normalized such that the variance of υ^i_j is 1. For the variances of the process and measurement noise, this results in

$$\begin{aligned} \xi^i_j &\sim \mathcal{N}(0, \lambda) \\ \mu^i_j &\sim \mathcal{N}(0, 1-\lambda), \quad 0 \leq \lambda \leq 1, \end{aligned} \tag{11}$$

assuming independence between ξ^i_j and μ^i_j. A value $\lambda = 1$ represents the case of encountering only process noise, whereas $\lambda = 0$ reflects the case where the noise is due to measurement only.

2.3 Independent vs. joint estimation

As the number of trials and measurements increases, more information about the system is collected, allowing an increasingly accurate estimate of the agents' constant noise terms d^i, $i \in \mathcal{I}$. Two limiting approaches might be taken when solving the estimation problem: (I) independent estimation, and (II) joint estimation.

In the case of independent estimation (I), each agent i individually estimates its disturbance d^i taking only its own measurements y^i_j, $j \in \{1, 2, \ldots\}$, into account. That is, information on the individually obtained measurements is not exchanged between the agents.

In the joint case (II), the acquired measurement data of each agent is made available to all other agents; that is, every agent receives the information about all other agents' measurements. Based on this global knowledge, we can design a joint estimation scheme that exploits the measurements of all agents and provides estimates d^i for every agent $i \in \mathcal{I}$. A vector D, reflecting the estimation objective in this case, is defined as: $D = [d^0, d^1, \ldots, d^N]^{\mathrm{T}} \in \mathbb{R}^{(N+1)}$. The measurements of all agents in the jth trial are combined in $Y_j = [y^1_j, y^2_j, \ldots, y^N_j]^{\mathrm{T}}$ and, analogously, the noise vector $V_j = [\upsilon^1_j, \upsilon^2_j, \ldots, \upsilon^N_j]^{\mathrm{T}}$ is introduced. Based on this representation, the joint estimation problem can be formulated as a Kalman filter problem, cf. [26, 27]:

$$\begin{aligned} D_j &= D_{j-1} \qquad \forall j \geq 1 \\ Y_j &= H D_j + V_j, \end{aligned} \tag{12}$$

where $H = [\mathbf{0}, I]$ is a matrix with zeros in the first column concatenated with an identity matrix of appropriate dimensions. The Kalman filter returns an unbiased state estimate $\widehat{D}_j$ for $j \geq 1$ that minimizes the trace of the error covariance matrix

$$P_j = E\left[(D_j - \widehat{D}_j)(D_j - \widehat{D}_j)^{\mathrm{T}}\right], \tag{13}$$

of trial j, taking measurements Y_m, $1 \leq m \leq j$, into account. $E[\cdot]$ denotes the expected value. The initial values are obtained from (10); in particular,

$$\widehat{D}_0 = [0, 0, \ldots, 0]^T \tag{14}$$

and the initial covariance matrix $P_0 = [p_0^{(k,l)}]$, $k, l \in \mathscr{K} = \{0, 1, \ldots, N\}$, is

$$P_0 = E\left[D_0 D_0^{\mathrm{T}}\right] \tag{15}$$

with

$$p_0^{(k,l)} = E\left[d^k d^l\right] = E\left[(d^0 + d^{k,ind})(d^0 + d^{l,ind})\right],$$

where $d^{0,ind} = 0$. Recalling the mutual independence of d^0 and $d^{i,ind}$ for all $i \in \mathscr{I}$, the initial covariance is given by

$$p_0^{(k,l)} = \begin{cases} \alpha + \beta & \text{for } k = l \geq 1 \\ \alpha & \text{otherwise.} \end{cases} \tag{16}$$

The variances α and β, which reflect the original noise characteristics (10), serve as initial values. Note that the above derivations do not place further assumptions or restrictions on how information is shared between agents; the information y_j^i of each agent is available to every other agent. In other words, we are investigating the ideal case of centralized, joint estimation within an optimal filtering context.

Equally important is that the independent estimation problem (I) is just a special case of the cooperative framework (II) with $N = 1$.

In both cases, (I) and (II), the variance of an individual's disturbance estimate at iteration j is given by

$$E\left[(d^i - \widehat{d}_j^i)^2\right] = p_j^{(i,i)} = p_j^{(1,1)}, \quad \forall i \in \mathscr{I}, \tag{17}$$

where $\widehat{D}_j = [\widehat{d}_j^i]$, $i \in \mathscr{I}$, and $P_j = [p_j^{(k,l)}]$, $k, l \in \mathscr{K}$. The variance is identical for all agents, since for each agent the same assumptions on the dynamics (9) and the initial noise characteristics (10) are made. The variance of an individual's disturbance (17) indicates the quality of the disturbance estimate. In the general case, (2)-(3), this value influences the effectiveness of the disturbance compensation, since the input update rule of the ILC algorithm is based on the current estimate $\widehat{d}_j^i$; for example, by a relation as follows (see [9]):

$$u_{j+1}^i = \arg\min_u \left|\left| F^i u + \widehat{d}_j^i \right|\right|. \tag{18}$$

Below, we distinguish between the individual disturbance variance $p_j^{(1,1)}$ in the case of joint and independent estimation, where the latter is given when evaluating $p_j^{(1,1)}$ for $N = 1$, i.e.

$$p_j^{(1,1)}\Big|_{N=1}. \tag{19}$$

Thus, the initial question can be reformulated: To what degree does joint estimation benefit the individual learning of an agent?

