

LQG Online Learning

Giorgio Gnecco

giorgio.gnecco@imtlucca.it

Alberto Bemporad

alberto.bemporad@imtlucca.it

DYSCO Research Unit, IMT School for Advanced Studies,

Piazza S. Francesco, 19-55110 Lucca, Italy

Marco Gori

marco@diism.unisi.it

DIISM Department, University of Siena, Via Roma, 56-53100 Siena, Italy

Marcello Sanguineti

marcello.sanguineti@unige.it

DIBRIS Department, University of Genoa, Via Opera Pia, 13-16145 Genova, Italy

Optimal control theory and machine learning techniques are combined to formulate and solve in closed form an optimal control formulation of online learning from supervised examples with regularization of the updates. The connections with the classical linear quadratic gaussian (LQG) optimal control problem, of which the proposed learning paradigm is a nontrivial variation as it involves random matrices, are investigated. The obtained optimal solutions are compared with the Kalman filter estimate of the parameter vector to be learned. It is shown that the proposed algorithm is less sensitive to outliers with respect to the Kalman estimate (thanks to the presence of the regularization term), thus providing smoother estimates with respect to time. The basic formulation of the proposed online learning framework refers to a discrete-time setting with a finite learning horizon and a linear model. Various extensions are investigated, including the infinite learning horizon and, via the so-called kernel trick, the case of nonlinear models.

1 Introduction

In recent years, the combination of techniques from the fields of optimization/control and machine learning has led to a successful interaction between the two disciplines. As an example, we mention reinforcement learning (Bertsekas, 1996b; Sutton & Barto, 1998). The cross-fertilization between these two fields shows itself in both directions.

1.1 Application of Machine Learning Techniques to Optimization/Control. Sparsity-inducing regularization techniques from machine learning have often been exploited to find suboptimal solutions to an initially unregularized optimization problem, having at the same time a sufficiently small number of nonzero arguments. For instance, the least absolute shrinkage and selection operator (LASSO) (Tibshirani, 1996) was applied in Gnecco, Morisi, and Bemporad (2015) to consensus problems and in Gallieri and Maciejowski (2012) to model predictive control (MPC).

Applications of machine learning techniques to control can be found, in Suykens, Vandewalle, & De Moor (2001) and in the series of papers (Gnecco & Sanguineti, 2010; Gaggero, Gnecco, & Sanguineti, 2013, 2014), where least squares support vector machines (LS-SVMs) and one-hidden-layer perceptron neural networks, respectively, were applied to find suboptimal solutions to optimal control problems. Mesbahi and Egerstedt (2010) applied spectral graph theory methods, already exploited successfully in machine learning problems (Belkin & Niyogi, 2006), to the control of multiagent dynamical systems.

Least squares support vector machines and spectral graph theory also have been applied, respectively, to system identification (Lataire, Piga, & Tóth, 2014) and control of epidemics (Chakrabarti, Wang, Wang, Leskovec, & Faloutsos, 2008).

1.2 Application of Optimization/Control Techniques to Machine Learning. This is the direction we followed in this letter: we develop an approach that exploits for machine learning techniques from optimization and control (particularly, optimal control). We focus on the problem of estimating online an unknown parameter vector modeling a noisy linear relationship between input examples and their output labels. The emphasis is on linear regression, but an extension to nonlinear regression is also detailed. Moreover, the model can be applied to linear (respectively, nonlinear) binary classification by interpreting a suitable nonlinear transformation (e.g., a sigmoidal mapping) of the output as the a posteriori probability (given the input example) of belonging to one of the two classes.

Specifically, we propose and solve in closed form an optimal control formulation of online learning with supervised examples and a regularization of the updates. In the online framework, the examples become available one by one as time passes and the training of the learning machine is performed continuously. Online learning problems have been investigated, in, for example, Zinkevich (2003), Ralaivola and d'Alché Buc (2005), Smale and Tao (2006), Ying and Pontil (2008), Liu, Park, Wang, & Principe (2009), Shalev-Shwartz (2012), but without using an approach based on optimal control theory. As suggested by the preliminary results that we obtained in Gnecco, Bemporad, Gori, Morisi, and Sanguineti (2015), such an approach can provide a strong theoretical foundation to the choice of a specific online learning algorithm by selecting the updates of the parameter estimate as the

outputs of a sequence of control laws that solve a suitable optimal control problem modeling online learning itself.¹ A distinguishing feature of our study is that we derive online learning algorithms as closed-form optimal solutions to suitable online learning problems. In contrast, typically, work in the literature proposes a certain algorithm and then investigates its properties, but does not analyze the optimality of such an algorithm with respect to a suitable online learning problem. A remarkable exception is Bertsekas (1996a), but it refers to a deterministic optimization problem and, different from our approach, does not contain any regularization of the updates of the parameter estimate.

In a nutshell, our contributions are the following:

- We make the machine learning community aware of a point of view that until now might have been overlooked.
- By exploiting such a viewpoint, we develop a novel machine learning paradigm, for which we provide closed-form solutions.
- We make connections between our results and other machine learning algorithms.

1.3 The Adopted Learning Model. The learning model that we adopt can be considered a nontrivial variation (due to the presence of suitable random matrices) of the linear quadratic (LQ) and linear quadratic gaussian (LQG) optimal control problems, which we briefly summarize. The LQ problem (Bertsekas, 1995) consists of the minimization—with respect to a set of control laws, one for each decision stage—of a convex quadratic cost related to the control of a linear dynamical system, which is decomposed into the summation of several convex quadratic per stage costs, associated with a priori given cost matrices. At each stage, a control law is applied. It is a function of an information vector, which collects all the information available to the controller up to that stage. More precisely, the information vector is formed by the sequence of controls applied to the dynamical system up to the previous stage and by the sequence of measures of the state of the dynamical system itself, acquired up to the current stage. A peculiarity of the LQ problem is that such measures are linearly related to the state, again through suitable a priori given measurement matrices. The measures may be corrupted by additive noise, with given covariance matrices. When all the noise vectors are gaussian, one obtains the LQG problem, for which closed-form optimal control laws in feedback form are known.² They are computed by solving recursively suitable Riccati equations and applying the Kalman filter (Sutton, 1992) to estimate the current state of the dynamical system.

¹The results from Gnecco, Bemporad et al. (2015) correspond, in this letter, to a subset of the results contained in section 3.

²Specifically, as functions of an estimate of the current state of the dynamical system.

The main difference between the LQ/LQG problems and the proposed formulation of online learning with supervised examples and regularization is the following. In our approach, both the cost and measurement matrices are random, being associated with the input examples, which become available as time goes on. It is worth mentioning that the randomness of some matrices in the context of the LQ optimal control problem was considered also in Bertsekas (1995), but in a way not directly applicable to the online learning problem investigated in this letter (see remark 7 in section 3.2 for further details). First, we consider a linear relationship between the input examples and their labels, possibly corrupted by additive noise, and collect into the state vector both the current estimate of the parameter vector modeling the input-output relationship and the parameter vector itself, which is unknown. Then we relax the linearity assumption and address a more general nonlinear context. The goal consists in finding an optimal online estimate of the parameter vector, on the basis of the information associated with the incoming examples, modeled in the simplest case as independent and identically distributed random vectors.³

Each decision stage corresponds to the presentation of one example to the learning machine, whereas one part of the convex per stage cost penalizes quadratically the difference between the observed output and the one predicted by the learning machine, by using the current estimate of the parameter vector. Causality in the updates (i.e., their independence on future examples, which is important for a truly “online” framework) is preserved by constraining the updates to depend only on the history of the observations and updates up to the decision time, likewise in the LQ/LQG problems.

At each stage, the error on future examples is taken into account through the conditional expectation of the summation of the associated per stage costs, conditioned on the current information vector. In the stationary case, the link between the examples used for training and the future examples is only in their common generation model. In order to reduce the influence of outliers on the online estimate of the parameter vector, its smoothness with respect to time is also enforced through the presence of a suitable regularization term in the functional to be optimized, weighted by a positive regularization parameter.⁴ This represents the second part of the convex and quadratic per stage cost. The optimal solution is obtained by applying dynamic programming (DP) and requires the solution of suitable Riccati equations. The difference already noted between the classical LQ/LQG

³This framework is also extended in the letter to other probability models for the generation of the examples (see section 8).

⁴We shall present a comparison with the sequence of Kalman filter estimates of the unknown parameter vector that shows the larger smoothness and less sensitivity to outliers of the sequences of estimates obtained solving the proposed optimal control formulations of online learning (see section 5).

Table 1: Some Correspondences between the Machine Learning Terminology and the Optimal Control One.

Update	Control
Updating function	Control function
Problem OLL (Online Learning)	Optimal control problem
Learning horizon	Optimization horizon
Learning functional	Cost functional
Average learning functional	Average cost functional

problems and the proposed online learning framework (i.e., the random nature of some matrices) determines two different forms for such equations—for the backward and forward phases of DP, respectively. When the optimization horizon is infinite, it is necessary to take into account the random nature of the matrices to perform a convergence analysis of the online estimate of the unknown parameter vector.

Table 1 provides the correspondence between the notations used for optimal control and the ones used for the proposed online learning framework.

1.4 Relationships with Other Machine Learning Techniques. The approaches to online learning most closely related to this work are Kalman filtering (Sutton, 1992; see also Bertsekas, 1995, appendix E) and its kernel version (Ralaivola & d’Alché Buc, 2005; Liu et al., 2009), in which, however, no penalization is made directly on the control (updating) variables. Indeed, one of our contributions consists in developing a theoretical framework in which such a penalization is taken into account and in providing in most cases closed-form results. Interestingly, the obtained solutions can be interpreted as smoothed versions (with respect to time) of the solution obtained applying the Kalman filter only. Most important, we show, both theoretically and numerically, that our solutions are less sensitive to outliers than the Kalman filter estimates. This is very useful (e.g., if one wants to give more importance to a whole set of most recently presented examples than to the current example) to obtain estimates that change more smoothly with respect to time (smoothness of an estimate is a desirable property, e.g., in applications to online system identification and control, in which one has also to control the system just identified). The issue of a smaller sensitivity to outliers is especially important in case of deviations from the gaussianity assumptions (which cannot be a priori excluded in practical applications), or of a mismatch between the true covariance matrices of the noise and the ones used to update the Kalman filter estimates. Other motivations for the presence in our formulation of a penalization of the updating variables are numerical benefits due to the regularization term (a typical way of dealing with ill-conditioning through regularization approaches) and better behavior when considering the extension of the online learning problem to the

case of a time-varying parameter vector. All of these issues are discussed later in the letter. It is also worth mentioning that various forms of regularization of the Kalman filter (and, in an offline context, of the Kalman smoother) were also proposed in the recent literature (e.g., Jay, Duvaut, Darolles, & Gouriéroux, 2011; Wang, Li, & Rizos, 2012; Li, Gui, Gu, Han, & Du, 2014; Liu & Chang, 2016; Khan, Bouaynaya, & Fathallah-Shaykh, 2014), where, however, an optimal control approach was not used.

