

Sensitivity of Joint Estimation in Multi-Agent Iterative Learning Control

Angela Schöllig* Raffaello D'Andrea*

* *Institute for Dynamic Systems and Control (IDSC), ETH Zurich,
Sonneggstr.3, 8092 Zurich, Switzerland (e-mail: aschoellig@ethz.ch,
rdandrea@ethz.ch).*

Abstract: We consider a group of agents that simultaneously learn the same task, and revisit a previously developed algorithm, where agents share their information and learn jointly. We have already shown that, as compared to an independent learning model that disregards the information of the other agents, and when assuming similarity between the agents, a joint algorithm improves the learning performance of an individual agent. We now revisit the joint learning algorithm to determine its sensitivity to the underlying assumption of similarity between agents. We note that an incorrect assumption about the agents' degree of similarity degrades the performance of the joint learning scheme. The degradation is particularly acute if we assume that the agents are more similar than they are in reality; in this case, a joint learning scheme can result in a poorer performance than the independent learning algorithm. In the worst case (when we assume that the agents are identical, but they are, in reality, not) the joint learning does not even converge to the correct value. We conclude that, when applying the joint algorithm, it is crucial not to overestimate the similarity of the agents; otherwise, a learning scheme that is independent of the similarity assumption is preferable.

Keywords: Agents, learning control, estimation, Kalman filters, sensitivity analysis.

1. INTRODUCTION

In most cases, multi-agent learning aims towards improving the joint performance of a group of agents that solve a complex or distributed task together. Through interaction and collaboration, the agents are able to jointly approach the common task and learn to work together in order to achieve the global objectives, cf. Panait and Luke [2005]. Reinforcement learning is a powerful tool to solve such cooperative problems. Here, agents are generally categorized to be either homogeneous or heterogeneous, see e.g. Schaefer et al. [1994], Mataric [1997]. Though the robustness of such learning schemes to parameter variations was suggested in recent publications (see Mannor and Shamma [2007], Morimoto and Doya [2005]), it has yet to be studied in detail.

In this paper, we focus on the potential for an individual agent to improve its performance when conducting a task alongside a group of similar agents conducting the same task. We extend the work in Schöllig et al. [2010]¹, where the learning performance of an individual agent is analyzed and compared for two scenarios: (i) the agent learns independently, disregarding the information of the other agents, and (ii) the agent has access to the knowledge of the other agents and optimally takes this information into account when learning the desired task. An iterative learning control (ILC) scheme was used to approach this problem (see Bien and Xu [1998], Bristow et al. [2006] for an introduction to ILC). ILC has been viewed as a




Fig. 1. Independent (left) vs. joint (right) estimation and learning.

two-step process of first identifying the unknown repetitive disturbances that corrupt the agent's performance and later compensating for the disturbances by adapting the input, cf. Phan and Longman [2002], Norrlöf [2004], Schöllig and D'Andrea [2009]. This scheme allowed us to reduce the previous question to an estimation problem, and we were able to show that, when assuming similarity between the agents, a joint estimation scheme that exploits the information of all agents is always beneficial.

The results in Schöllig et al. [2010] were obtained under the assumption that we know the degree of similarity between the agents; the goal of this work is to study the sensitivity of joint estimation to the underlying similarity assumption. We analyze the effects of an assumption error on the joint estimation result in order to determine whether it is possible that an incorrect assumption regarding the agents' similarity cause the joint estimation scheme to

¹ Paper and additional material may be found on the project webpage: www.idsc.ethz.ch/Downloads/multiagentILC.

perform more poorly than the independent estimation scheme.

The paper is organized as follows: Sec. 2 recalls the dynamic equations that define the multi-agent iterative learning problem as previously derived in Schöllig et al. [2010], and introduces the parameter which defines the similarity between agents. Sec. 3 solves the estimation problem under the assumption that the degree of similarity between the agents is not known precisely. Using the results from Sec. 3 as a basis, the sensitivity of the joint estimation scheme to the assumption errors is studied in Sec. 4. To help readers visualize the analytical results derived herein, we present several numerical examples in Sec. 5. The work is summarized in Sec. 6. Proofs are presented in Appendix A, with additional files available at www.idsc.ethz.ch/Downloads/multiagentILC.

2. PROBLEM STATEMENT

We extend the work in Schöllig et al. [2010], where we considered a group of N agents that simultaneously and repeatedly perform the same task. The execution of the task is corrupted by an unknown, repetitive disturbance that is constant across iterations. In this context, we assume that the agents are similar in the sense that they have the same nominal dynamics and share a common iteration-independent, repetitive disturbance component. In addition, process noise acts on the agent's dynamics, and varies from trial to trial. Our goal is to improve the agents' performance by estimating the repetitive disturbance from past measurements; once the disturbance is known, an adapted input trajectory can be created to compensate for it.

2.1 Agent Dynamics

The dynamics of an agent $i \in \mathcal{I} = \{1, 2, \dots, N\}$ during a single execution of the task are represented in the lifted domain, cf. Phan and Longman [1988], Tousain et al. [2001], Bamieh et al. [1991]. A given discrete-time input signal

$$u_j^i = (u^i(0), u^i(1), \dots, u^i(T)) \quad (1)$$

at iteration $j \in \{1, 2, \dots\}$ is mapped to the corresponding lifted states x^i via a constant matrix F , which represents the nominal dynamics of the agents,

$$x_j^i = F u_j^i + d^i + \xi_j^i, \quad (2)$$

In this context, $(T + 1)$ samples represent a single run. The vector d^i represents the repetitive disturbance and ξ_j^i accounts for the trial-uncorrelated process noise. The vectors x_j^i and u^i are defined as the deviation from the desired task trajectory and the corresponding nominal input, see for example Schöllig and D'Andrea [2009]. The agents' output y_j^i (also defined as the deviation from the nominal output) is corrupted by measurement noise and similarly represented in the lifted domain,

$$y_j^i = x_j^i + \mu_j^i. \quad (3)$$

Differences between the agents are captured in the disturbance vector d^i , which is composed of a common part d^0 identical for all agents, and an individual part $d^{i,ind}$,

$$d^i = d^0 + d^{i,ind} \quad \forall i \in \mathcal{I}. \quad (4)$$

For a more detailed introduction to the lifted system representation refer to Bristow et al. [2006], Hätonen et al. [2006], Barton and Alleyne [2008], Butcher et al. [2008]

In the above context, the goal of the iterative learning algorithm is to reduce x_j^i (that is, the deviation from the desired task trajectory) with an increasing number of iterations j . In Schöllig et al. [2010], we showed that the learning problem can be divided into two steps: (i) estimating the disturbance vector d^i based on all measurements from previous iterations, and (ii) determining an appropriate open-loop input for the next trial that compensates for the disturbance, see Fig. 1. We saw that the characteristics of a joint learning scheme can be studied by focusing on the estimation problem; compensating input for each agent is found by solving an optimization problem for each agent once the disturbance estimate of d^i is updated, cf. Fig. 1.