III. RESULT

We compared the learning performance based on (I) independent and (II) joint estimation, via the variance of the state x_j^i given all past measurements. This value indicates the accuracy of the tracking behavior in each iteration j. We investigated the benefits of information sharing and used, as our basis for the investigation, two limiting cases of (8): (i) encountering pure process noise, and (ii) dealing with measurement noise only. From these benchmark examples, we were able to deduce properties for the general mixed-noise case in Section IV and draw conclusions about the advantages of passing information in an ILC framework.

In order to compare the independent estimation result (I) with the joint estimation result (II), we derived an analytical expression for $p_j^{(1,1)}$.

Proposition 1. The error variance of an agent's disturbance $p_j^{(1,1)}$ can be expressed in terms of the initial variances α and β, the number of agents N, and the iteration j,

$$p_j^{(1,1)} = \frac{\alpha + \beta + j\beta^2 + jN\alpha\beta}{(1 + j\beta)(1 + j\beta + jN\alpha)}. \tag{20}$$

The result is obtained by solving the Kalman filter equations for (12) with initial conditions (14) and (16).

A detailed proof is found in the Appendix.

Next, we use the relation (20) to derive an upper bound on the performance improvement due to joint estimation. Two limiting cases are considered: (i) pure process noise and (ii) pure measurement noise.

3.1 Pure process noise

We assumed perfect measurements, i.e. $\mu_j^i = 0$ in (7) and $\lambda = 1$ in (11). The noise υ_j^i is interpreted as pure process noise, $\upsilon_j^i = \xi_j^i$. The performance of independent (I) vs. joint (II) estimation is analyzed through the variance of the state estimate. As mentioned in Section 2.1, the goal of ILC is to reduce the value x_j^i. This is achieved best if the variance in the estimate of x_j^i is small. That is, the variance of the state estimate can be used as a measure of learning performance. Given (6) and (10), the best estimate of the state $\widehat{x}_j^i$ at iteration j is equal to the current disturbance estimate $\widehat{d}_j^i$,

$$\widehat{x}_j^i = \widehat{d}_j^i, \tag{21}$$

since the noise ξ_j^i has zero mean. Recalling the noise characteristics (10) and the previous assumption of mutual independence between d^i and υ_j^i, we obtain the variance of state estimate from the sum of the variance of the estimate $\widehat{d}_j^i$ and the variance of ξ_j^i. That is, with (17) and (11),

$$\begin{aligned} E\left[(x_j^i - \widehat{x}_j^i)^2\right] &= E\left[(d^i + \xi_j^i - \widehat{d}_j^i)^2\right] \\ &= p_j^{(1,1)} + \lambda, \end{aligned} \tag{22}$$

where $\lambda = 1$ in the pure process noise case. We introduce the performance index (for the pure process noise case) as the ratio of the state variance in the independent case vs. the joint case,

$$R^{\text{proc}} = \frac{p_j^{(1,1)}\Big|_{N=1} + 1}{p_j^{(1,1)} + 1}, \tag{23}$$

using the notation of (19).

The following theorem can be stated:

Theorem 1. The bounds on the performance improvement due to joint estimation (vs. independent estimation) are given by

$$1 \le R^{\text{proc}} \le \frac{1+j}{j} \quad \forall \alpha, \beta, N, j, \tag{24}$$

where the best performance improvement occurs when $N \to \infty$, $\alpha \to \infty$, and $\beta = 0$. In this case, $R^{\text{proc}} = (1+j)/j$.

Interpretation of the result:

- The performance improvement due to joint estimation has an upper bound which is valid for all possible combinations of α, β, N, and j.
- Joint estimation is most beneficial if the agents' common disturbance component dominates and the individual noise component is negligible compared to the process noise; this corresponds to a large common noise variance α and a small individual component $\beta \ll 1$.
- The largest possible improvement in performance is a factor of 2, which is obtained only in the first iteration. With more iterations, the performance index rapidly decays to 1. In other words, the more often the agents perform a task, the less beneficial the exchange of information is.
- Intuitively, the result shows that if the agents are different, the measurements of the other agents do not provide significant information for an individual's performance improvement. If the agents are almost identical, however, 'averaging' the measurements of the agents via a joint estimation still has no 'visible' effect, since the process noise directly corrupts the value of interest, x_j^i; see (6).