For each decision stage, the updating formula that provides the solution to the proposed learning paradigm is similar to the ones of other online estimates obtained through various machine learning techniques such as stochastic gradient descent. However, there is a substantial difference: we derive it as the optimal solution of an optimization problem modeling online learning, and this allows us to prove various interesting properties. We believe that this approach could also be fruitfully applied to other machine learning techniques used in online learning.

A number of extensions of the basic problem formulation is also described with some detail at the end of the letter, showing the generality of our theoretical framework, highlighting the large potentialities of the application of optimal control theory to machine learning, and providing hints to stimulate further interdisciplinary research on this topic. A connection with neuroscience is discussed too.

1.5 Organization of the Letter. Section 2 is a nontechnical overview of the main results derived in the letter, written to allow readers who are not familiar with optimal control, but work in the field of machine learning, to appreciate the nature of our approach and its contributions. At the same time, it provides a summary of the main results of the paper. Section 3 introduces and solves the proposed model of online learning as an LQ optimal control problem with random matrices and finite optimization horizon and provides closed-form expressions for the optimal solutions in the LQG case. Section 4 extends the analysis to the infinite-horizon framework. Section 5 investigates convergence properties of the algorithm, whereas section 6 compares the proposed online approach with average regret minimization and the Kalman filter estimates, both theoretically and numerically. Section 7 extends the analysis to nonlinear models (kernel methods). Other extensions are described in section 8. Section 9 is a conclusive discussion. To improve the readability, most technical proofs are contained in the appendix.

2 Overview of the Main Results

In the following, we summarize the main results with links to the parts of the letter in which they are presented:

- *We derive closed-form optimal solutions* for the proposed optimal control formulation of online learning and for some of its extensions. They are expressed in terms of two Riccati equations (see section 3), associated, respectively, with the backward phase of DP (to determine the gain matrix of the optimal controller) and the determination of the gain matrix in the Kalman filter (in the case of gaussian random vectors). Different from the LQG problem, the two Riccati equations have different natures: one involves expectations of random matrices (so it may be called an average Riccati equation), whereas the other involves realizations of random matrices (so it may be called a stochastic Riccati equation). As a consequence, a specific study, detailed in the letter, is needed to study properties of their solutions, which confirms that the proposed problem is not a trivial application of the LQG problem to an online learning framework.

- *We analyze both theoretically and numerically the role of the regularization parameter* (see section 3.3).

- *In the infinite-horizon case, we investigate the existence of a stationary (and linear) optimal updating function* (see section 4), and *stability issues, that is, the convergence to 0 of the mean square error between the parameter vector and its online estimate* when the number of examples goes to infinity (see section 5). In this context, another nontrivial difference with respect to the classical LQG problem arises. When computing certain expectations conditioned on the current information vector, one has to take into account that the information vector at a generic stage has additional components deriving from the realizations of the output measurement matrices up to the stage itself (as these are random matrices, associated with the input examples). As a consequence, the Kalman gain matrix, which is shown in the letter to be embedded in the optimal solution, is not only stage dependent but also a random matrix (although it becomes deterministic when conditioned on the input examples already presented to the learning machine). This motivates the investigation of issues such as its convergence in probability and the convergence of its expectation when the number of examples goes to infinity.

- *We discuss the connection of the proposed online learning framework with average regret minimization.* We prove that the sequence of our estimates minimizes the average regret functional (see section 6.1).

- *We investigate the connections among our solution, the Kalman filter estimate, and stochastic gradient descent* (see remark 3 and sections 6.2–6.4). We prove that our solution can be interpreted as a smoothed Kalman filter estimate, with time-varying gain matrix, and we show that it outperforms the latter in terms of its larger smoothness (section 6.3; see also section 8 *e* and its smaller sensitivity to outliers, section 6.4).

- *We discuss cases in which the proposed solution can be computed efficiently* (see, e.g., the comments presented after proposition 2, remark 10, and the numerical results reported in section 6.3).

Table 2: Acronyms of Frequent Use.

Problem OLL_γ^N	Online learning problem over finite horizon N and regularization parameter γ
Problem OLL_γ^∞	Online learning problem over infinite horizon and with regularization parameter γ
OLL estimate	Estimate obtained solving problem OLL_γ^N or OLL_γ^∞
LQ	Linear quadratic
LQG	Linear quadratic gaussian
LQR	Linear quadratic regulator
ARE	Average Riccati equation
SRE	Stochastic Riccati equation
KF	Kalman filter
MSE	Mean square error

• *We address the case of nonlinear input-output relationships, modeled using the kernel trick of kernel machines (see section 7).* As is well known, the kernel trick is based on a preliminary (in general nonlinear) mapping of the input space to a larger-dimensional feature space, to which the original linear model is applied in a second step. The kernel trick, which consists of computing certain inner products in the auxiliary feature space through a suitable function called “kernel,” can be applied in our context since we show that the optimal solution can be expressed in terms of inner products in the feature space that can be computed through a kernel.

• *We describe various other extensions (see section 8),* such as the case of a time-varying parameter vector to be learned, the introduction of a discount factor, the inclusion of additional regularization terms, a continuous-time extension framework, and a possible extension of the problem formulation through techniques from robust estimation and control.

Table 2 collects some acronyms of frequent use in the letter.

3 The Basic Case: Discrete Time, Finite Horizon, and Linear Model

For simplicity, we consider first a discrete time setting with a finite learning horizon and a linear model. Then we address the extensions to an infinite learning horizon and nonlinear models.

3.1 Problem Formulation.

Assumption 1 (linear data generation model). At each time $k = 0, 1, \dots$, a learning machine can observe the supervised pair (x_k, y_k) , where $x_k \in \mathbb{R}^d$ is a column vector and $y_k \in \mathbb{R}$. The output y_k is generated from the input x_k according to the following linear model:

$$y_k = w'x_k + \varepsilon_k, \quad (3.1)$$

where $\varepsilon_k \in \mathbb{R}$ is a measurement noise and $w \in \mathbb{R}^d$ is a random column vector, unknown to the learning machine, and to be estimated by the learning machine itself by using the sequence of examples (x_k, y_k) as they become available.

Assumption 2 (random variables). The random variables w , $\{x_k\}$, $\{\varepsilon_k\}$ are mutually independent⁵ and (only for simplicity of notation and without any loss of generality) have mean 0. The random variables ε_k have the same finite variance σ_ε^2 , and each x_k has finite covariance matrix $\mathbb{E}\{x_k x_k'\}$.

Assumption 3 (learning machine). Starting from the initial estimate $\hat{w}_0 := 0$ of w , at each time $k + 1 = 1, 2, \dots$, the learning machine builds an estimate \hat{w}_k of w , generated according to

$$\hat{w}_{k+1} = \hat{w}_k + u_k, \quad (3.2)$$

where u_k is the update of the estimate of w at the time k (to be optimized according to a suitable optimality criterion, defined later on).

Remark 1. It is important to observe that model 3.1 is time invariant in the sense that the same w is used to generate every y_k , starting from every x_k and every ε_k .⁶ So once a realization of the random vector w has been generated, this can be interpreted as a fixed vector, to be estimated by the learning machine using the online supervised examples.

To analyze the time evolution of the estimate, one has to consider the following dynamical system (see Alessandri, Sanguineti, & Maggiore, 2002 for a similar approach), with state vector (w_k', \hat{w}_k') and initial conditions $w_0 := w$ and $\hat{w}_0 := 0$:

$$\begin{cases} w_{k+1} = w_k, \\ \hat{w}_{k+1} = \hat{w}_k + u_k, \end{cases} \quad (3.3)$$

together with the measures

$$y_k = C_k w_k + \varepsilon_k, \quad (3.4)$$

where $C_k := x_k'$.

⁵As an extension, one could consider the case in which the inputs x_k are generated by the learning machine as the states of another controlled dynamical system. This, together with the optimization of a suitable learning functional similar to the one presented in equation 3.5, would model the problem of online active learning, as the learning machine would have an influence even on the choice of the sequence of input examples (see item n in section 8).

⁶An extension to the case of a (slowly) time-varying parameter vector is discussed in item e of section 8.

Assumption 4 (available information and updating functions). The update u_k at the time k is chosen as a function $u_k(I_k)$, called the updating function, of the information vector I_k at the same time, which collects the “history” up to the time k and is defined as

$$I_k := \{(x_j, y_j) \text{ for } j = 0, \dots, k, \text{ and } u_j \text{ for } j = 0, \dots, k - 1\}$$

for $k = 1, 2, \dots$, and

$$I_0 := \{(x_0, y_0)\}.$$

Hence, the update u_k depends only on the sequence of examples (x_j, y_j) observed up to the current stage k and on the sequence of previous updates u_j (or, equivalently, since $\hat{w}_0 = 0$, on the sequence of previous updates of the estimate of w).

In our model, the updating functions u_k are chosen in order to minimize a learning functional over a finite learning horizon, defined as follows:

Definition 1 (learning functional over horizon N). Let N be a positive integer, $\gamma > 0$, $Q_k := x_k x_k'$, and

$$\begin{aligned} & J_\gamma^N \left(\{u_k(I_k)\}_{k=0}^{N-1} \right) \\ & := \mathbb{E}_{w, \{x_k\}_{k=0}^N, \{\varepsilon_k\}_{k=0}^{N-1}} \left\{ \sum_{k=0}^{N-1} \left[((\hat{w}_k - w_k)' x_k)^2 + \gamma u_k' u_k \right] + ((\hat{w}_N - w_N)' x_N)^2 \right\} \\ & = \mathbb{E}_{w, \{x_k\}_{k=0}^N, \{\varepsilon_k\}_{k=0}^{N-1}} \left\{ \sum_{k=0}^{N-1} \left[(\hat{w}_k - w_k)' Q_k (\hat{w}_k - w_k) + \gamma u_k' u_k \right] \right. \\ & \quad \left. + (\hat{w}_N - w_N)' Q_N (\hat{w}_N - w_N) \right\}. \end{aligned} \tag{3.5}$$

We state the following online learning problem (in this letter, the symbol “ \circ ” is used to denote optimality):

Problem OLL $_\gamma^N$ (online learning over a finite horizon). Given the finite learning horizon N , the examples (x_k, y_k) generated at each time instant $k = 0, 1, \dots, N$ according to the model defined by assumptions 1 and 2, and the learning machine defined by assumption 3, find the finite sequence $u_0^\circ(I_0), \dots, u_{N-1}^\circ(I_{N-1})$ of optimal updating functions with the structure defined by assumption 4, that minimizes the learning functional 3.5.

Problem OLL $_\gamma^N$ can be considered a parameter identification problem or an optimal estimation problem, as the final goal consists of estimating optimally the parameter vector w relating the input examples to their outputs,

given the current subsequence of examples and the adopted optimality criterion. It can also be considered an optimal control problem, interpreting the updating function u_k as a control function for the dynamical system 3.3. Although this last interpretation may seem less natural than the first two, it is motivated by the fact that problem OLL_γ^N and its variations presented later in the letter can be investigated using optimal control techniques, as is done in the following.