2.2 Simplified Model

Focusing on the estimation problem, we consider a condensed form of the above multi-agent system representation (2)-(3),

$$x_j^i = d^i + \xi_j^i \quad (5)$$

$$y_j^i = x_j^i + \mu_j^i, \quad (6)$$

which features the key noise and disturbance characteristics, but omits the *known* part $F u_j^i$, without loss of generality. Equations (5) and (6) are summarized by

$$y_j^i = d^i + v_j^i, \quad (7)$$

where $v_j^i = \xi_j^i + \mu_j^i$ captures both process and measurement noise.

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

$$y_j^i = d^0 + d^{i,ind} + v_j^i, \quad (8)$$

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

$$\begin{aligned} d^0 &\sim \mathcal{N}(0, \alpha) \\ d^{i,ind} &\sim \mathcal{N}(0, \beta) \end{aligned} \quad (9)$$

$$v_j^i \sim \mathcal{N}(0, 1), \quad \alpha, \beta \geq 0,$$

where all quantities, v_j^i , $i \in \mathcal{I}$, $j \in \{1, 2, \dots\}$, $d^{i,ind}$, $i \in \mathcal{I}$, 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 (9), 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 v_j^i is 1. For the variances of the process and measurement noise, this means

$$\begin{aligned} \xi_j^i &\sim \mathcal{N}(0, \lambda) \\ \mu_j^i &\sim \mathcal{N}(0, 1 - \lambda), \quad 0 \leq \lambda \leq 1 \end{aligned} \quad (10)$$

assuming independence between ξ_j^i and μ_j^i . 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 Similarity Assumption

In our previous work Schöllig et al. [2010], we assumed that the variances of the individual and common disturbance, α and β , are known. In reality, however, these values are difficult to determine. While the sum ($\alpha + \beta$) may be approximated with reasonable accuracy (it indicates the magnitude of the agent's disturbance d^i , or more precisely, the probability of having larger values for d^i), a prior partitioning of the disturbance d^i into an individual and a common component is almost impossible. In other words, when facing a real multi-agent learning problem, the determined ratio between α and β , is subject to error.

For the subsequent analysis, we assume that the sum,

$$\gamma = \alpha + \beta, \quad (11)$$

is known precisely. With respect to the partitioning of d^i into the individual and common disturbance component, we distinguish between the nominal values,

$$\bar{\alpha} = \bar{\epsilon} \gamma \quad \text{and} \quad \bar{\beta} = (1 - \bar{\epsilon}) \gamma, \quad (12)$$

and the real variances α and β , defined analogously by ϵ , where $0 \leq \bar{\epsilon}, \epsilon \leq 1$. The nominal values represent our assumption on the individual and common disturbance component. The real ratio ϵ of the multi-agent system is unknown. The assumption error δ defines the difference between the real disturbance ratio and our assumed partitioning,

$$\delta = \epsilon - \bar{\epsilon}. \quad (13)$$

Below we study the effects of the assumption error δ on the performance of the joint learning algorithm. Our goal is to determine the degree to which joint estimation is affected by incorrect assumptions of similarity between agents.

3. ESTIMATION PROBLEM

Analogously to Schöllig et al. [2010], we consider two limiting approaches when solving the estimation problem: (I) independent estimation, and (II) joint estimation, see Fig. 1.

In the case of independent estimation (I), each agent i individually estimates its disturbance d^i , taking only its own measurements y_j^i , $j \in \{1, 2, \dots\}$, into account.

In the joint case (II), every agent has access to the measurements of all other agents. 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 , which reflects the estimation objective in this case, is defined as: $D = (d^0, d^1, \dots, d^N) \in \mathbb{R}^{(N+1)}$. The measurements of all agents in the j th trial are combined in $Y_j = (y_j^1, y_j^2, \dots, y_j^N)$, and analogously, the noise vector is $V_j = (v_j^1, v_j^2, \dots, v_j^N)$. Based on this representation, the joint estimation problem can be formulated as a Kalman filter problem, cf. Chui and Chen [1998], Verhaegen and Verdult [2007]:

$$\begin{aligned} D_j &= D_{j-1} & \forall j \geq 1 \\ Y_j &= H D_j + V_j, \end{aligned} \quad (14)$$

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 \hat{D}_j for $j \geq 1$ that minimizes the error covariance matrix

$$S_j = E \left[(D_j - \hat{D}_j)(D_j - \hat{D}_j)^T \right], \quad (15)$$

of trial j , taking measurements Y_m , $1 \leq m \leq j$, into account. $E[\cdot]$ denotes the expected value.

The recursive algorithm is based on the stochastic characteristics of the noise terms v_j^i , defined by (9), and relies on a given initial covariance matrix S_0 , which reflects the characteristics of the disturbances d^i . The initial disturbance estimate is obtained from (9),

$$\hat{D}_0 = (0, 0, \dots, 0), \quad (16)$$

and the initial covariance $S_0 = [s_0^{(k,l)}]$, $k, l \in \mathcal{K} = \{0, 1, \dots, N\}$ is given by

$$S_0 = E [D_0 D_0^T] \quad (17)$$

and with (4),

$$s_0^{(k,l)} = E [d^k d^l] = E [(d^0 + d^{k,ind}) (d^0 + d^{l,ind})], \quad (18)$$

where $d^{0,ind} = 0$.