Moreover, independent estimation and learning (I) is robust to uncertainties in the initial noise assumptions (10). Note that the variance of an individual's disturbance in the independent case depends solely on the sum $(\alpha + \beta)$, cf. (20) with $N = 1$. In other words, the assumption on how the disturbance d^i is decomposed in d^0 and $d^{i,ind}$, does not enter the result. It does, however, affect the joint estimation.

To conclude, there is little benefit of sharing information in the case of pure process noise.

Proof. Based on the closed-form representation in (20), Theorem 1 is proven by introducing R^{proc} as a function of j, α, β, and N,

$$R_j^{\text{proc}}(\alpha, \beta, N) = \frac{p_j^{(1,1)}(\alpha, \beta, 1) + 1}{p_j^{(1,1)}(\alpha, \beta, N) + 1}. \tag{25}$$

Recalling the properties

$$\alpha, \beta \ge 0 \quad \text{and} \quad j, N \in \{1, 2, \ldots\}, \tag{26}$$

we note that $p_j^{(1,1)}(\alpha, \beta, N) \ge 0$ for all possible arguments. By taking partial derivatives of $R_j^{\text{proc}}(\alpha, \beta, N)$, it can be shown that

$$\frac{\partial R_j^{\text{proc}}(\alpha, \beta, N)}{\partial N} \ge 0 \tag{27}$$

and $R_j^{\text{proc}}(\alpha, \beta, N)$ is bounded by

$$R_j^{\text{proc}}(\alpha, \beta, \infty) := \lim_{N \to \infty} R_j^{\text{proc}}(\alpha, \beta, N) \tag{28}$$

with

$$R_j^{\text{proc}}(\alpha, \beta, \infty) = \left(1 + \frac{\alpha + \beta}{1 + j(\alpha + \beta)}\right)\left(1 + \frac{\beta}{1 + j\beta}\right)^{-1}.$$

Secondly, it is shown that

$$\frac{\partial R_j^{\text{proc}}(\alpha, \beta, \infty)}{\partial \alpha} \ge 0 \tag{29}$$

with

$$R_j^{\text{proc}}(\infty, \beta, \infty) := \lim_{\alpha\to\infty} R_j^{\text{proc}}(\alpha, \beta, \infty)$$
$$= \left(1+\frac{1}{j}\right)\left(1+\frac{\beta}{1+j\beta}\right)^{-1},$$

that is $R_j^{\text{proc}}(\alpha, \beta, N) \le R_j^{\text{proc}}(\infty, \beta, \infty)$. Finally, with

$$\frac{\partial R_j^{\text{proc}}(\infty, \beta, \infty)}{\partial \beta} \ge 0, \tag{30}$$

and

$$R_j^{\text{proc}}(\infty, 0, \infty) = 1 + \frac{1}{j},$$

statement (24) is proven,

$$R_j^{\text{proc}}(\alpha, \beta, N) \le R_j^{\text{proc}}(\infty, 0, \infty)$$

for all α, β, N, and j. The lower bound is obtained for $N=1$, cf. (27). Matlab and Mathematica files for reproducing the results below are available at www.idsc.ethz.ch/Downloads/multiagentILC. □

3.2 Pure measurement noise

We studied the system properties under the assumption of zero process noise, i.e. $\xi_j^i = 0$ in (6) and $\lambda = 0$ in (11), and interpreted υ_j^i as pure measurement noise, $\upsilon_j^i = \mu_j^i$. Following the derivation (22), the ratio of the state variances (for the pure measurement noise case) is given by

$$R^{\text{meas}} = \frac{p_j^{(1,1)}\Big|_{N=1}}{p_j^{(1,1)}}. \tag{31}$$

The following theorem can be stated:

Theorem 2. The bounds on the performance improvement due to joint estimation (vs. independent estimation) are given by

$$1 \le R^{\text{meas}} \le N \quad \forall \alpha, \beta, N, j, \tag{32}$$

where the best performance improvement occurs when $\alpha \to \infty$ and $\beta = 0$, for all N, and j. In this case, $R^{\text{meas}} = N$.

Interpretation of the result:

- Again, an upper bound of the performance index is found which is valid for all possible combinations of α, β, N, and j. However, the upper bound does not depend on the number of iterations.
- Joint estimation is most beneficial if the agents' common disturbance component dominates and the individual noise component is negligible compared to the measurement noise; this corresponds to a large common noise variance α and a negligible individual component $\beta \ll 1$. The largest possible improvement in performance is a factor of N.
- Intuitively, the result shows that if the agents are very similar ($\beta \ll 1$), joint estimation has a 'visible' effect. The measurement noise is 'averaged out' and, moreover, does not corrupt the performance result, x_j^i, directly; see (7). A significant improvement in the individual's performance can be achieved.

Joint estimation is beneficial when considering a large group of almost identical agents, where the individual disturbance is small compared to the measurement noise.