For every $k = 0, 1, \dots, N$, we shall call \hat{w}_k *online estimate* (*OLL estimate*, for short).

Remark 2. The term $((\hat{w}_k - w_k)'x_k)^2$ in the learning functional, equation 3.5, penalizes a large deviation of the learning machine estimate $\hat{w}_k'x_k$ of the label y_k from its best estimate (in a mean square sense) $w_k'x_k = w'x_k$, which would have been obtained if w were known, whereas the term $u_k'u_k$ penalizes a large square of the norm of the update u_k of the estimate of w , and γ is a regularization term, which measures the trade-off between the two terms.

Remark 3. “Shifted” Kalman filter estimates correspond to the limit case $\gamma = 0$ in the formulation of problem OLL_γ^N . Indeed, in such a case, each term

$$\mathbb{E}_{w, \{x_i\}_{i=0}^k, \{\varepsilon_i\}_{i=0}^{k-1}} \{(\hat{w}_k - w_k)'Q_k(\hat{w}_k - w_k)\}$$

in equation 3.5 is minimized when \hat{w}_k is the conditional expectation of w_k given I_{k-1} , that is, when it is the Kalman filter estimate of w_k at time $k - 1$ (see, e.g., Bertsekas, 1995).⁷ It is worth mentioning that since the parameter vector to be learned is constant and the data generation model is described by equation 3.1, the specific Kalman estimation problem is equivalent to recursive least squares (see Sayed, 2003 for a proof of this equivalence).

In sections 6.3 and 6.4, we discuss some relationships of the proposed learning framework with the classical Kalman filter (Sutton, 1992). As it will be shown by proposition 8 and by the numerical results in Figure 1, the presence of the regularization term in the learning functional 3.5 can make the resulting sequence of optimal estimates of w smoother with respect to the time index, and less sensitive to outliers, than the sequence of estimates obtained by using the classical Kalman filter, under gaussian assumptions on the random variables w and ε_k .

Remark 4. The constraint that each update u_k (hence, also each optimal updating function u_k°) depends only on the sequence of examples (x_j, y_j)

⁷ Note that Bertsekas (1995, proposition E.1) is formulated in terms of the square of the Euclidean norm of the error vector, which is $\hat{w}_k - w_k$ in our case. However, the proposition can be still applied if one moves from the square of the Euclidean norm to the square of the (semi)norm induced by Q_k (as in the present case).

observed up to the current stage k and on the sequence of previous updates u_j implies that no future examples are taken into account to update the current estimate of w . Hence, the proposed solution is actually a model of online learning. Instead, batch learning corresponds to the case where one assumes that all the sequence $\{(x_j, y_j), j = 0, \dots, N\}$ of examples is known to the learning machine, starting from time $k = 0$.

Remark 5. An alternative definition of the learning functional can be obtained by replacing the term $((\hat{w}_k - w_k)'x_k)^2$ in 3.5 by $(\hat{w}'_k x_k - y_k)^2$, that is, by the square of the difference between the label estimated by the learning machine before measuring y_k (but knowing x_k) and the label y_k generated by model 3.1 at the time k (note that, different from the term $w'_k x_k$, both are observable at time k). However, by taking expectations and recalling that ε_k has mean 0 and is mutually independent from x_k, w_k , and \hat{w}_k , one obtains

$$\begin{aligned}
 & J_{\gamma, y}^N (\{u_k(I_k)\}_{k=0}^{N-1}) \\
 & := \mathbb{E}_{w, \{x_k\}_{k=0}^N, \{\varepsilon_k\}_{k=0}^N} \left\{ \sum_{k=0}^{N-1} [(\hat{w}'_k x_k - y_k)^2 + \gamma u'_k u_k] + (\hat{w}'_N x_N - y_N)^2 \right\} \\
 & = \mathbb{E}_{w, \{x_k\}_{k=0}^N, \{\varepsilon_k\}_{k=0}^N} \left\{ \sum_{k=0}^{N-1} [(\hat{w}_k - w_k)' Q_k (\hat{w}_k - w_k) + \gamma u'_k u_k] \right. \\
 & \quad \left. + (\hat{w}_N - w_N)' Q_N (\hat{w}_N - w_N) \right\} + (N + 1) \sigma_\varepsilon^2. \tag{3.6}
 \end{aligned}$$

Hence, since the last term in equation 3.6 is a constant, the learning functionals 3.5 and 3.6 have the same sequence of optimal updating functions. It is worth noting that in both formulas 3.5 and 3.6, in order to generate the estimates \hat{w}_k , one uses only the probability distribution of w_k conditioned on the already available observations.

The statement of problem OLL_γ^N can be simplified by defining the learning error,

$$e_k := \hat{w}_k - w_k,$$

which evolves according to

$$e_{k+1} = e_k + u_k, \tag{3.7}$$

where

$$e_0 := -w_0 = -w.$$

Of course, $e_k \simeq 0$ means $\hat{w}_k \simeq w_k = w$. Moreover, since both \hat{w}_k and x_k are known at time k , one can replace the measures y_k by

$$\tilde{y}_k := \hat{w}'_k x_k - y_k,$$

hence obtaining the measurement equation

$$\tilde{y}_k = C_k e_k + \tilde{\varepsilon}_k, \quad (3.8)$$

where $\tilde{\varepsilon}_k := -\varepsilon_k$, and has the same variance σ_ε^2 as ε_k . In this case, the “history” of the learning machine, measures, and past updates up to time k is collected in the new information vector \tilde{I}_k , defined as

$$\tilde{I}_k := \{(x_j, \tilde{y}_j) \text{ for } j = 0, \dots, k, \text{ and } u_j \text{ for } j = 0, \dots, k-1\}$$

for $k = 1, 2, \dots$ and

$$\tilde{I}_0 := \{(x_0, \tilde{y}_0)\}.$$

There is a one-to-one correspondence between the information vectors I_k and \tilde{I}_k . So, the optimization of the learning functional 3.5 assuming that the dynamical system evolves according to equation 3.3, the sequence of measures is provided by equation 3.4, and the updating functions u_k have the form $u_k(I_k)$, is equivalent to the optimization of the following learning functional:

$$\begin{aligned} & \tilde{J}_\gamma^N \left(\{u_k(\tilde{I}_k)\}_{k=0}^{N-1} \right) \\ & := \mathbb{E}_{e_0, \{x_k\}_{k=0}^N, \{\tilde{\varepsilon}_k\}_{k=0}^{N-1}} \left\{ \sum_{k=0}^{N-1} [(e'_k x_k)^2 + \gamma u'_k u_k] + (e'_N x_N)^2 \right\} \\ & = \mathbb{E}_{e_0, \{x_k\}_{k=0}^N, \{\tilde{\varepsilon}_k\}_{k=0}^{N-1}} \left\{ \sum_{k=0}^{N-1} [e'_k Q_k e_k + \gamma u'_k u_k] + e'_N Q_N e_N \right\}, \quad (3.9) \end{aligned}$$

assuming that the error vector evolves according to equation 3.7, the sequence of measures is provided by equation 3.8, and the update u_k is now a function $u_k(\tilde{I}_k)$ of the information vector \tilde{I}_k . Such a problem is a nontrivial variation of the classical LQ problem (Bertsekas, 1995). Whereas in the latter, the matrices C_k and Q_k are deterministic, in the proposed formulation of online learning, they are random, since they depend on the input examples x_k . Another difference is that for $j = 0, \dots, k$, the information vector \tilde{I}_k includes the realizations of the inputs x_j , hence also of the matrices C_j and Q_j .

3.2 Solution of the Finite Horizon Online Learning Problem. To solve problem OLL_γ^N , we make extensive use of the concept of the cost-to-go function from the theory of dynamic programming (Bertsekas, 1995). In our context, the cost-to-go function \tilde{J}_k° at the time stage $k = 0, \dots, N - 1$ is defined as

$$\tilde{J}_k^\circ(\tilde{I}_k) := \inf_{\{u_j(\tilde{I}_j)\}_{j=k}^{N-1}} \mathbb{E}_{e_k, \{x_j\}_{j=k+1}^N, \{\tilde{\varepsilon}_j\}_{j=k+1}^{N-1}} \left\{ \sum_{j=k}^{N-1} [e'_j Q_j e_j + \gamma u'_j u_j] + e'_N Q_N e_N \middle| \tilde{I}_k \right\}, \quad (3.10)$$

whereas

$$\tilde{J}_N^\circ(\tilde{I}_N) = \mathbb{E}_{e_N} \{e'_N Q_N e_N | \tilde{I}_N\}. \quad (3.11)$$

Finally, the optimal value of the learning functional equation 3.9, is

$$\tilde{J}_0^\circ = \mathbb{E}_{\tilde{I}_0} \{\tilde{J}_0^\circ(\tilde{I}_0)\}.$$

Under mild regularity conditions (see remark 6), the cost-to-go functions can be determined recursively by solving the Bellman Equations,

$$\tilde{J}_k^\circ(\tilde{I}_k) = \inf_{u_k \in \mathbb{R}^d} \mathbb{E}_{e_k, \tilde{I}_{k+1}} \{e'_k Q_k e_k + \gamma u'_k u_k + \tilde{J}_{k+1}^\circ(\tilde{I}_{k+1}) | \tilde{I}_k, u_k\}, \quad (3.12)$$

for $k = N - 1, \dots, 0$.

Remark 6. The regularity conditions mentioned above are satisfied, for example, in the case (studied in this letter), where the random vectors w and ε_k are gaussian. Indeed, in such a context, the optimal updating functions that will be provided by equation 3.13 are linear with respect to the information vector (Bertsekas, 1995; Bertsekas & Shrieve, 1978). In the following, the mild regularity conditions mentioned above will be implicitly assumed throughout the letter.

Equations 3.11 and 3.12 are similar to those for the cost-to-go functions in the LQ problem (see, e.g., Bertsekas, 1995), with the difference that in the present context, the matrices Q_k and C_k are random. Moreover, both matrices become known to the learning machine at time k , as they can be derived from the information vector \tilde{I}_k . In the following, we sometimes use the superscript “ \circ ” not only for the optimal updating functions but also to denote vectors (e.g, \hat{w}_k and e_k) evaluated when the sequence of optimal updating functions, equation 3.13, is applied.