When solving the filter equations, we distinguish between the real variance values of the system denoted by α, β and the nominal values $\bar{\alpha}, \bar{\beta}$ that represent our *assumption* on the individual and common disturbance component, see Sec. 2.3. The Kalman filter derivation below is based on the nominal values $\bar{\alpha}, \bar{\beta}$. The real values α, β are unknown and difficult to identify a priori. Note that in Schöllig et al. [2010], we derived the estimation problem under the assumption $\alpha = \bar{\alpha}$ and $\beta = \bar{\beta}$.

The Kalman filter proceeds in two steps:

Step 1 The Kalman gains K_j are calculated prior to the experiment based on the nominal values $\bar{\alpha}, \bar{\beta}$ by solving the filter equations

$$\begin{aligned} Q_j &= H S_{j-1} H^T + I \\ K_j &= S_{j-1} H^T Q_j^{-1} \\ S_j &= (I - K_j H) S_{j-1} \end{aligned} \quad (19)$$

with initial covariance S_0 , see (17) and (18). Recalling the mutual independence of d^0 and $d^{i,ind}$ for all $i \in \mathcal{I}$, the initial covariance is given by

$$s_0^{(k,l)} = \begin{cases} \bar{\alpha} + \bar{\beta} & \text{for } k = l \geq 1 \\ \bar{\alpha} & \text{otherwise.} \end{cases} \quad (20)$$

A closed-form representation of K_j that depends only on $\bar{\alpha}, \bar{\beta}, N$ and j was derived in Schöllig et al. [2010]. The values are explicitly stated in (A.3) and (A.4) (in terms of γ and $\bar{\epsilon}$) and are used to derive the results in Sec. 4.

Step 2 The disturbance estimate is updated in each iteration j based on the measurement Y_j ,

$$\hat{D}_j = \hat{D}_{j-1} + K_j (Y_j - H \hat{D}_{j-1}), \quad (21)$$

where \hat{D}_0 is given by (16).

Important to note is that the independent estimation problem (I) is simply a special case of the cooperative framework (II) with $N = 1$. We compare the performance of both estimation schemes via the variance of the disturbance estimate.

3.1 Variance of Disturbance Estimate

We determine the variance of the disturbance estimate when applying the above Kalman filter equations to the real system. The covariance matrix of the real system is denoted by $P_j = [p_j^{(k,l)}]$, $k, l \in \mathcal{K}$. A recursive equation for calculating P_j is derived from (21),

$$P_j = (I - K_j H) P_{j-1} (I - K_j H)^T + K_j K_j^T \quad (22)$$

where K_j represent the Kalman gains calculated from the nominal covariance matrices S_j , based on the assumed variance values $\bar{\alpha}, \bar{\beta}$. The initial covariance matrix P_0 is defined analogously to S_0 , see (20), but is based on the unknown system variances α, β ,

$$p_0^{(k,l)} = \begin{cases} \alpha + \beta & \text{for } k = l \geq 1 \\ \alpha & \text{otherwise.} \end{cases} \quad (23)$$

In brief, the estimation algorithm is run based on our assumed variance values $\bar{\alpha}, \bar{\beta}$ and yields the Kalman gain K_j used in (21); to determine the variance of the estimate when running the estimation on the real system, we have to take the real variance parameters α, β into account, via (23) and (21). This step is artificial, since the real values are not known; however, it allows us to study the effects of incorrect variance assumptions on the estimation result, see Sec. 4. In the ideal case when $\bar{\alpha} = \alpha$ and $\bar{\beta} = \beta$, $S_j = P_j$ for all $j \in \{1, 2, \dots\}$. This scenario was studied in Schöllig et al. [2010].

When we compare the performance of the independent (I) and the joint (II) estimation, we use the variance of an individual's disturbance estimate, which in both cases is given by

$$E \left[(d^i - \hat{d}_j^i)^2 \right] = p_j^{(i,i)} = p_j^{(1,1)}, \quad \forall i \in \mathcal{I}, \quad (24)$$

where $\hat{D}_j = [\hat{d}_j^i]$, $i \in \mathcal{I}$, and $P_j = [p_j^{(k,l)}]$, $k, l \in \mathcal{K}$. The variance is identical for all agents, since the same assumptions on the dynamics (8) and the initial noise characteristics (9) hold for every agent. The variance of an individual's disturbance (24) is a measure for the effectiveness of the disturbance compensation, since in the general ILC framework, cf. (2)-(3), the input update rule is based on the current estimate \hat{d}_j^i . See for example Schöllig and D'Andrea [2009].

Below, we distinguish between the individual disturbance variance (24) in the cases of joint and independent estimation, where the latter is given when evaluating the disturbance variance for $N = 1$, i.e.

$$p_j^{(1,1)} \Big|_{N=1}. \quad (25)$$

Thus, the initial question can be reformulated: To what degree does joint estimation benefit the individual learning of an agent? How does an incorrect assumption on the initial variances affect the learning performance?

3.2 Performance Index

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 , cf. (5). This is best achieved if the variance

in the estimate of x_j^i is small; in other words, the variance of the state estimate can be used as a measure of learning performance, see Schöllig et al. [2010].

Given (5) and (9), the best estimate of the state \hat{x}_j^i at iteration j is equal to the current disturbance estimate,

$$\hat{x}_j^i = \hat{d}_j^i, \quad (26)$$

since the noise ξ_j^i has zero mean. Recalling the noise characteristics (9) and the previous assumption of mutual independence between d^i and ξ_j^i , we obtain the variance of state estimate from the sum of the variance of the estimate \hat{d}_j^i and the variance of ξ_j^i . That is, with (24) and (10),

$$\begin{aligned} E \left[(x_j^i - \hat{x}_j^i)^2 \right] &= E \left[(d^i + \xi_j^i - \hat{d}_j^i)^2 \right] \\ &= p_j^{(1,1)} + \lambda. \end{aligned} \quad (27)$$

We introduce the performance index as the ratio of the state variance in the independent case vs. the joint case,

$$R = \frac{p_j^{(1,1)} \Big|_{N=1} + \lambda}{p_j^{(1,1)} + \lambda}, \quad (28)$$

using the notation of (25). Given this definition, a value $R > 1$ indicates that the joint estimation scheme is more beneficial than an independent estimation, while a value $R < 1$ means that the independent estimation yields a better performance. The larger the value R , the more beneficial the joint estimation algorithm.