Proof. The proof of Theorem 2 proceeds similarly to the proof in Section 3.1. With (20), the performance index R^{meas} is given as a function of j, α, β, and N,

$$R_j^{\text{meas}}(\alpha, \beta, N) = \frac{p_j^{(1,1)}(\alpha, \beta, 1)}{p_j^{(1,1)}(\alpha, \beta, N)}. \tag{33}$$

Partial derivatives are directly computed, where

$$\frac{\partial R^{\text{meas}}}{\partial \beta} \le 0, \quad \frac{\partial R^{\text{meas}}}{\partial \alpha} \ge 0, \quad \frac{\partial R^{\text{meas}}}{\partial N} \ge 0, \tag{34}$$

with (26). In addition, the limiting property for $\beta = 0$ is

$$R_j^{\text{meas}}(\alpha, 0, N) = \frac{1+\alpha j N}{1+\alpha j}$$

and

$$\lim_{\alpha\to\infty} R_j^{\text{meas}}(\alpha, 0, N) = N \quad \forall j, N.$$

The lower bound is obtained for $N=1$, cf. (34). Matlab and Mathematica files for reproducing the results are available at www.idsc.ethz.ch/Downloads/multiagentILC. □

IV. NUMERICAL EXAMPLES

Characteristic features of the performance indices, R^{proc} and R^{meas}, are highlighted by showing selected numerical examples. In the subsequent considerations, the general mixed-noise case is included and put in context with the results for (i) pure process and (ii) pure measurement noise.

The performance index for the mixed-noise case with $0 \leq \lambda \leq 1$, cf. (11), is derived analogously to (22) and (23):

$$R^{\mathrm{mix}} = \frac{p_j^{(1,1)}\Big|_{N=1} + \lambda}{p_j^{(1,1)} + \lambda}. \tag{35}$$

A comparison of the three performance indices, R^{proc}, R^{mix}, and R^{meas}, shows that

$$R^{\mathrm{proc}} \leq R^{\mathrm{mix}} \leq R^{\mathrm{meas}} \quad \forall \alpha, \beta, N, j, \tag{36}$$

since $\lambda \in [0, 1]$ and, from (20),

$$p_j^{(1,1)} \leq p_j^{(1,1)}\Big|_{N=1} \quad \forall \alpha, \beta, N, j. \tag{37}$$

An intuitive explanation for (36) is that the process noise ξ_j^i directly corrupts the value of interest, state x_j^i, which represents the deviation from the desired trajectory and is aimed to be small. The measurement noise μ_j^i acts on the output y_j^i. In this case, multiple agents are beneficial in order to average out the measurement noise.

In Figs 1 to 3, the evolution of the performance indices, R^{proc}, R^{mix}, and R^{meas}, is shown for different pairs of α and β. A group of $N = 10$ agents is considered and $\lambda = 0.1$ is chosen for the mixed-noise case. Even for this small value of λ, we observe a noticeable degradation of the performance index R^{mix} (compared with R^{meas}). Note that the scaling of the vertical axis changes in the plots presented. The figures highlight the relationship (36). For the limiting case $j \to \infty$, $\beta > 0$, we observe

$$\lim_{j \to \infty} R^{\mathrm{mix}} = 1. \tag{38}$$

Note that the mixed-noise case includes the special cases of pure process and pure measurement noise. In fact, if the number of iterations increases, the benefits of joint estimation become negligible. This result can be derived analytically from (20) and (35). As stated in Theorem 2, for the limiting case $\beta = 0$,

$$\lim_{j \to \infty} R^{\mathrm{meas}} = N. \tag{39}$$

Thus, we observe that joint estimation yields a significant improvement in performance only if the agents are identical, $\beta = 0$, and the system dynamics are not corrupted by process noise, see Fig. 1(a). Moreover, Fig. 1 shows how the behavior of R^{proc}, R^{mix}, and R^{meas} change, if the individual disturbance component β is gradually increased from zero for a given common disturbance α. The relation

$$\frac{\partial R^{\mathrm{meas}}}{\partial j} > 0 \quad \Leftrightarrow \quad \beta(\beta + \alpha N) < \frac{1}{j^2}, \tag{40}$$