Proposition 1 (optimal updating functions and average Riccati equation). *Let assumptions 1–4 be satisfied. Then the updating functions that solve problem OLL_γ^N are given, for $k = N - 1, \dots, 0$, by*

$$u_k^\circ(\tilde{I}_k) = L_k \mathbb{E}_{e_k^\circ} \{e_k^\circ | \tilde{I}_k\}, \quad (3.13)$$

where

$$L_k := -(\bar{K}_{k+1} + \gamma I)^{-1} \bar{K}_{k+1}, \quad (3.14)$$

and the matrices

$$K_k := \bar{K}_{k+1} - \bar{K}_{k+1}(\bar{K}_{k+1} + \gamma I)^{-1} \bar{K}_{k+1} + Q_k, \quad (3.15)$$

$$F_k := \bar{K}_{k+1}(\bar{K}_{k+1} + \gamma I)^{-1} \bar{K}_{k+1}, \quad (3.16)$$

and

$$\bar{K}_k := \mathbb{E}_{K_k} \{K_k\} = \bar{K}_{k+1} - \bar{K}_{k+1}(\bar{K}_{k+1} + \gamma I)^{-1} \bar{K}_{k+1} + \mathbb{E}_{Q_k} \{Q_k\} \quad (3.17)$$

are symmetric positive-semidefinite. The recursions above are initialized by

$$K_N := Q_N \quad (3.18)$$

and

$$\bar{K}_N := \mathbb{E}_{K_N} \{K_N\}. \quad (3.19)$$

Equation 3.17 can be called an average Riccati equation (ARE) since it contains the expectation term $\mathbb{E}_{Q_k} \{Q_k\}$. In practice, it can be solved offline, as can the classical deterministic Riccati equation of the linear quadratic regulator (LQR) subproblem (Bertsekas, 1995), simply by replacing Q_k (which is deterministic in the LQ problem) by $\mathbb{E}_{Q_k} \{Q_k\}$. It is worth remarking that solving the ARE, equation 3.17, does not require the knowledge of future input examples and that all the matrices L_k in equation 3.14 have spectral radius⁸ $|\lambda|_{\max}(L_k)$ strictly smaller than 1. Finally, the matrices F_k are reported in formula 3.16 because they are used to express $\tilde{f}_k^\circ(\tilde{I}_k)$ (see formula (A.18) in the appendix). They are also used in the infinite-horizon version of problem

⁸For a square matrix M , we denote by $|\lambda|_{\max}(M)$ its spectral radius.

OLL $_v^N$ (problem OLL $_v^\infty$), to reduce one part of the proof of proposition 4 in section 4 to the finite-horizon case.

Due to equation 3.13, in order to generate the optimal update $u_k^\circ(\tilde{I}_k)$ at time k , one has to compute $\mathbb{E}_{e_k^\circ} \{e_k^\circ | \tilde{I}_k\}$. Let us now make the following additional assumption:

Assumption 5 (gaussian random variables). The random variables w and ε_k are gaussian.

The next proposition shows that when the additional assumption 5 is satisfied, the optimal estimate \hat{w}_k° of the proposed framework tracks the (usually time-varying) Kalman filter estimate. Indeed, inspection of its proof shows that

$$\hat{e}_k^{\circ,\dagger} := \mathbb{E}_{e_k^\circ} \{e_k^\circ | \tilde{I}_k\}$$

is the Kalman filter (KF estimate) of the error vector e_k° at the time k , based on the information vector \tilde{I}_k , thus getting a Kalman filter recursion scheme.

In the following, we denote by

$$\hat{w}_k^\dagger := \hat{w}_k^\circ - \hat{e}_k^{\circ,\dagger}$$

the KF estimate of $w_k = w$ at the time k , based on the information vector I_k (or, equivalently, on the corresponding information vector \tilde{I}_k). Moreover, for $k = 0, 1, \dots$, let

$$\Sigma_k := \mathbb{E}_{e_k} \{ (e_k - \mathbb{E}_{e_k} \{e_k | \tilde{I}_k\}) (e_k - \mathbb{E}_{e_k} \{e_k | \tilde{I}_k\})' | \tilde{I}_k \} \quad (3.20)$$

be the (conditional) covariance matrix⁹ of e_k , conditioned on the information vector \tilde{I}_k , and

$$\Sigma_{-1} := \mathbb{E}_{e_0} \left\{ \left(e_0 - \mathbb{E}_{e_0} \{e_0\} \right) \left(e_0 - \mathbb{E}_{e_0} \{e_0\} \right)' \right\} = \Sigma_w \quad (3.21)$$

the (unconditional) covariance matrix¹⁰ of e_0 , which is equal to the (unconditional) covariance matrix of w .

⁹Here, the superscript \circ is omitted to highlight that expression 3.20 (and other expressions presented later, such as 3.25) holds also when e_k is not evaluated in correspondence of the sequence of optimal updating functions, equation 3.13.

¹⁰Proceeding as in Bertsekas (1995, appendix E.4), one could use the symbol $\Sigma_{k|k}$ to denote the (conditional) covariance matrix Σ_k just introduced to distinguish it from the (conditional) covariance matrix of e_{k+1} , conditioned on the information vector \tilde{I}_k and denoted by $\Sigma_{k+1|k}$. However, in the specific case, they are equal, so both are denoted by Σ_k .

Proposition 2 (*optimal online estimate and Stochastic Riccati Equation*). *Let assumptions 1–5 be satisfied. Then*

$$\begin{aligned}\hat{w}_{k+1}^\circ &= \hat{w}_k^\circ + L_k \left(\hat{w}_k^\circ - \frac{\mathbb{E}\{w|I_k\}}{w} \right) \\ &= \hat{w}_k^\circ + L_k(\hat{w}_k^\circ - \hat{w}_k^\dagger) = \hat{w}_k^\circ + L_k(\hat{e}_k^\circ - \hat{e}_k^{\circ,\dagger}),\end{aligned}\quad (3.22)$$

where, for $k = -1, 0, \dots$

$$\hat{w}_{k+1}^\dagger = \hat{w}_k^\dagger + H_{k+1}(y_{k+1} - C_{k+1}\hat{w}_k^\dagger), \quad (3.23)$$

$$H_{k+1} := \Sigma_{k+1}C'_{k+1}(\sigma_\varepsilon^2)^{-1}, \quad (3.24)$$

and, for $k = 0, 1, \dots$,

$$\Sigma_k = \Sigma_{k-1} - \Sigma_{k-1}C'_k(C_k\Sigma_{k-1}C'_k + \sigma_\varepsilon^2)^{-1}C_k\Sigma_{k-1}, \quad (3.25)$$

with the initializations

$$\hat{w}_{-1}^\dagger = 0, \quad (3.26)$$

$$\hat{w}_{-1}^\circ = 0, \quad (3.27)$$

and

$$L_{-1} = -(\bar{K}_0 + \gamma I)^{-1} \bar{K}_0. \quad (3.28)$$

Equation 3.25 has the form of the Riccati equation of the well-known Kalman filter. Due to the stochastic nature of C_k , it can be called a stochastic Riccati equation (SRE). From a computational point of view, solving equation 3.25 is easy even in a high-dimensional setting (i.e., when the dimension d of the input space is large). Indeed, $C_k\Sigma_{k-1}C'_k + \sigma_\varepsilon^2$ (which needs to be inverted in equation 3.25) is a scalar. Similarly, in formula 3.24, one has to invert the scalar σ_ε^2 . In other applications of the Kalman filter, one has to invert matrices. Moreover, differently from the average Riccati equation, the stochastic Riccati equation has to be solved online.

Remark 7. It is worth mentioning that also Bertsekas (1995) investigates an LQ optimal control problem with random matrices. In that case, however, there is perfect information on the state, and the randomness is limited to the system dynamics. For that problem, a suitable average Riccati equation is obtained, but no stochastic Riccati equation. Hence, that formulation, though inspiring for our work presented in this letter, cannot be applied directly to our online learning framework.

Remark 8. Equations 3.13 and 3.23 show that the classical separation principle of control and estimation holds also for problem OLL_γ^N . More precisely,

it is reduced to two subproblems, which can be solved independently: the determination of the matrices L_k (solution of the LQR subproblem) and the determination of the Kalman gain matrices H_k (solution of the Kalman filter estimation subproblem). One might wonder why, in problem OLL_γ^N one gets, instead of the classical Riccati equation, two different kinds of equations for the two subproblems, the ARE, equation 3.17, and the SRE, equation 3.25, in spite of the well-known duality between the LQR subproblem and the Kalman filter estimation problem (Söderström, 2002). The reason is that when moving from the LQR subproblem to the Kalman filter estimation subproblem, the roles of the matrices

$$A_k := I, B_k := I, Q_k, R_k := \gamma I$$

in the primal problem (i.e., the LQR subproblem) are played, respectively, by the following matrices of the dual problem (i.e., the Kalman filter estimation problem):

$$A_k^{\text{dual}} := A_k' = I, B_k^{\text{dual}} := C_k', Q_k^{\text{dual}} := 0, R_k^{\text{dual}} := \sigma_\varepsilon^2,$$

where Q_k^{dual} is the covariance matrix of the system noise (a kind of noise that is not present in model 3.7), hence it is an all-0s matrix. Now, in our specific context, in the primal problem, the matrix Q_k is random, whereas in the dual problem, the matrix Q_k^{dual} is deterministic. Similarly, in the primal problem, the matrix B_k is deterministic, whereas in the dual problem, the matrix B_k^{dual} is random. This lack of symmetry is the reason that the two Riccati equations 3.17 and 3.25, have different forms.

The next proposition states some properties of the solution to the SRE. For two symmetric square matrices S_1 and S_2 of the same dimension, $S_1 \leq S_2$ means that $S_2 - S_1$ is symmetric and positive-semidefinite. When it is evident from the context, we use the symbol 0 to denote a matrix whose elements are all equal to 0.

Proposition 3 (properties of the solution to the SRE). *Let assumptions 1–5 be satisfied. Then*

i.

$$0 \leq \Sigma_{k+1} \leq \Sigma_k \tag{3.29}$$

(i.e., the sequence is “nonnegative” and monotonic “nonincreasing” in a generalized sense, according to \leq) for all the realizations of the random matrices Σ_{k+1} and Σ_k .

ii. For all the realizations of these random matrices,

$$0 \leq \text{Tr}\{\Sigma_{k+1}\} \leq \text{Tr}\{\Sigma_k\} \tag{3.30}$$

and

$$0 \leq \text{Tr}\{\Sigma_{k+1}^2\} \leq \text{Tr}\{\Sigma_k^2\}. \quad (3.31)$$

iii. There exists a symmetric and positive-semidefinite matrix $\bar{\Sigma}$ such that

$$\lim_{k \rightarrow +\infty} \mathbb{E}_{\Sigma_k}\{\Sigma_k\} = \bar{\Sigma}. \quad (3.32)$$

iv. If all the input examples x_k have a common probability distribution with bounded support and the same positive-definite covariance matrix \bar{Q} in such a way that

$$\mathbb{E}_{Q_k}\{Q_k\} = \bar{Q} \succeq 0 \quad (3.33)$$

for all k , then with a priori probability 1, one has

$$\lim_{k \rightarrow +\infty} \mathbb{E}_{\Sigma_k}\{\Sigma_k\} = \bar{\Sigma} = 0. \quad (3.34)$$

When equation 3.34 holds, then

$$\lim_{k \rightarrow +\infty} \text{Tr}\left\{\mathbb{E}_{\Sigma_k}\{\Sigma_k\}\right\} = \text{Tr}\{\bar{\Sigma}\} = 0. \quad (3.35)$$

v. For every $k = -1, 0, 1, 2, \dots$, and all the realizations of the random matrices,

$$\text{Tr}\{F_{k+1}\Sigma_{k+1}\} \leq \text{Tr}\{F_{k+1}\Sigma_k\} \leq \dots \leq \text{Tr}\{F_{k+1}\Sigma_{-1}\}. \quad (3.36)$$

An intuitive explanation of the second bound in equation 3.30 is the following. When the time index moves from k to $k + 1$, the new information acquired at the time $k + 1$ cannot deteriorate, on the average, the quality of the KF estimate, which is in accordance with its optimality properties (Bertsekas, 1995). Equations 3.29, 3.30, and 3.36 will be used, together with equations 3.34 and 3.35, in the convergence analysis of the proposed method for $k \rightarrow +\infty$ (see section 4 and the appendix).