In Sec. 4, we analyze the independent and joint estimation schemes for their sensitivity to inaccurate disturbance assumptions, cf. (9). In this context, the performance index (28) allows us to compare the performance of the two estimation schemes and to determine in which cases a joint estimation is more beneficial.

4. SENSITIVITY ANALYSIS

In our previous work Schöllig et al. [2010], we studied independent and joint estimation under the assumption that the variances of the individual and common disturbance, α and β , are known. These values were used when solving for the Kalman gains K_j . They also served as initial conditions when calculating the variance of the disturbance estimate (23) in each iteration j .

Below, the effects of incorrect variance assumptions on the performance of the estimation algorithm are studied for both independent and joint estimation, and compared with the result derived in Schöllig et al. [2010]. This analysis allows us to deduce rules on how to choose $\bar{\alpha}$ and $\bar{\beta}$ in order to achieve robustness to assumption errors.

4.1 Variance of Disturbance Estimate

We derive the individual's disturbance variance $p_j^{(1,1)}$ given the assumptions in Sec. 2.3. In this context, we introduce the notation $*(\cdot)$ to represent a respective quantity assuming perfect knowledge, i.e.

$$\bar{\epsilon} := \epsilon. \quad (29)$$

The ideal value of the individual's disturbance variance was derived in Schöllig et al. [2010],

$$*p_j^{(1,1)} = \frac{\alpha + \beta + j\beta^2 + jN\alpha\beta}{(1 + j\beta)(1 + j\beta + jN\alpha)}, \quad (30)$$

under the assumption that the disturbance characteristics of $d^{i,ind}$ and d^0 are known, i.e. $\delta = 0$. For the following sensitivity analysis, we express (30) in terms of the total disturbance variance γ , see (11), and the disturbance ratio ϵ ,

$$*p_j^{(1,1)} = \frac{\gamma + j\gamma^2(1 - \epsilon)(1 + \epsilon(N - 1))}{m_j(\gamma, \epsilon)n_j(\gamma, \epsilon, N)}, \quad (31)$$

$$:= f_j(\gamma, \epsilon, N) \quad (32)$$

where

$$\begin{aligned} m_j(\gamma, \epsilon) &= 1 + j\gamma(1 - \epsilon) \\ n_j(\gamma, \epsilon, N) &= 1 + j\gamma(1 + \epsilon(N - 1)). \end{aligned} \quad (33)$$

The variance $*p_j^{(1,1)}$ serves as reference value for determining the robustness of the joint estimation scheme to an incorrect similarity assumption, $\bar{\epsilon} \neq \epsilon$.

We derive an analytical expression for the variance of an agent's disturbance estimate for the general case, where ϵ is not assumed to be known.

Proposition 1. The variance of an agent's disturbance estimate can be expressed in terms of the combined variance γ , the nominal disturbance ratio $\bar{\epsilon}$, the assumption error δ , the number of agents N , and the iteration j ,

$$p_j^{(1,1)} = f_j(\gamma, \bar{\epsilon}, N) - \delta g_j(\gamma, \bar{\epsilon}, N), \quad (34)$$

where

$$g_j(\gamma, \bar{\epsilon}, N) = \frac{j\gamma^2\bar{\epsilon}(N - 1)(2 + j\gamma(2 + \bar{\epsilon}(N - 2)))}{m_j(\gamma, \bar{\epsilon})^2 n_j(\gamma, \bar{\epsilon}, N)^2} \quad (35)$$

and $f_j(\gamma, \bar{\epsilon}, N)$, $m_j(\gamma, \bar{\epsilon})$, $n_j(\gamma, \bar{\epsilon}, N)$ are defined by (32) and (33). Recalling (13), the assumption error is bounded by

$$-\bar{\epsilon} \leq \delta \leq 1 - \bar{\epsilon}. \quad (36)$$

The result is obtained by first solving the Kalman filter equations (19) for the Kalman gain K_j , given the initial conditions (20). Finally, the recursive equation (22) yields the above result, given the starting values (23). A more detailed proof is found in the Appendix A.

The goal of the following analysis is to compare a real scenario (where $\bar{\epsilon} \neq \epsilon$) to an ideal case, where we have perfect system knowledge ($\bar{\epsilon} = \epsilon$).

Independent Estimation If every agent estimates its disturbance d^i independently, a partitioning of d^i into a common part and an individual part, cf. (4), is not necessary. Consequently, an incorrect assumption $\bar{\epsilon} \neq \epsilon$ has no effect on the disturbance estimate; that is,

$$p_j^{(1,1)}|_{N=1} = *p_j^{(1,1)}|_{N=1}. \quad (37)$$

This is also reflected by the equations (34)-(35) and (31)-(32), where

$$g_j(\gamma, \bar{\epsilon}, 1) = 0 \quad (38)$$

and

$$f_j(\gamma, \bar{\epsilon}, 1) = f_j(\gamma, \epsilon, 1) = \frac{\gamma}{1 + j\gamma} \quad (39)$$

depends only on the sum γ , which is assumed to be known precisely, $\gamma = \alpha + \beta = \bar{\alpha} + \bar{\beta}$, see (11).

In the limit case when $j \rightarrow \infty$, the variance of the disturbance estimate approaches zero monotonically,

$$\lim_{j \rightarrow \infty} p_j^{(1,1)}|_{N=1} = 0 \quad \text{with} \quad \frac{\partial p_j^{(1,1)}|_{N=1}}{\partial j} \leq 0. \quad (40)$$

Fig. 2 illustrates the evolution of the variance $p_j^{(1,1)}|_{N=1}$ for the example introduced in Sec. 5.

Joint Estimation If N agents jointly estimate their disturbance d^i , the assumption on the disturbance partitioning is crucial to the estimation performance. We compare the variance $p_j^{(1,1)}$ with the ideal value by using the relation $\epsilon = \bar{\epsilon} + \delta$ in (31) and subtracting (31) from (34). This yields

$$*p_j^{(1,1)} < p_j^{(1,1)}, \quad (41)$$

for $N > 1$ and $\bar{\epsilon} \neq \epsilon$. The performance of the joint estimation algorithm when assuming perfect knowledge is better (i.e. results in a smaller variance) than in the realistic scenario, where the disturbance partitioning is not accurately known.