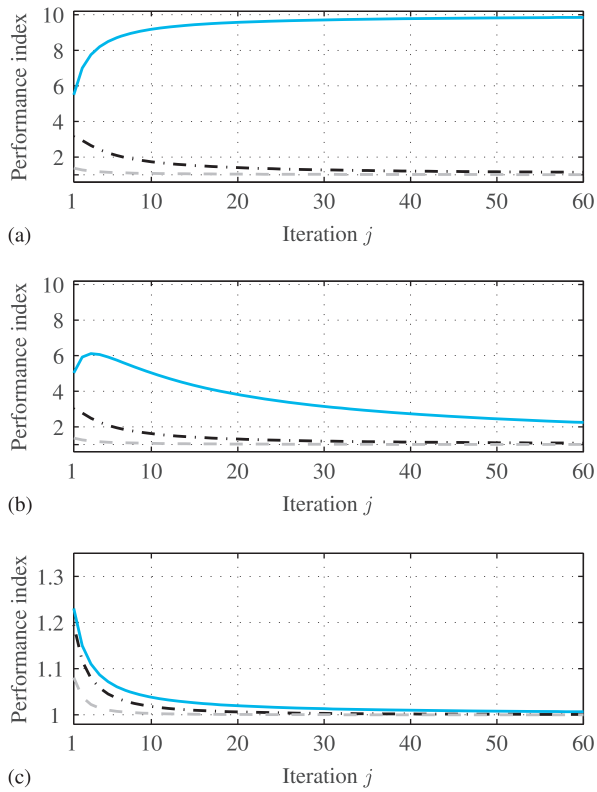

Fig. 1. Evolution of the performance indices for $N = 10$ agents: pure measurement noise R^{meas} (solid line), mixed noise case R^{mix} with $\lambda = 0.1$ (dashed-dotted line), and pure process noise R^{proc} (dashed line): (a) $\alpha = 1$, $\beta = 0$; (b) $\alpha = 1$, $\beta = 0.01$; and (c) $\alpha = 1$, $\beta = 1$.

derived from (20) and (31), provides insight in the evolution of the performance index R^{meas}. Figure 1 shows that the performance indices, which represent the performance improvement due to joint estimation, decrease with larger values β. The relation,

$$\frac{\partial R^{\mathrm{mix}}}{\partial \beta} \leq 1, \tag{41}$$

is derived from (20) and (35), where the mixed-noise case includes the special cases of pure process and pure measurement noise. The variances of the common and individual disturbance component, α and β, must be interpreted in relation to the variance of the noise υ_j^i which is normalized to 1; see (10). Figure 1(c) shows the evolution of the performance indices if the disturbance variances have the same value as the noise variance. If both disturbance variances, α and β, are smaller than the noise variance, a behavior as shown in Fig. 2(a) is obtained. Figure 2(b) shows both disturbance variances being larger than the noise variance.

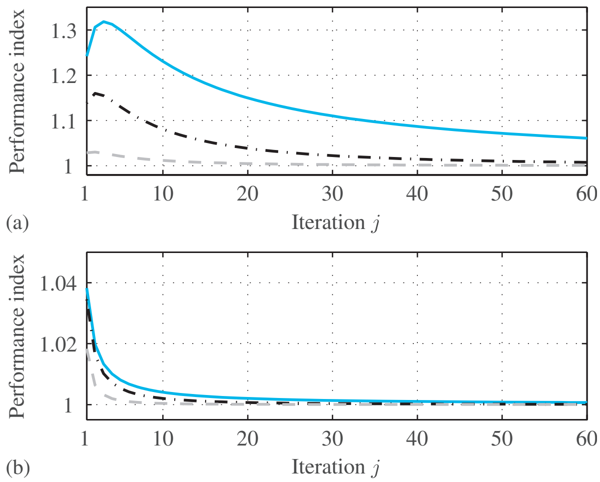

Fig. 2. Evolution of the performance indices for $N=10$ agents: pure measurement noise R^{meas} (solid line), mixed noise case R^{mix} with $\lambda=0.1$ (dashed-dotted line), and pure process noise R^{proc} (dashed line): (a) $\alpha=0.1$, $\beta=0.1$ and (b) $\alpha=10$, $\beta=10$.

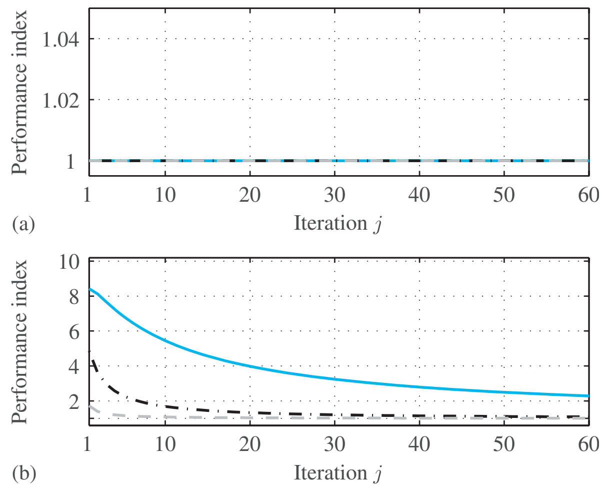

Fig. 3. Evolution of the performance indices for $N=10$ agents: pure measurement noise R^{meas} (solid line), mixed noise case R^{mix} with $\lambda=0.1$ (dashed-dotted line), and pure process noise R^{proc} (dashed line): (a) $\alpha=0.01$, $\beta=10$ and (b) $\alpha=10$, $\beta=0.01$.