Remark 9. An important assumption that is needed in the proof of proposition 3.iv is that the common covariance matrix \bar{Q} of the input examples is positive-definite. When this is not the case, this means that with probability 1, all the input examples belong to a finite-dimensional subspace \mathcal{S} of \mathbb{R}^d ; hence, with probability 1, it is not possible to extract from the input-output pairs (x_k, y_k) any information about the component of w that it is orthogonal to that subspace unless such a component is correlated with the projection of w on \mathcal{S} . However, one still has the convergence of both the KF estimate and the OLL estimate of w to the projection of w on \mathcal{S} , as can be shown by

setting the problem directly on S . Moreover, the possible absence of information in the data about the component of w that it is orthogonal to S has no negative consequences on the estimation process, in the sense that in order to compute $w'x$ for a possibly unseen input x , one needs, with probability 1, to know only the component of w that belongs to the subspace S .

3.3 Role of the Regularization Parameter. Let us investigate the behavior of the optimal updating functions provided by equations 3.13 and 3.14 for the two limit cases, $\gamma \simeq 0$ and $\gamma \rightarrow +\infty$, and for intermediate values of γ .

3.3.1 *The Case $\gamma \simeq 0$.* The penalization of the update u_k in the learning functional, equation 3.9, becomes negligible, and one obtains $L_k \simeq -I$ from equation 3.14, and

$$u_k^\circ \simeq -\mathbb{E}_{e_k^\circ} \{e_k^\circ | \tilde{I}_k\} \tag{3.37}$$

from equation 3.13. Hence, one gets (from equations 3.7 and 3.37),

$$e_{k+1}^\circ \simeq e_k^\circ - \mathbb{E}_{e_k^\circ} \{e_k^\circ | \tilde{I}_k\}.$$

Equivalently, in terms of the unknown vector w and its optimal estimates $\hat{w}_k^\circ, \hat{w}_{k+1}^\circ$ at the times k and $k + 1$, respectively, one obtains

$$(\hat{w}_{k+1}^\circ - w) \simeq (\hat{w}_k^\circ - w) - \left(\hat{w}_k^\circ - \mathbb{E}_w \{w | I_k\} \right),$$

hence

$$\hat{w}_{k+1}^\circ \simeq \mathbb{E}_w \{w | I_k\},$$

which is just the KF estimate of w at the time k , based on the information vector I_k .

3.3.2 *The Case $\gamma \rightarrow +\infty$.* The penalization of the update u_k in the learning functional 3.9 becomes larger and larger. Indeed, for γ large enough, one obtains $L_k \simeq 0$ from equation 3.14, and

$$u_k^\circ \simeq 0$$

from equation 3.13. Hence, one gets

$$e_{k+1}^\circ \simeq e_k^\circ$$

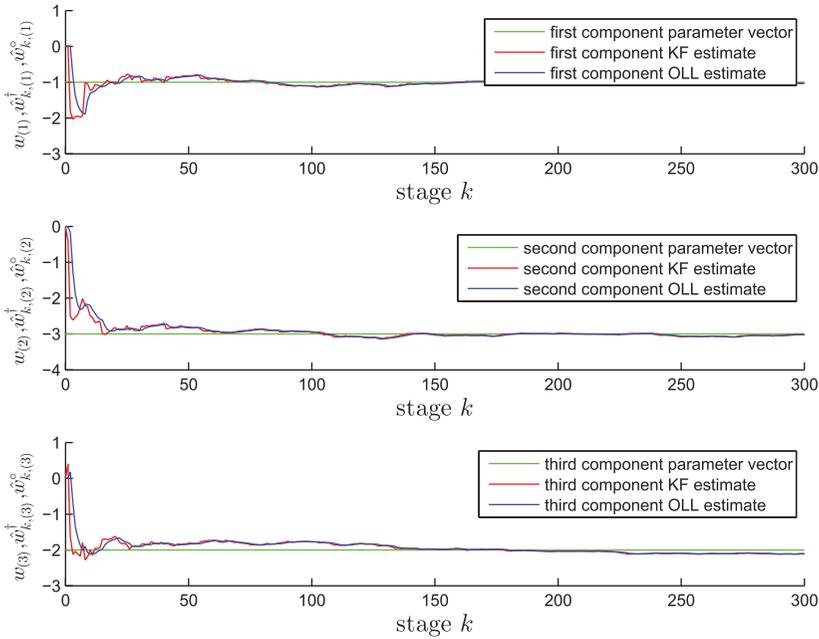


Figure 1: A comparison between the components of the OLL estimate \hat{w}_k° and the KF estimate \hat{w}_k^\dagger . A three-dimensional case has been considered, with the realization $w = (-1, -3, -2)'$, and $N + 1 = 301$ online examples (x_k, y_k) have been used to train the learning machine. The input examples have been generated with components mutually independent and uniformly distributed in $[-1, 1]$, whereas the covariance matrix Σ_w of w has been chosen to be diagonal with diagonal entries equal to 4, and the variance σ_ε^2 of the measurement noise is equal to 1, as is the regularization parameter γ .

and

$$\hat{w}_{k+1}^\circ \simeq \hat{w}_k^\circ \simeq \dots \simeq \hat{w}_0^\circ = 0.$$

3.3.3 *Intermediate Values of γ .* In this case the estimate \hat{w}_k° enjoys convergence properties similar to the ones of the KF estimate \hat{w}_k^\dagger , as illustrated numerically in Figure 1. Moreover, \hat{w}_k° is a smoothed version of the estimate \hat{w}_k^\dagger . The sequence of estimates \hat{w}_k° is smoother and less sensitive to outliers than the sequence of estimates \hat{w}_k^\dagger , as a large change in the estimate when moving from \hat{w}_k° to \hat{w}_{k+1}° is penalized by the presence of the term $\gamma u_k' u_k$ in the cost functional 3.9. This can be seen also by formula 3.22, as equation 3.14 implies that all the eigenvalues of the symmetric matrix L_k are inside the unit circle. A deeper investigation of these three issues (convergence,

smoothness, and sensitivity to outliers) is made in sections 5, 6.3, and 6.4, respectively.

4 LQG Learning Over an Infinite Horizon

To address the infinite-horizon case, we remove the final-stage cost $e'_N Q_N e_N$ (or, equivalently, we assume $x_N = 0$ with probability 1, hence also $Q_N = 0$ with probability 1) and let $N \rightarrow +\infty$ (the precise formulation is provided later in this section).

Assumption 6 (identical distributions of the input examples). The random variables $\{x_k\}$ are identically distributed and have the same positive-definite covariance matrix,

$$\bar{Q}_k := \mathbb{E}_{x_k} \{x_k x'_k\} = \bar{Q}$$

for every $k = 0, 1, \dots$. Moreover, the common probability distribution has bounded support.

Due to assumption 6, the analysis has some similarities with the one for the optimal solution to the LQG problem performed, for example, in Bertsekas (1995). We denote by $\bar{Q}^{-1/2}$ a symmetric and positive-definite square root of \bar{Q} . As one can check directly from the definitions of reachability and observability¹¹ (Antsaklis & Michel, 2007), we observe that the pair

$$(A, B) := (I, I)$$

is reachable,¹² whereas the pair

$$(A, C) := (I, \bar{Q}^{-1/2})$$

¹¹ Given a discrete-time and time-invariant linear dynamical system of the form

$$\begin{cases} z_{t+1} = Az_t + Bv_t, \\ \xi_t = Cz_t + Dv_t, \end{cases} \tag{4.1}$$

where $z_t \in \mathbb{R}^n$, $v_t \in \mathbb{R}^m$, $\xi_t \in \mathbb{R}^p$, the pair (A, B) is reachable if and only if, starting from any initial state, any other state can be reached at some subsequent finite time \bar{t} by choosing an appropriate sequence of \bar{t} controls. Moreover, the pair (A, C) is said to be observable if and only if, given any sequence of measures $\xi_0, \dots, \xi_{\bar{t}-1}$ and applied controls $v_0, \dots, v_{\bar{t}-1}$ for $\bar{t} \geq 1$ sufficiently long, it is possible to determine exactly the initial state $z_0 \in \mathbb{R}^n$ of the dynamical system, equation 4.1.

¹² Actually, the expression used in (Bertsekas, 1995, sec. 4.1, definition 1.1) for this situation is *controllable pair*, but the definition provided there is actually the one of “reachable pair” reported in note 11.

is observable. Hence, one can apply proposition 4.1 of Bertsekas (1995), from which it follows that the ARE, equation 3.17 admits a stationary solution \bar{K} , associated with the two stationary matrices

$$L := -(\bar{K} + \gamma I)^{-1} \bar{K} \tag{4.2}$$

and

$$F := \bar{K}(\bar{K} + \gamma I)^{-1} \bar{K}$$

(see 3.16). Moreover, by reversing the time indices in equation 3.17, (i.e., setting $t := N - k$ and $P_t := \bar{K}_{N-k}$), the solution P_{t+1} of the ARE

$$P_{t+1} = P_t - P_t(P_t + \gamma I)^{-1} P_t + \bar{Q} \tag{4.3}$$

(which is equivalent to equation 3.17), converges to \bar{K} for any initialization of the positive-semidefinite matrix P_0 , still by Bertsekas (1995, proposition 4.1).

The (average) learning functional over infinite horizon is defined as follows:

Definition 2 (average learning functional over infinite horizon). Let $\gamma > 0$, and

$$\begin{aligned} J_\gamma^\infty (\{u_k(\tilde{I}_k)\}_{k=0}^\infty) &:= \liminf_{N \rightarrow +\infty} \left(\frac{1}{N} \mathbb{E}_{w, \{x_k\}_{k=0}^{N-1}, \{\varepsilon_k\}_{k=0}^{N-2}} \left\{ \sum_{k=0}^{N-1} [e'_k Q_k e_k + \gamma u'_k u_k] \right\} \right) \\ &= \liminf_{N \rightarrow +\infty} \left(\frac{1}{N} \mathbb{E}_{w, \{x_k\}_{k=0}^{N-1}, \{\varepsilon_k\}_{k=0}^{N-2}} \left\{ \sum_{k=0}^{N-1} [(e'_k x_k)^2 + \gamma u'_k u_k] \right\} \right). \end{aligned} \tag{4.4}$$

Problem OLL $_\gamma^\infty$ (online learning over infinite horizon). Given the examples (x_k, y_k) generated at each time instant $k = 0, 1, \dots$, according to the model defined by assumptions 1 and 2, and the learning machine defined by assumption 3, find the infinite sequence $u_0^\circ(I_0), u_1^\circ(I_1), \dots$, of optimal updating functions with the structure defined by assumption 4 that minimizes the average learning functional 4.4.