When analyzing (34), (35) with respect to the assumption error δ , we obtain

$$\frac{\partial p_j^{(1,1)}}{\partial \delta} \leq 0, \quad \frac{\partial p_j^{(1,1)}}{\partial \epsilon} \leq 0. \quad (42)$$

For a given ratio $\bar{\epsilon}$, the performance of the joint estimation scheme improves (i.e. smaller variance) if the real ratio ϵ increases. In other words, no matter how wrong our assumption on the disturbance partitioning $\bar{\epsilon}$ is, the joint estimation scheme becomes more effective if the agents show an increasing similarity in reality (and as long as $\bar{\epsilon} \neq 0$). However, if the nominal ratio is chosen to be zero ($\bar{\epsilon} = 0$, assuming the agents are completely different), the joint variance (34) corresponds to the individual variance (39) and is not improved by the agent's actual similarity, see Fig. 2.

An interesting next step is to study the limit behavior of the disturbance variance for $j \rightarrow \infty$. We first consider $f_j(\gamma, \bar{\epsilon}, N)$ in (34) and state

$$\lim_{j \rightarrow \infty} f_j(\gamma, \bar{\epsilon}, N) = 0 \quad \text{with} \quad \frac{\partial f_j(\gamma, \bar{\epsilon}, N)}{\partial j} \leq 0. \quad (43)$$

The function $f_j(\gamma, \epsilon, N)$ also defines the perfect variance and its limit, cf. (30), and hence,

$$\lim_{j \rightarrow \infty} *p_j^{(1,1)} = 0 \quad \text{with} \quad \frac{\partial *p_j^{(1,1)}}{\partial j} \leq 0. \quad (44)$$

This means that, in the perfect knowledge case, the disturbance is accurately estimated after a large number of iterations. We keep this in mind when analyzing the limit behavior of the variance $p_j^{(1,1)}$. Two cases are distinguished:

- (1) If we assume that the agents are not perfectly identical, that is $\bar{\epsilon} \neq 1$, the variance converges to zero,

$$\lim_{j \rightarrow \infty} p_j^{(1,1)} = 0. \quad (45)$$

The joint estimation algorithm provides an increasingly accurate estimate of the disturbance with each

additional iteration. Even if our assumption is wrong, and $\delta \neq 0$, the estimation algorithm provides us with the correct disturbance estimate in the limit case when the number of iterations approaches infinity.²

- (2) If we assume that the agents are identical, $\bar{\epsilon} = 1$, the limit behavior for $j \rightarrow \infty$ is

$$\lim_{j \rightarrow \infty} p_j^{(1,1)} = -\delta \gamma \frac{N-1}{N} := \ell(\gamma, \delta, N), \quad (46)$$

where $-1 \leq \delta \leq 0$, cf. (36). The variance of the disturbance estimate has a finite, non-zero limit value (if $\delta \neq 0$, which is equivalent to saying $\epsilon \neq 1$). In other words, when we assume the agents are identical and they are (in fact) not, the joint estimation scheme does not provide an accurate estimate, even after a large number of iterations. The variance in the limit $j \rightarrow \infty$ depends on the total disturbance level γ , the number of agents N , and the assumption error δ , where

$$\frac{\partial \ell(\gamma, \delta, N)}{\partial N} \geq 0, \quad \frac{\partial \ell(\gamma, \delta, N)}{\partial \delta} \leq 0, \quad \frac{\partial \ell(\gamma, \delta, N)}{\partial \gamma} \geq 0. \quad (47)$$

The limiting variance grows with an increasing number of agents, an increasing difference $|\delta|$ between the assumed and real ratio (note that $\delta \leq 0$), and an increasing overall disturbance level. When $\epsilon = 0$ (meaning that the agents, in reality, have no common disturbance component), the difference between the assumed and real ratio is largest, $\delta = -1$, and

$$\lim_{N \rightarrow \infty} \ell(\gamma, -1, N) = \gamma. \quad (48)$$

Based on the above results, first conclusions can be drawn on how our assumption on the agents' similarity affects the joint estimate, and on how to choose the nominal ratio $\bar{\epsilon}$:

- Generally, an incorrect assumption about the agents' similarity ($\bar{\epsilon} \neq \epsilon$) decreases the accuracy of the disturbance estimate and increases the variance of the disturbance estimate, cf. (41).
- In the worst case, the variance does not approach zero in the limit when the number of iterations goes to infinity. This happens if we assume that the agents are perfectly identical, $\bar{\epsilon} = 1$, but they are, in fact, not. In this case, the limit value of the variance for $j \rightarrow \infty$ increases with the number of agents and with an increasing total disturbance. The variance is worst if the agents share no common disturbance component in the real scenario, $\epsilon = 0$ or $\delta = -1$. *To assure that the variance converges to zero for any ratio ϵ , we must choose $\bar{\epsilon} \neq 1$.*
- As shown in (42), the joint estimation scheme is more beneficial if the agents are more similar in reality. This is independent of the assumed value and holds for all $\bar{\epsilon} \neq 0$. Hence, if we want to benefit from this characteristic of the joint estimation scheme, and if we expect a certain degree of similarity between the agents, the ratio $\bar{\epsilon}$ should not be set to zero.

Fig. 2 summarizes the characteristics derived above: the perfect knowledge case results in the smallest variance values and, in the limit case, variances approach zero

² Mathematica files including the presented results are available at www.idsc.ethz.ch/Downloads/multiagentILC.

except for the case $\bar{\epsilon} = 1$, where the limit value is obtained from (46), cf. Sec. 5.

Thus far we have compared the disturbance estimate $p_j^{(1,1)}$ with the perfect value $*p_j^{(1,1)}$. What remains is to compare the independent and joint estimation schemes given an insufficient knowledge $\bar{\epsilon} \neq \epsilon$. In the following section, we attempt to determine whether, in light of our new findings, the joint estimation continues to perform better than the independent estimation.