In Fig. 3, the two cases, $\alpha \ll 1$, $\beta \gg 1$ and $\beta \ll 1$, $\alpha \gg 1$, are depicted. Figure 3(a) underlines the fact that, in the general case,

$$\lim_{\alpha \to 0} R^{\text{mix}} = 1, \quad \lim_{\beta \to \infty} R^{\text{mix}} = 1. \tag{42}$$

In contrast, information exchange and joint estimation is most beneficial when $\alpha \to \infty$ and $\beta = 0$; see Theorems 1 and 2. Fig. 3(b) shows a corresponding setting.

V. CONCLUSION

In this paper we considered a group of agents which share the same dynamics and a common iteration-independent disturbance, but differ in an additional individual error component. In the context of having these agents learn to perform an identical task, we asked: How beneficial is it to exchange experience in order to improve an individual agent's learning performance? We considered two cases: (I) independent learning without information exchange and (II) learning based on full information exchange between agents. In the proposed framework, the question can be reduced to the comparison of the disturbance estimate in the case of independent estimation (I) and when solving a global estimation problem for (II). An upper bound for the performance improvement due to information exchange is derived analytically and reflects the limited benefit of sharing information in the given setup. In the best case – where the noise is due to measurement noise only, the agent's common disturbance dominates, and the individual disturbance component is small compared to the noise – joint estimation improves the performance by a factor equal to the number of agents. That is, instead of one agent performing a task N times, N agents performing the task once results in the same accuracy for the disturbance estimate. For the general case and, in particular, in the presence of process noise or a large individual disturbance component, the benefits are shown to be limited.

APPENDIX A

We derive an explicit representation of the variance $p_j^{(1,1)}$ that depends only on α, β, j, and N as presented in Proposition 1. Matlab and Mathematica files for reproducing the results below are available at www.idsc.ethz.ch/Downloads/multiagentILC.

Proof. A closed form of the covariance matrix P_j is derived, cf. (16) and (17). Since noise is assumed to have the same characteristics for each agent, by symmetry,

$$p_j^{(k,l)} = \begin{cases} p_j^{(0,0)} & \text{if } k=l=0 \\ p_j^{(0,1)} & \text{if } kl=0 \text{ and } k \neq l \\ p_j^{(1,1)} & \text{if } k=l \neq 0 \\ p_j^{(1,2)} & \text{otherwise.} \end{cases} \tag{A1}$$

We obtain the previous values by solving the filter equations, cf. [26, 27],

$$\begin{aligned} Q_j &= HP_{j-1}H^T + I \\ K_j &= P_{j-1}H^T Q_j^{-1} \\ P_j &= (I - K_j H)P_{j-1}, \end{aligned} \tag{A2}$$

where $Q_j = [q_j^{(k,l)}]$, $k, l \in \mathscr{I}$ and $K_j = [k_j^{(k,l)}]$, $k \in \mathscr{K}$, $l \in \mathscr{I}$. With (A1) and (A2), the matrix Q_j and its inverse $Q_j^{-1} = [m_j^{(k,l)}]$ are directly computed,

$$\begin{aligned} q_j^{(k,l)} &= \begin{cases} 1 + p_{j-1}^{(1,1)} & \text{if } k = l \\ p_{j-1}^{(1,2)} & \text{otherwise} \end{cases} \\ m_j^{(k,l)} &= \begin{cases} m_j^{(1,1)} & \text{if } k = l \\ m_j^{(1,2)} & \text{otherwise,} \end{cases} \end{aligned} \tag{A3}$$

where

$$\begin{aligned} m_j^{(1,1)} &= \frac{1 + p_{j-1}^{(1,1)} + (N-2)p_{j-1}^{(1,2)}}{n_1 n_2} \\ m_j^{(1,2)} &= \frac{-p_{j-1}^{(1,2)}}{n_1 n_2} \end{aligned} \tag{A4}$$

with

$$\begin{aligned} n_1 &= \left(1 + p_{j-1}^{(1,1)} - p_{j-1}^{(1,2)}\right) \\ n_2 &= \left(1 + p_{j-1}^{(1,1)} + (N-1)p_{j-1}^{(1,2)}\right). \end{aligned} \tag{A5}$$

With this, the filtering matrix K_j is given by

$$k_j^{(k,l)} = \begin{cases} k_j^{(0,1)} & \text{if } k = 0 \\ k_j^{(1,1)} & \text{if } k = l \\ k_j^{(1,2)} & \text{otherwise,} \end{cases} \tag{A6}$$

where

$$\begin{aligned} k_j^{(0,1)} &= p_{j-1}^{(0,1)} \left(m_j^{(1,1)} + (N-1)m_j^{(1,2)}\right) \\ k_j^{(1,1)} &= p_{j-1}^{(1,1)} m_j^{(1,1)} + (N-1)p_{j-1}^{(1,2)} m_j^{(1,2)} \\ k_j^{(1,2)} &= p_{j-1}^{(1,1)} m_j^{(1,2)} + p_{j-1}^{(1,2)} m_j^{(1,1)} \\ &\quad + (N-2)p_{j-1}^{(1,2)} m_j^{(1,2)}. \end{aligned} \tag{A7}$$