As with problem OLL $_\gamma^N$, for every $k = 0, 1, \dots$, we shall call \hat{w}_k online estimate (OLL estimate, for short). In the following, we consider directly the LQG case, with identical distributions for the input examples.

Proposition 4 (optimal updating functions and ARE). Let assumptions 1–6 be satisfied. Then, updating functions that solve problem OLL $_\gamma^\infty$ are given, for $k = 0, 1, \dots$, by

$$u_k^\circ(\tilde{l}_k) = L \mathbb{E}_{e_k^\circ} \{e_k^\circ | \tilde{l}_k\}, \tag{4.5}$$

where L is defined in equation 4.2.¹³

Remark 10. A nice feature of the AREs, equation 3.17 (with $\mathbb{E}\{Q_k\} = \bar{Q}$) and equation 4.3 is that their common stationary solution \bar{K} can be easily expressed in terms of the eigenvalues/eigenvectors of the matrix \bar{Q} , which can be useful from a computational point of view. Indeed, let us express \bar{Q} as

$$\bar{Q} = U \Lambda_{\bar{Q}} U',$$

where U is a basis of orthogonal unit-norm eigenvectors of \bar{Q} (hence, $U' = U^{-1}$) and Λ is a diagonal matrix collecting the corresponding positive eigenvalues. We recall that \bar{K} satisfies

$$\bar{K} = \bar{K} - \bar{K}(\bar{K} + \gamma I)^{-1} \bar{K} + \bar{Q},$$

that is,

$$\bar{K}(\bar{K} + \gamma I)^{-1} \bar{K} = \bar{Q}. \tag{4.6}$$

Now, $\bar{K}(\bar{K} + \gamma I)^{-1} \bar{K}$ has the same eigenvectors as \bar{K} . Hence, \bar{K} and \bar{Q} also have the same eigenvectors, so \bar{K} is expressed as

$$\bar{K} = U \Lambda_{\bar{K}} U', \tag{4.7}$$

where $\Lambda_{\bar{K}}$ is a suitable diagonal matrix, with positive eigenvalues. Due to equation 4.6, the diagonal elements $(\Lambda_{\bar{K}})_{(i,i)}$ and $(\Lambda_{\bar{Q}})_{(i,i)}$ are related by

$$(\Lambda_{\bar{K}})_{(i,i)}^2 ((\Lambda_{\bar{K}})_{(i,i)} + \gamma)^{-1} = (\Lambda_{\bar{Q}})_{(i,i)}.$$

Hence, by the positiveness of $(\Lambda_{\bar{K}})_{(i,i)}$, one gets

$$(\Lambda_{\bar{K}})_{(i,i)} = \frac{(\Lambda_{\bar{Q}})_{(i,i)} + \sqrt{(\Lambda_{\bar{Q}})_{(i,i)}^2 + 4\gamma(\Lambda_{\bar{Q}})_{(i,i)}}}{2}.$$

Similarly, the stationary matrix L can be expressed as

$$L = U \Lambda_L U', \tag{4.8}$$

¹³ Here, we recall that $\mathbb{E}\{e_k^\circ | \tilde{l}_k\}$ is the KF estimate of the error vector e_k° at the time k , based on the information vector \tilde{l}_k .

where the elements $(\Lambda_L)_{(i,i)}$ of the diagonal matrix Λ_L are

$$(\Lambda_L)_{(i,i)} = -\frac{(\Lambda_{\bar{K}})_{(i,i)}}{(\Lambda_{\bar{K}})_{(i,i)} + \gamma} = -\frac{(\Lambda_{\bar{Q}})_{(i,i)} + \sqrt{(\Lambda_{\bar{Q}})_{(i,i)}^2 + 4\gamma(\Lambda_{\bar{Q}})_{(i,i)}}}{(\Lambda_{\bar{Q}})_{(i,i)} + \sqrt{(\Lambda_{\bar{Q}})_{(i,i)}^2 + 4\gamma(\Lambda_{\bar{Q}})_{(i,i)} + 2\gamma}}. \tag{4.9}$$

A particularly simple case occurs when the matrix \bar{Q} is diagonal; hence, one can choose the matrices U and U' as the identity matrix I , so \bar{K} and L are also diagonal, by formulas 4.7 and 4.8. Moreover, if \bar{Q} is proportional to the identity matrix I , \bar{K} and L are also proportional to I . Finally, similar remarks also hold in the finite-horizon case for the matrices \bar{K}_k and L_k in case the matrices $\mathbb{E}_{Q_k}\{Q_k\}$ commute (this happens, for example, when all the matrices $\mathbb{E}_{Q_k}\{Q_k\}$ are equal to \bar{Q}).

5 Convergence Properties of the Online Learning Estimates in Terms of Mean Square Errors

Let us denote by

$$MSE_k^\dagger := \mathbb{E}_{w, \hat{w}_k^\dagger} \left\{ (w - \hat{w}_k^\dagger)' (w - \hat{w}_k^\dagger) \right\}$$

the mean-square error of the KF estimate at time k . Different from the LQG case detailed in Bertsekas (1995), the expectation of Σ_k is needed in the following, as Σ_k depends on the sequence of random matrices C_0, \dots, C_k . Under the assumptions of proposition 3.iv, this converges to the 0 matrix as k tends to $+\infty$ (see formula 3.34).

Proposition 5 (convergence of the MSE of the KF estimate). *Let assumptions 1–6 be satisfied. Then the following hold:*

i.

$$MSE_k^\dagger = Tr \left\{ \mathbb{E}_{\Sigma_k} \{ \Sigma_k \} \right\}.$$

ii. For every $k = 1, 2, \dots$,

$$Tr \left\{ \mathbb{E}_{\Sigma_k} \{ \Sigma_k \} \right\} \leq \sqrt{\frac{(c_1 + \sigma_\varepsilon^2) d Tr \{ \mathbb{E}_{\Sigma_0} \{ \Sigma_0 \} \}}{k \lambda_{\min}(\bar{Q})}}, \tag{5.1}$$

where c_1 is a positive constant such that $C_{k+1} \Sigma_{-1} C'_{k+1} \leq c_1$ with a priori probability 1. Moreover, $\lim_{k \rightarrow +\infty} MSE_k^\dagger = 0$.

iii.

$$\lim_{k \rightarrow +\infty} \mathbb{E}\{H_k\} = 0. \tag{5.2}$$

iv. Every element $H_{k,(h,l)}$ of H_k converges to 0 also in probability, that is, for every $\delta > 0$,

$$\lim_{k \rightarrow +\infty} \Pr\{|H_{k,(h,l)}| > \delta\} = 0. \tag{5.3}$$

Note that the upper bound in proposition 5.ii provides for the convergence to 0 of the mean square error of the KF estimate of w at the time k , a rate of order $O(\sqrt{1/k})$. As for proposition 5.iv, an intuitive explanation is the following. As the parameter w to be learned does not change in time, after a sufficiently large number of “good” examples, the learning machine has practically learned w , and future examples are practically not needed (of course, this holds in the case, considered so far, in which the parameter w does not change with time; see section 8 for a relaxation of this assumption).

Now, let

$$\text{MSE}_k^\circ := \mathbb{E}_{w, \hat{w}_k^\circ} \{(w - \hat{w}_k^\circ)' (w - \hat{w}_k^\circ)\}$$

denote the mean square error of the OLL estimate at time k . The next proposition provides a recursion to compute and bound from above MSE_k° and states its convergence to 0. We refer to remark 18 in the appendix for a possible way to derive estimates of the associated rate of convergence.

Let

$$e_k^\dagger := \hat{w}_k^\dagger - w,$$

and denote by

$$\Sigma_{e_k^\dagger} := \mathbb{E}_{e_k^\dagger} \left\{ \left(e_k^\dagger - \mathbb{E}\{e_k^\dagger\} \right) \left(e_k^\dagger - \mathbb{E}\{e_k^\dagger\} \right)' \right\} = \mathbb{E}_{e_k^\dagger} \left\{ (e_k^\dagger) (e_k^\dagger)' \right\} \tag{5.4}$$

the (unconditional) covariance matrix of e_k^\dagger . Moreover, we denote by

$$\begin{aligned} \Sigma_{e_k^\circ} &:= \mathbb{E}_{e_k^\circ} \left\{ \left(e_k^\circ - \mathbb{E}\{e_k^\circ\} \right) \left(e_k^\circ - \mathbb{E}\{e_k^\circ\} \right)' \right\} \\ &= \mathbb{E}_{e_k^\circ} \left\{ (e_k^\circ) (e_k^\circ)' \right\} \end{aligned} \tag{5.5}$$

the (unconditional) covariance matrix of e_k° .

Proposition 6 (convergence of the MSE of the OLL estimate). *Let assumptions 1–6 be satisfied. Then the following hold:*

i.

$$MSE_k^\circ = \text{Tr} \{ \Sigma_{e_k^\circ} \},$$

ii.

$$MSE_k^\circ \leq \text{Tr} \left\{ (I + L_{k-1}) \Sigma_{e_{k-1}^\circ} (I + L_{k-1})' \right\} + \text{Tr} \left\{ L_{k-1} \Sigma_{e_{k-1}^\circ} L_{k-1}' \right\} \\ + 2\sqrt{\text{Tr} \left\{ (I + L_{k-1}) \Sigma_{e_{k-1}^\circ} (I + L_{k-1})' \right\} \text{Tr} \left\{ L_{k-1} \Sigma_{e_{k-1}^\circ} L_{k-1}' \right\}}.$$

iii. Under the assumptions made in section 4, one has

$$\lim_{k \rightarrow +\infty} MSE_k^\circ = \lim_{k \rightarrow +\infty} \mathbb{E} \{ (w - \hat{w}_k^\circ)' (w - \hat{w}_k^\circ) \} = 0. \tag{5.6}$$

6 Comparisons with Other Machine Learning Techniques

In this section, we present some connections and comparisons between the solutions to our learning paradigm and machine learning techniques such as average regret minimization (section 6.1), stochastic gradient descent (section 6.2), and Kalman-based estimates (sections 6.3 and 6.4).