4.2 Performance Index

We introduced the performance index R as the ratio of the state variance in the independent case vs. the joint case, see Sec. 3.2 and Schöllig et al. [2010]. As shown in Schöllig et al. [2010], in the nominal case, assuming the real values α, β are known, joint estimation is always beneficial and yields a performance index

$$1 \leq *R = \frac{*p_j^{(1,1)}|_{N=1} + \lambda}{*p_j^{(1,1)} + \lambda}. \quad (49)$$

When the number of iterations increases, the performance index shows the following limit behavior,

$$\lim_{j \rightarrow \infty} *R = \begin{cases} N & \text{for } \epsilon = 1 \text{ and } \lambda = 0, \\ 1 & \text{otherwise.} \end{cases} \quad (50)$$

If the agents are not identical ($\epsilon \neq 1$), the performance index $*R$ converges to one. The same limit behavior is observed if process noise acts on the system, $\lambda \neq 0$. Only if (i) the agents are identical, and (ii) noise is due to measurement only, is the limit value of the performance index equal to N , see Schöllig et al. [2010] for a detailed analysis.

From Sec. 4.1 we know that a mismatch between the real disturbance characteristics and the nominal values does not affect the independent estimate, see (37), but it does corrupt the disturbance estimate when jointly estimating, see (41). In this context, it is interesting to ask: Is it possible that the performance index becomes smaller than one, meaning that the individual estimation performs better than the joint estimation? Taking into account the disturbance variances derived in Sec. 4.1, we answer this question below.

First, we compare the performance index in the ideal case, where $\bar{\epsilon} = \epsilon$, with the performance index R of the realistic scenario, $\bar{\epsilon} \neq \epsilon$. Given the definition of the performance index in (28) and the derived characteristics of the disturbance variance, cf. (37) and (41), we conclude that

$$R < *R. \quad (51)$$

The ideal performance index represents an upper bound to R and, recalling (49), is larger or equal to one.

Second, before studying the evolution of R with j , we focus on the limiting behavior of performance index R as $j \rightarrow \infty$. Taking into account the limit values of the disturbance variance derived in Sec. 4.1 and explicitly stated in (40), (45) and (46), the following statement is derived:

Lemma 2. As $j \rightarrow \infty$, the performance index R , defined by (28) and (34), shows two distinct limiting behaviors:

(1) If $\bar{\epsilon} \neq 1$, the performance index converges to one,

$$\lim_{j \rightarrow \infty} R = 1 \quad (52)$$

for all possible values $N, j > 1, \gamma \geq 0$, and $\epsilon, \lambda \in [0, 1]$, where $0 \leq \bar{\epsilon} < 1$.

(2) If $\bar{\epsilon} = 1$, the limit behavior for $j \rightarrow \infty$ is

$$\lim_{j \rightarrow \infty} R = \frac{\lambda N}{\lambda N - \delta \gamma (N - 1)} \quad (53)$$

$$:= \ell_R(\gamma, N, \delta, \lambda) \leq 1,$$

for all possible values $N, j > 1, \gamma > 0$, and $\lambda \in [0, 1]$, where $-1 \leq \delta < 0$.

Since

$$\frac{\partial \ell_R}{\partial N} \leq 0 \quad \text{and} \quad \frac{\partial \ell_R}{\partial \delta} \geq 0, \quad (54)$$

the limit value ℓ_R reaches its minimum for $\delta = -1 \Leftrightarrow \epsilon = 0$, and $N \rightarrow \infty$. In this case,

$$\lim_{N \rightarrow \infty} \ell_R(\gamma, N, -1, \lambda) = \frac{\lambda}{\lambda + \gamma}. \quad (55)$$

Interpretation of the result. If we assume the agents are not identical, $\bar{\epsilon} \neq 1$, the joint estimation algorithm converges and provides us with an accurate disturbance estimate for $j \rightarrow \infty$, cf. (45) and (52). In other words, with respect to the convergence of the disturbance variance, the joint estimation scheme is robust to assumption errors as long as $\bar{\epsilon} \neq 1$. If we choose $\bar{\epsilon} = 1$ and $\delta \neq 0$, however, we lose this property, cf. (46) and (55). That is, under the assumption that all agents are identical, the joint estimation algorithm is highly sensitive to assumption errors. In the case of pure measurement noise $\lambda = 0$, (53) is zero. Note that the case, $\lambda = 0$ and $\bar{\epsilon} = 1$, yields the best performance improvement in the ideal case, see (50), but is the most sensitive to assumption errors. *In brief, when building upon a joint estimation scheme, the disturbance ratio $\bar{\epsilon}$ should not be set to one.* Moreover, we have shown that the performance index can be less than one (see statement (2) in Lemma 2). In these cases, an independent estimation is preferable.

The above result guarantees that we eventually obtain a precise estimate of the disturbance, but it does not provide insight into the transient performance of the joint estimation scheme as compared to an independent algorithm.

As a last step, we perform a more detailed analysis and identify parameter combinations that support an application of the joint estimation scheme (where $R > 1$). At an iteration j , the joint estimation is beneficial, if the inequality $R \geq 1$ is satisfied, which is equivalent to saying

$$f_j(\gamma, \bar{\epsilon}, 1) \geq f_j(\gamma, \bar{\epsilon}, N) - \delta g_j(\gamma, \bar{\epsilon}, N) \quad (56)$$

for a given number of agents N , a known nominal disturbance ratio $\bar{\epsilon}$, and an overall disturbance γ , see (28) and (34). The functions f_j and g_j are non-negative for all possible values $N, j \geq 1, \gamma \geq 0$, and $\bar{\epsilon} \in [0, 1]$, cf. (32) and (35),

$$f_j(\gamma, \bar{\epsilon}, N) \geq 0 \quad \text{and} \quad g_j(\gamma, \bar{\epsilon}, N) \geq 0, \quad (57)$$

and

$$\frac{\partial f_j(\gamma, \bar{\epsilon}, N)}{\partial N} = \frac{-\epsilon^2 \gamma^2 j}{m_j(\gamma, \epsilon) n_j(\gamma, \epsilon, N)^2} \leq 0. \quad (58)$$

Combining (56) with (57),(58) yields the following lemma:

Lemma 3. A sufficient condition for the joint estimation to yield a better (or equal) learning performance than the independent estimation is a disturbance ratio ϵ that is larger than (or equal to) the assumed one. The following implication holds,

$$\epsilon \geq \bar{\epsilon} \Rightarrow R \geq 1, \quad (59)$$

for all possible values $N, j \geq 1, \gamma \geq 0$, and $\bar{\epsilon} \in [0, 1]$.