From (A2), the following values for P_j are found,

$$p_j^{(k,l)} = \begin{cases} p_j^{(0,0)} & \text{if } k = l = 0 \\ p_j^{(0,1)} & \text{if } kl = 0 \text{ and } l \neq k \\ p_j^{(1,1)} & \text{if } k = l \neq 0 \\ p_j^{(1,2)} & \text{otherwise} \end{cases} \tag{A8}$$

with

$$\begin{aligned} p_j^{(0,0)} &= p_{j-1}^{(0,0)} - Np_{j-1}^{(0,1)} k_j^{(0,1)} \\ p_j^{(0,1)} &= p_{j-1}^{(0,1)} + p_{j-1}^{(1,1)} k_j^{(0,1)} \\ &\quad - (N-1)\, p_{j-1}^{(1,2)} k_j^{(0,1)} \\ p_j^{(1,1)} &= (1 - k_j^{(1,1)})\, p_{j-1}^{(1,1)} \\ &\quad - (N-1)\, p_{j-1}^{(1,2)} k_j^{(1,2)} \\ p_j^{(1,2)} &= (1 - k_j^{(1,1)})\, p_{j-1}^{(1,2)} - k_j^{(1,2)} p_{j-1}^{(1,1)} \\ &\quad - (N-2)p_{j-1}^{(1,2)} k_j^{(1,2)}. \end{aligned} \tag{A9}$$

We prove the desired symmetry and obtain the following values for (A1) by induction, using (A8) with starting condition (16):

$$\begin{aligned} p_j^{(0,0)} &= \frac{(1 + j\beta)\alpha}{1 + j\beta + jN\alpha} \\ p_j^{(0,1)} &= \frac{\alpha}{1 + j\beta + jN\alpha} \\ p_j^{(1,1)} &= \frac{\alpha + \beta + j\beta^2 + jN\alpha\beta}{(1 + j\beta)(1 + j\beta + jN\alpha)} \\ p_j^{(1,2)} &= \frac{\alpha}{(1 + j\beta)(1 + j\beta + jN\alpha)}. \end{aligned} \tag{A10}$$