6.1 Connections with Average Regret Minimization. As for the finite-horizon case, the minimization of the average learning functional, equation 4.4 is equivalent to the minimization of the alternative learning functional

$$\liminf_{N \rightarrow +\infty} \left(\frac{1}{N} \mathbb{E}_{w, \{x_k\}_{k=0}^{N-1}, \{\varepsilon_k\}_{k=0}^{N-2}} \left\{ \sum_{k=0}^{N-1} [(\hat{w}'_k x_k - y_k)^2 + \gamma u'_k u_k] \right\} \right), \tag{6.1}$$

since equation 6.1 is just equal to

$$J_\gamma^\infty (\{u_k(\tilde{I}_k)\}_{k=0}^\infty) + \sigma_\varepsilon^2 \tag{6.2}$$

(see formula 3.6). We now consider the limit case $\gamma = 0$, denoting by J_0^∞ the corresponding average learning functional. Then, observing that the following equality holds (due to the independence of the disturbance noises ε_k),

$$\sigma_\varepsilon^2 = \liminf_{N \rightarrow +\infty} \left(\frac{1}{N} \mathbb{E}_{w, \{x_k\}_{k=0}^{N-1}, \{\varepsilon_k\}_{k=0}^{N-2}} \left\{ \sum_{k=0}^{N-1} [(w' x_k - y_k)^2] \right\} \right), \tag{6.3}$$

we obtain

$$J_0^\infty (\{u_k(\tilde{I}_k)\}_{k=0}^\infty) = \liminf_{N \rightarrow +\infty} \left(\frac{1}{N} \mathbb{E}_{w, \{x_k\}_{k=0}^{N-1}, \{e_k\}_{k=0}^{N-2}} \left\{ \sum_{k=0}^{N-1} [(\hat{w}'_k x_k - y_k)^2] \right\} \right) \\ - \liminf_{N \rightarrow +\infty} \left(\frac{1}{N} \mathbb{E}_{w, \{x_k\}_{k=0}^{N-1}, \{e_k\}_{k=0}^{N-2}} \left\{ \sum_{k=0}^{N-1} [(w' x_k - y_k)^2] \right\} \right), \quad (6.4)$$

which can be interpreted as an average regret functional (Zinkevich, 2003). Under the assumptions of proposition 4 (with $\gamma = 0$), the sequence of KF estimates minimizes the functional J_0^∞ , hence also the average regret functional, equation 6.4. Moreover, the minimum average regret is 0 because, by proposition 5.ii, one has

$$\lim_{k \rightarrow +\infty} \text{MSE}_k^\dagger = \lim_{k \rightarrow +\infty} \mathbb{E}_{w, \hat{w}_k^\dagger} \left\{ (w - \hat{w}_k^\dagger)' (w - \hat{w}_k^\dagger) \right\} \\ = \lim_{k \rightarrow +\infty} \mathbb{E}_{e_k^\dagger} \left\{ (e_k^\dagger)' (e_k^\dagger) \right\} = 0; \quad (6.5)$$

then one gets also

$$\lim_{k \rightarrow +\infty} \mathbb{E}_{e_k^\dagger, \bar{Q}_k} \left\{ (e_k^\dagger)' Q_k (e_k^\dagger) \right\} = \lim_{k \rightarrow +\infty} \mathbb{E}_{e_k^\dagger} \left\{ (e_k^\dagger)' \bar{Q} (e_k^\dagger) \right\} \\ \leq \lambda_{\max}(\bar{Q}) \lim_{k \rightarrow +\infty} \mathbb{E}_{e_k^\dagger} \left\{ (e_k^\dagger)' (e_k^\dagger) \right\} \\ = 0. \quad (6.6)$$

Nevertheless, the next proposition shows that for any $\gamma > 0$, also the sequence of OLL estimates minimizes the average regret functional 6.4.

Proposition 7. *For any $\gamma > 0$, under the assumptions of proposition 4, the sequence of OLL estimates minimizes the average regret functional, equation 6.4.*

6.2 Connections with KF and Stochastic Gradient Descent. At each time $k + 1$, our OLL estimate \hat{w}_{k+1}° of the parameter vector w associated with the data generation model has the following recursive form (see, e.g., the statement of proposition 2):

$$\hat{w}_{k+1}^\circ = \hat{w}_k^\circ + L_k(\hat{w}_k^\circ - \hat{w}_k^\dagger), \quad (6.7)$$

where L_k is a suitable square matrix and $\hat{w}_k^\dagger := \mathbb{E}\{w|I_k\}$ is the Kalman filter estimate of w at time k , based on the vector I_k that $\overset{w}{}$ collects all the information available to the learning machine up to time k . Hence, it follows from equation 6.7 that our estimates are obtained from the Kalman filter estimates

through an additional smoothing step. The form of equation 6.7 is similar to the ones of other online estimates obtained through various machine learning techniques, such as stochastic gradient descent (Spall, 2003). However, there is a substantial difference: we derive equation 6.7 as the optimal solution of a suitable optimal control/estimation problem, showing various interesting consequences of that in the letter, made possible by our use of optimal control/estimation techniques. We believe that this approach could also be fruitfully applied to other machine learning techniques used in online learning. Moreover, we offer a principled way to construct the matrix L_k in equation 6.7 as the solution of a suitable Riccati equation.

6.3 Outperformance with Respect to KF in Terms of Smoothness. For simplicity, we consider the finite-horizon case. The extension to the infinite-horizon case can be performed by a limit process, as in section 4. The next proposition shows that the OLL estimates are smoother than the KF estimates in the sense that the value of

$$\mathbb{E}_{\{u_k\}_{k=0}^{N-1}} \left\{ \sum_{k=0}^{N-1} u'_k u_k \right\}$$

when $\gamma > 0$ and the updates are generated by equation 3.13 is smaller than or equal to the corresponding value obtained when the updates are generated by equation 3.37 with \simeq replaced by $=$ (i.e., in the limit $\gamma \rightarrow 0$). The limit problem obtained when γ tends to 0 is just the Kalman estimation problem, whose optimal sequence of updates is

$$u_k^\dagger := - \mathbb{E}_{e_k^\dagger} \{ e_k^\dagger | \tilde{I}_k \},$$

where

$$e_k^\dagger := w_k^\dagger - w_k.$$

Proposition 8. *Let assumptions 1–5 be satisfied. Then*

$$\mathbb{E}_{\{u_k^\circ\}_{k=0}^{N-1}} \left\{ \sum_{k=0}^{N-1} [(u_k^\circ)^\dagger (u_k^\circ)] \right\} \leq \mathbb{E}_{\{u_k^\dagger\}_{k=0}^{N-1}} \left\{ \sum_{k=0}^{N-1} [(u_k^\dagger)^\dagger (u_k^\dagger)] \right\}.$$

The simulation results shown in Figure 2, which refers to a setup similar to the one of Figure 1, are in line with the result from proposition 8. The figure suggests that for every $k = 0, \dots, N - 1$, the stronger result,

$$\mathbb{E}_{u_k^\circ} \{ [(u_k^\circ)^\dagger (u_k^\circ)] \} \leq \mathbb{E}_{u_k^\dagger} \{ [(u_k^\dagger)^\dagger (u_k^\dagger)] \}, \tag{6.8}$$

may also hold.

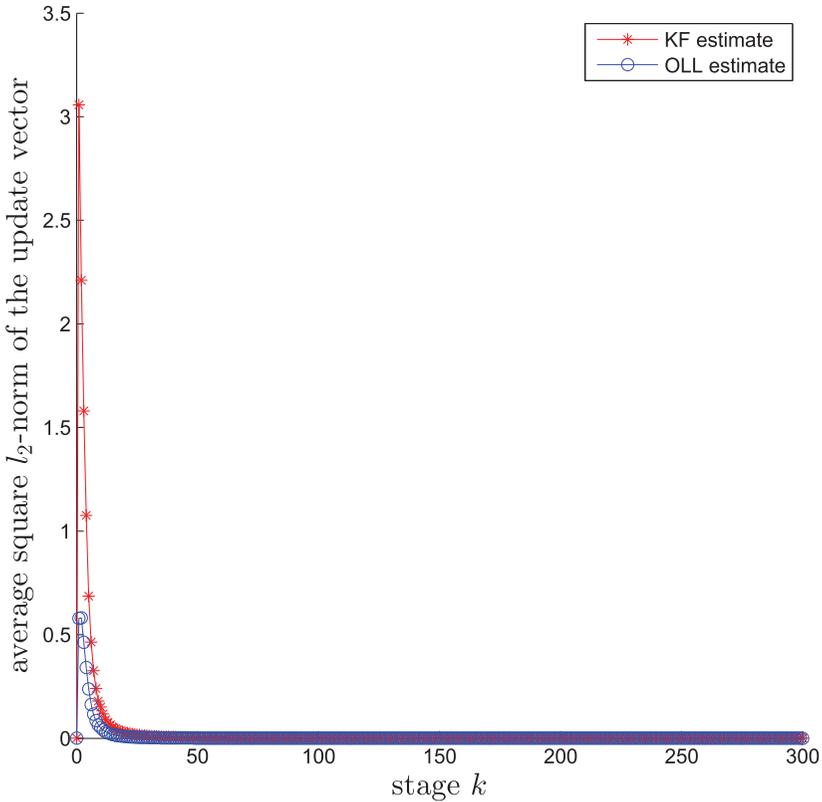


Figure 2: A comparison between the empirical averages of the square l_2 -norms of the vectors of updates used to generate, respectively, the OLL estimates \hat{w}_k° and the KF estimates \hat{w}_k^\dagger . The parameters are the same as in Figure 1, apart from w , which is generated according to a gaussian distribution, with mean $(0, 0, 0)'$ and covariance matrix $\Sigma_w = 4I$. The empirical averages of the square l_2 -norms have been computed by considering 10,000 independent simulations.