Interpretation of the result. If the agents are more similar in reality than assumed, it is beneficial to jointly estimate the repetitive disturbances. We should avoid an overestimation of the similarity between the agents, since in this case an independent scheme would actually be more beneficial. However, if we underestimate the similarity for a given situation defined by ϵ , we increase the variance of the disturbance estimate, cf. (42). Moreover, from (28) and (34) with $\delta = \epsilon - \bar{\epsilon}$, we obtain

$$\frac{\partial R}{\partial \bar{\epsilon}} \geq 0 \quad \text{if} \quad \epsilon \geq \bar{\epsilon}, \quad (60)$$

which means that, if we underestimate the similarity of the agents (reducing $\bar{\epsilon}$), we reduce the benefit of the joint estimation vs. the independent estimation.² *A design rule for $\bar{\epsilon}$ is consequently: Given a priori knowledge about the multi-agent system, make $\bar{\epsilon}$ as large as possible while, at the same time, ensuring that it is less than the real value ϵ .* In this case, applying the joint estimation guarantees a better learning performance than the independent estimation scheme; however, the benefits of joint estimation remain marginal, as shown in Schöllig et al. [2010].

Note that if $\delta \geq 0 \Leftrightarrow \epsilon \leq \bar{\epsilon}$, there are also cases for which $R \geq 1$, see Fig. 3. From (56) with (13), we derive the necessary and sufficient condition:²

Lemma 4. In the proposed multi-agent learning framework, joint estimation yields a better (or equal) learning performance as the independent estimation *if and only if* the following inequality is satisfied for given values $N, j > 1, \gamma, \bar{\epsilon} > 0$ and ϵ :

$$\epsilon \geq \bar{\epsilon} \frac{1 + \gamma j (2 + \gamma j (1 + \bar{\epsilon}^2 (N - 1)))}{(1 + \gamma j) (2 + \gamma j (2 + \bar{\epsilon} (N - 2)))} \quad (61)$$

$$:= \bar{\epsilon} h_j(\gamma, \bar{\epsilon}, N)$$

with $h_j(\gamma, \bar{\epsilon}, N) < 1$. The inequality becomes less restrictive for an increasing number of agents, $\partial h_j / \partial N < 0$, and reduces to:

$$\epsilon \geq \bar{\epsilon} \frac{\bar{\epsilon} \gamma j}{1 + \gamma j} \quad \text{for} \quad N \rightarrow \infty, \quad (62)$$

and, in the limit case for $j \rightarrow \infty$, to:

$$\epsilon \geq \bar{\epsilon}^2 \quad \text{for} \quad j, N \rightarrow \infty. \quad (63)$$

Fig. 3 illustrates the previously derived characteristics for $\epsilon = 0.5$ and the nominal values $\bar{\epsilon}$ ranging from 0 to 1. Note that for the case $\bar{\epsilon} = 0.75$, the performance index crosses the $R = 1$ line and finally approaches one.

The previous analysis focused on identifying the cases for which the joint estimation is beneficial despite an incorrect assumption $\bar{\epsilon} \neq \epsilon$ and, in turn, defined the cases where an incorrect assumption $\bar{\epsilon} \neq \epsilon$ corrupts the performance of the joint estimation to such an extent that




Fig. 2. The variance of an individual's disturbance estimate for $N = 10$ agents, $\epsilon = 0.5$, $\gamma = 0.1$ and $\lambda = 0$ (pure measurement noise).

an independent estimation is more effective. The numerical examples below summarize the results of this section and highlight the sensitivity of the joint estimation scheme to incorrect noise assumptions.

5. NUMERICAL EXAMPLES

We consider a group of $N = 10$ agents with a similarity of $\epsilon = 0.5$. The noise is due to measurement only ($\lambda = 0$) and the overall disturbance γ is 0.1. Fig. 2 and Fig. 3 show the evolution of the variance and the performance index for various nominal values $\bar{\epsilon}$ ranging from 0 to 1. Note that different intervals of j are chosen in Fig. 2 and Fig. 3 to emphasize the main characteristics. Both figures show that an incorrect similarity assumption results in a worse performance, i.e. in a higher variance in Fig. 2 and a lower performance index in Fig. 3. The limiting behavior in the case of $\bar{\epsilon} = 1$, see (46) and (53), is for the given scenario

$$\lim_{j \rightarrow \infty} p_j^{(1,1)} = 0.5 \cdot 0.1 \cdot \frac{9}{10} = 0.045 \quad \text{and} \quad \lim_{j \rightarrow \infty} R = 0.$$

In all other cases, the variance approaches 0 (see Fig. 2) and the performance index 1 (see Fig. 3) as $j \rightarrow \infty$. For the two cases where $\bar{\epsilon} > \epsilon$, the performance index is (partly) smaller than one.

6. CONCLUSION

We analyzed the sensitivity of joint estimation to the underlying assumption of similarity between agents. The analysis was driven by our previous results, which showed that the learning performance of an agent is improved by exchanging information with other agents that are learning the same task.

While previous results assumed perfect knowledge about the degree of similarity between the agents, this paper studied the effects of an incorrect similarity assumption. We found that an incorrect assumption not only degrades the performance of the joint estimation scheme (when compared to the case of perfect knowledge), but that, for some problems, the joint estimation performs worse than an independent estimation scheme. This is particularly




Fig. 3. The performance index for $N = 10$ agents, $\epsilon = 0.5$, $\gamma = 0.1$ and $\lambda = 0$ (pure measurement noise).

true if we overestimate the similarity between the agents. As a consequence, it is not advisable to assume that agents are identical because, if agents are not identical in reality, the joint estimation does not even converge to the correct disturbance value. However, note that from our previous analysis, we know that the case of identical agents (and no process noise) provided the best performance improvement of joint vs. independent estimation – an improvement of a factor equal to the number of agents. In other words, the case with the highest performance improvement (due to joint estimation) shows the highest sensitivity to assumption errors.

To conclude: in order to guarantee improved performance over an individual learning scheme, a joint estimation scheme must not overestimate the similarity between the agents.