The only value of interest is $p_j^{(1,1)}$. □

REFERENCES

1. Arimoto, S., S. Kawamura, and F. Miyazaki, "Bettering operation of robots by learning," *J. Robot. Syst.*, Vol. 1, No. 2, pp. 123–140 (1984).
2. Bien, Z. and J. X. Xu, *Iterative Learning Control*: *Analysis*, *Design*, *Integration and Applications*, Kluwer Academic Publishers, Norwell, MA (1998).
3. Bristow, D. A., M. Tharayil, and A. G. Alleyne, "A survey of iterative learning control," *IEEE Control Syst. Mag.*, Vol. 26, No. 3, pp. 96–114 (2006).
4. Moore, K. L., "An iterative learning control algorithm for systems with measurement noise," *Proc. 38th IEEE Conf. Decis. Control*, Phoenix, AZ, Vol. 1, pp. 270–275 (1999).
5. Lee, J. H., 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).
6. Phan, M. and R. Longman, "Higher-order iterative learning control by pole placement and noise filtering," *Proc. IFAC World Congr.*, Barcelona, Spain, pp. 1899–1904 (2002).
7. Norrlöf, M., "An adaptive iterative learning control algorithm with experiments on an industrial robot," *IEEE Trans. Robot. Autom.*, Vol. 18, No. 2, pp. 245–251 (2002).
8. Norrlöf, M., "Disturbance rejection using an ILC algorithm with iteration varying filters," *Asian J. Control*, Vol. 6, pp. 432–438 (2004).
9. Schoellig, A. P. and R. D'Andrea, "Optimization-based iterative learning control for trajectory tracking," *Proc. European Control Conf.*, Budapest, Hungary, pp. 1505–1510 (2009).
10. Lee, K. S., J. Lee, I. Chin, J. Choi, and J. H. Lee, "Control of wafer temperature uniformity in rapid thermal processing using an optimal iterative learning control technique," *Ind. Eng. Chem. Res.*, Vol. 40, No. 7, pp. 1661–1672 (2001).
11. Chin, I., S. J. Qin, K. S. Lee, and M. Cho, "A two-stage iterative learning control technique combined with real-time feedback for independent disturbance rejection," *Automatica*, Vol. 40, No. 11, pp. 1913–1922 (2004).
12. Cho, M., Y. Lee, S. Joo, and K. S. Lee, "Semi-empirical model-based multivariable iterative learning control of an RTP system," *IEEE Trans. Semicond. Manuf.*, Vol. 18, No. 3, pp. 430–439 (2005).
13. Rice, J. K. and M. Verhaegen, "A structured matrix approach to efficient calculation of LQG repetitive learning controllers in the lifted setting," *Int. J. Control*, Vol. 83, No. 6, pp. 1265–1276 (2010).
14. Guizzo, E., "Three engineers, hundreds of robots, one warehouse," *IEEE Spectr.*, Vol. 45, No. 7, pp. 22–29 (2008).
15. Wurman, P., R. D'Andrea, and M. Mountz, "Co-ordinating hundreds of cooperative, autonomous vehicles in warehouses," *AI Mag.*, Vol. 29, No. 1, pp. 9–19 (2008).
16. Ahn, H.-S., K. L. Moore, and Y. Chen, *Iterative Learning Control*: *Robustness and Monotonic Convergence for Interval Systems* (*Communications and Control Engineering*), Springer, London, UK, (2007).
17. Ahn, H.-S. and Y. Chen, "Iterative learning control for multi-agent formation," *Proc. Int. Joint Conf. Control Autom. Syst. Soc. Instrum. Control Eng. Inst. Control*, *Robot. Syst.*, Fukuoka, Japan, pp. 3111–3116 (2009).
18. Schoellig, A. P., J. Alonso-Mora, and R. D'Andrea, "Independent vs. joint estimation in multi-agent iterative learning control," *Proc. 49th IEEE Conf. Decis. Control*, pp. 6949–6954 (2010).
19. Tan, M., "Multi-agent reinforcement learning: independent vs. cooperative agents," *Proc. 10th Int. Conf. Mach. Learn.*, Amherst, MA, pp. 330–337 (1993).
20. Phan, M. Q. and R. W. Longman, "A mathematical theory of learning control for linear discrete multivariable systems," *Proc. AIAA/AAS Astrodynamics Conf.*, Boston, MA, pp. 740–746 (1988).
21. Tousain, R., E. van der Meche, and O. Bosgra, "Design strategy for iterative learning control based on optimal control," *Proc. 40th IEEE Conf. Decis. Control*, Orlando, FL, Vol. 5, pp. 4463–4468 (2001).
22. Bamieh, B., J. B. Pearson, B. A. Francis, and A. Tannenbaum, "A lifting technique for linear periodic systems with applications to sampled-data control," *Syst. Control Lett.*, Vol. 17, No. 2, pp. 79–88 (1991).
23. Hätönen, J., D. Owens, and K. Feng, "Basis functions and parameter optimisation in high-order iterative learning control," *Automatica*, Vol. 42, No. 2, pp. 287–294 (2006).
24. Barton, K. and A. Alleyne, "A cross-coupled iterative learning control design for precision motion control," *IEEE Trans. Control Syst. Technol.*, Vol. 16, No. 6, pp. 1218–1231 (2008).
25. Butcher, M., A. Karimi, and R. Longchamp, "Iterative learning control based on stochastic approximation," *Proc. 17th IFAC World Congr.*, Seoul, Korea, pp. 1478–1483 (2008).
26. Chui, C. K. and G. Chen, *Kalman Filtering*: *with Real-Time Applications* (Springer Series in Information Sciences), Springer, Heidelberg, Germany (1998).
27. Verhaegen, M. and V. Verdult, *Filtering and System Identification*: *A Least Squares Approach*. Cambridge University Press, Cambridge, MA (2007).

Angela P. Schoellig is a Ph.D. candidate in Dynamic Systems and Control at ETH Zurich. She graduated with a M.Sc. in Engineering Science and Mechanics from the Georgia Institute of Technology in 2007 and received her Masters degree in Engineering Cybernetics from the University of Stuttgart in 2008. Her past research

experience includes optimal control of hybrid systems and stability analysis of networked systems with communication delays. She has worked at the European Aeronautic Defence and Space Company (Germany), where she developed distributed estimation and control algorithms for satellites. For her graduate studies, Angela received scholarships from the German National Academic Foundation, the Cusanuswerk, and the German Academic Exchange Service. She was finalist of the 2008 IEEE Fellowship in Robotics and Automation, which supports prospective leaders in this field.

Javier Alonso-Mora received his Diploma in Mathematics and his M.Sc. in Engineering from the Technical University of Catalonia UPC in 2008 and 2010, respectively. In 2010 he graduated with honours from ETH Zurich with his M.Sc. in Robotics, Systems and Control. During his studies, he has received scholarships from CSIC (Spain), CFIS (Spain), DAAD (Germany) and ESKAS (Switzerland). He is currently a Ph.D. Candidate at the Institute of Robotics and Intelligent Systems at ETH Zurich, and is associated with Disney Research Zurich.

Raffaello D'Andrea received his B.Sc. degree in Engineering Science from the University of Toronto in 1991, and his M.S. and Ph.D. degrees in Electrical Engineering from the California Institute of Technology in 1992 and 1997, respectively. He held positions as Assistant Professor, and, later, Associate Professor at Cornell University from 1997 to 2007. He is currently a Full Professor of Dynamic Systems and Control at ETH Zurich, and Technical Co-founder and Chief Technical Advisor at Kiva Systems. A creator of dynamic sculpture, his work has appeared at various international venues, including the National Gallery of Canada, the Venice Biennale, Ars Electronica, and ideaCity.