ACKNOWLEDGEMENTS

The authors would like to thank Javier Alonso-Mora for many fruitful discussions.

REFERENCES

- B. Bamieh, J.B. Pearson, B.A. Francis, and A. Tannenbaum. A lifting technique for linear periodic systems with applications to sampled-data control. *Systems & Control Letters*, 17(2):79–88, 1991.
- K.L. Barton and A.G. Alleyne. A cross-coupled iterative learning control design for precision motion control. *IEEE Transactions on Control Systems Technology*, 16(6):1218–1231, 2008.
- Z. Bien and J.X. Xu. *Iterative learning control: analysis, design, integration and applications*. Kluwer Academic Publishers Norwell, MA, USA, 1998.
- D.A. Bristow, M. Tharayil, and A.G. Alleyne. A survey of iterative learning control. *IEEE Control Systems Magazine*, 26(3):96–114, 2006.
- M. Butcher, A. Karimi, and R. Longchamp. Iterative learning control based on stochastic approximation. In *Proceedings of the 17th IFAC World Congress*, pages 1478–1483, 2008.

- C.K. Chui and G. Chen. *Kalman Filtering: with Real-Time Applications (Springer Series in Information Sciences)*. Springer, 1998.
- J. Hättönen, D.H. Owens, and K. Feng. Basis functions and parameter optimisation in high-order iterative learning control. *Automatica*, 42(2):287–294, 2006.
- S. Mannor and J.S. Shamma. Multi-agent learning for engineers. *Artificial Intelligence*, 171(7):417–422, 2007.
- M.J. Matarić. Reinforcement learning in the multi-robot domain. *Autonomous Robots*, 4:73–83, 1997.
- J. Morimoto and K. Doya. Robust reinforcement learning. *Neural computation*, 17(2):335–359, 2005.
- M. Norrlöf. Disturbance rejection using an ILC algorithm with iteration varying filters. *Asian Journal of Control*, 6:432–438, 2004.
- L. Panait and S. Luke. Cooperative multi-agent learning: The state of the art. *Autonomous Agents and Multi-Agent Systems*, 11(3):387–434, 2005.
- M.Q. Phan and R.W. Longman. A mathematical theory of learning control for linear discrete multivariable systems. In *Proceedings of the AIAA/AAS Astrodynamics Conference*, pages 740–746, 1988.
- M.Q. Phan and R.W. Longman. Higher-order iterative learning control by pole placement and noise filtering. In *Proceedings of the IFAC World Congress*, pages 1899–1904, 2002.
- A. Schaerf, Y. Shoham, and M. Tennenholtz. Adaptive load balancing: a study in multi-agent learning. *Journal of Artificial Intelligence Research*, 2(1):475–500, 1994.
- A. Schöllig and R. D’Andrea. Optimization-based iterative learning control for trajectory tracking. In *Proceedings of the European Control Conference*, pages 1505–1510, 2009.
- A. Schöllig, Javier Alonso-Mora, and R. D’Andrea. Independent vs. joint estimation in multi-agent iterative learning control. In *Proceedings of the 49th IEEE Conference on Decision and Control*, pages 6949–6954, 2010.
- R. Tousain, E. van der Meche, and O. Bosgra. Design strategy for iterative learning control based on optimal control. In *Proceedings of the 40th IEEE Conference on Decision and Control*, volume 5, pages 4463–4468, 2001.
- M. Verhaegen and V. Verdult. *Filtering and System Identification: A Least Squares Approach*. Cambridge University Press, 2007.

Appendix A. PROOF OF PROPOSITION 1

We derive an explicit representation of the variance $p_j^{(1,1)}$, as presented in Proposition 1, that depends on the total variance γ , the nominal similarity factor $\bar{\epsilon}$, the assumption error δ , the iteration j , and the number of agents N .

The proof proceeds similarly to the derivation of the ideal variance (31) in Schöllig et al. [2010]. Note that in Schöllig et al. [2010], we assumed that the real noise characteristics (9) were known precisely and

$$p_j^{(1,1)} \text{ simply denoted the ideal variance } *p_j^{(1,1)}. \quad (\text{A.1})$$

Matlab and Mathematica files for reproducing the results below are available on the project webpage².

Proof. For the general case, we derive a closed form of the covariance matrix P_j , with the assumptions in Sec. 2.3. Since the disturbance and noise characteristics are the same for each agent, the covariance matrix is of the following symmetric structure,

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

We derive a recursive relationship for the values in (A.2) based on (21) and the closed-form representation of K_j (derived in Schöllig et al. [2010]). The Kalman gains K_j are calculated based on the nominal disturbance variances $\bar{\alpha} = \bar{\epsilon}\gamma$ and $\bar{\beta} = (1 - \bar{\epsilon})\gamma$ and are given by

$$k_j^{(k,l)} = \begin{cases} k_j^{(0,1)} & \text{if } k = 0 \\ k_j^{(1,1)} & \text{if } k = l \\ k_j^{(1,2)} & \text{otherwise,} \end{cases} \quad (\text{A.3})$$

where

$$\begin{aligned} k_j^{(0,1)} &= \frac{\bar{\epsilon}\gamma}{n_j(\gamma, \bar{\epsilon}, N)}, & k_j^{(1,2)} &= \frac{\bar{\epsilon}\gamma}{m_j(\gamma, \bar{\epsilon})n_j(\gamma, \bar{\epsilon}, N)}, \\ k_j^{(1,1)} &= f_j(\gamma, \bar{\epsilon}, N), \end{aligned} \quad (\text{A.4})$$

From (21) with (A.2) and (A.4), we obtain recursive equations for

$$p_j^{(0,0)}, p_j^{(0,1)}, p_j^{(1,1)}, p_j^{(1,1)} \quad (\text{A.5})$$

that depend only on the $(j - 1)$ th values of (A.5), on the Kalman gains (A.4) and the number of agents N . A proof by induction using the recursive equations for (A.5) with initial condition (23) verifies the closed-form expression in (34). For the proof, the closed-form representations of all values (A.5) were needed and derived. However, the only value of interest is $p_j^{(1,1)}$.

Mathematica files with the recursive equations for (A.5) and the closed-form representations of all quantities (A.5) are available on the project webpage².