Finally, Figures 3 and 4 show that both approaches are also suitable for parameter vectors of much larger dimension. Indeed, it refers to the case of $d = 100$ and $N + 1 = 1001$ online examples. Figure 3 reports the l_2 -norm of the error vector associated with the KF estimate and with the OLL estimate, respectively, at the generic stage k , for a single realization of the sequence of the input-output examples. The running time of such a simulation (whose code was written in Matlab R2013, as for all of the other simulations) was about 29 seconds, on a notebook with a 1.40 GHz CPU and 4 GB of RAM. Figure 4 reports a similar comparison for the l_2 -norm of the update vector, for the same realization. It is worth noting that for a better comparison of the

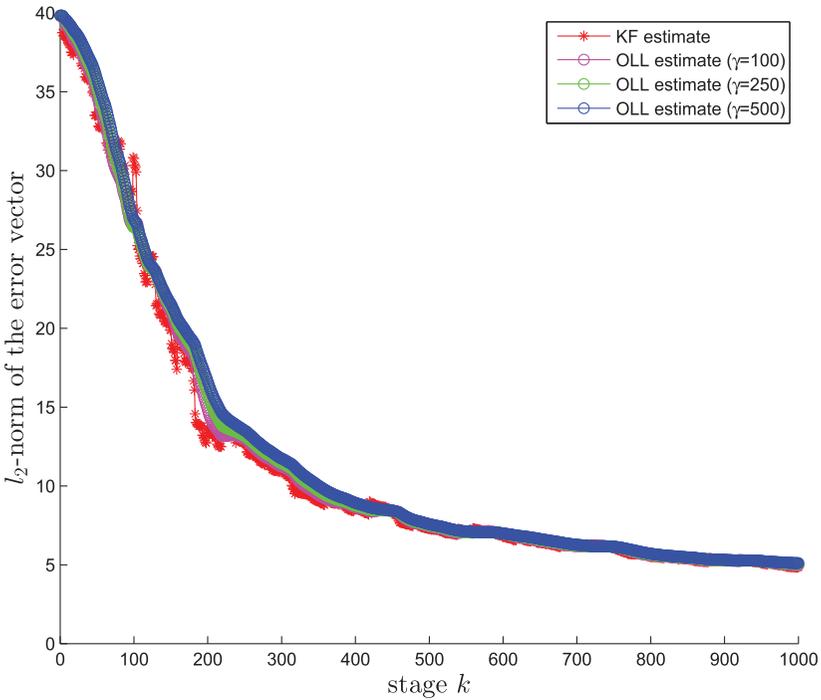


Figure 3: For a setup similar to the one of Figure 2 but with $d = 100$, $N + 1 = 1001$ online examples, a single realization of the sequence of input-output examples and three choices ($\gamma = 100, 250, 500$) for the regularization parameter: l_2 -norm of the error vector associated with the KF estimate and OLL estimate at the generic stage k . The square root of the mean square l_2 -norm of the error vector (averaged with respect to all the stages) is about 14.4134 for the KF estimate and, respectively, 14.6413, 14.8192, and 15.0421 for the OLL estimates with $\gamma = 100, 250, 500$.

two figures, l_2 -norms and not square l_2 -norms have been considered therein (in Figure 2, the square l_2 -norm of the update vector was considered, for its relationship with proposition 8). To compare the two approaches, various choices for the regularization parameter have been considered in Figures 3 and 4 (see the captions for details). In more detail, Figure 3 shows that for this case of a time-invariant parameter vector, the KF estimate and the OLL estimates have a similar error (in a root mean square sense, slightly smaller for the KF estimate, which is, however, much less smooth with respect to the time index k). Moreover, the OLL estimates have much smaller norms of the update vectors (in Figure 4, they are practically negligible compared to the ones associated with the KF estimate, for all the values considered for the regularization parameter). In item e.3 of section 8, we show that for

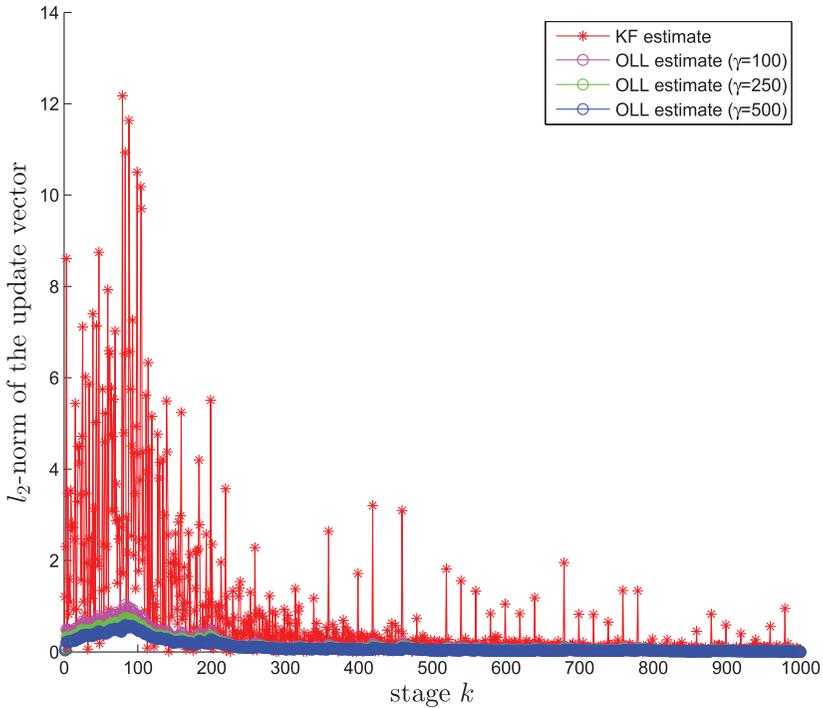


Figure 4: For the same setup and the same realization of the input-output examples as in Figure 3: l_2 -norm of the update vector used to generate, respectively, the KF estimate and the OLL estimate at the generic stage k .

the case of a slowly time-varying parameter vector, the OLL estimate can achieve even a much smaller error than the KF estimate, under a suitable periodic reinitialization of the matrices Σ_k (see Figure 8).

6.4 Outperformance with Respect to KF in Terms of Sensitivity to Outliers. Here we further compare numerically the KF and OLL estimates, now in terms of their different sensitivity to outliers. To this end, we periodically alter the output data perturbation model, choosing the disturbance ε_k to be equal to a positive constant z_1 when k is a multiple of some positive integer Z ; otherwise, it is equal to a negative constant z_2 when k is not a multiple of Z . The two constants z_1 and z_2 are chosen in such a way that the empirical mean (over any time window of duration Z) of the ε_k s is 0 (i.e., the condition $z_1 + (Z - 1)z_2 = 0$ is imposed), and the empirical variance (over the same time-window) is σ_ε^2 (i.e., the condition $\frac{z_1^2 + (Z-1)z_2^2}{Z-1} = \sigma_\varepsilon^2$ is imposed). Hence, $z_1 = \frac{(Z-1)\sigma_\varepsilon}{\sqrt{Z}}$ and $z_2 = -\frac{\sigma_\varepsilon}{\sqrt{Z}}$ are obtained. Moreover, for

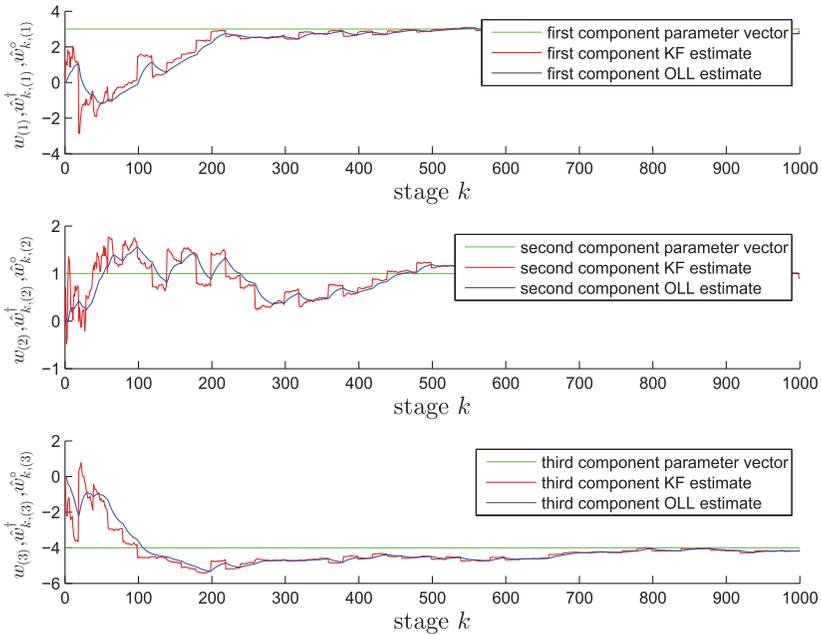


Figure 5: For a setup similar to the one of Figure 1, but choosing $\sigma_\varepsilon = 10, \gamma = 50, N + 1 = 1001$ examples, the diagonal entries of Σ_w equal to 10, and the disturbance ε_k equal to $z_1 = \frac{(Z-1)\sigma_\varepsilon}{\sqrt{Z}}$ when k is a multiple of $Z = 20$, otherwise equal to $z_2 = -\frac{\sigma_\varepsilon}{\sqrt{Z}}$: comparison between the components of the OLL estimate \hat{w}_k^\dagger and of the KF estimate \hat{w}_k° .

a fair comparison with the KF estimate, this modified assumption on the output data-perturbation model is not included in the optimization problem producing the OLL estimates.¹⁴ In other words, that knowledge is not provided to the learning machine.

The numerical results reported in Figure 5 show clearly the much smaller sensitivity to outliers of the OLL estimates with respect to the KF ones (details about the parameter choices are reported in the figure caption). This is ultimately due to the larger smoothness of the OLL estimates with respect to the time index.

An additional theoretical motivation for the smaller sensitivity to outliers of our OLL estimates is obtained by an inspection of formula 3.22

¹⁴Nevertheless, it is worth observing that the sequence of measures generated in this way is still an admissible sequence of measures for the original gaussian disturbance model.

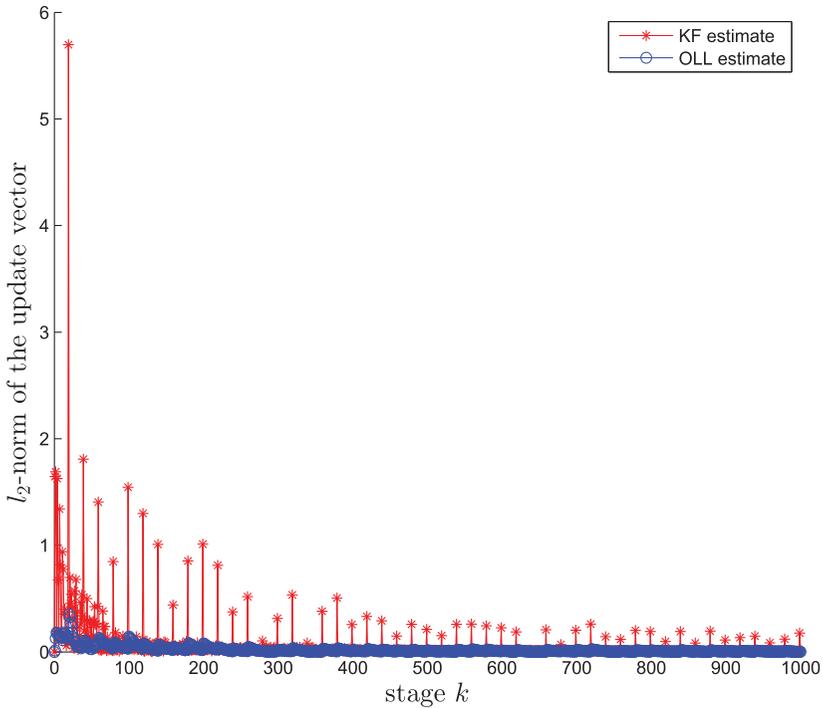


Figure 6: For the example in Figure 5: comparison between the l_2 -norms of the vectors of updates used to generate, respectively, the OLL estimates \hat{w}_k° and the KF estimates \hat{w}_k^\dagger . (Compared to Figure 2, the l_2 -norm is considered here instead of its square for graphical reasons to reduce the difference in range between the two graphs.)

in proposition 2. Limiting for simplicity of the analysis to the first OLL updates of the parameter vector, it follows by that formula and by $\hat{w}_{-1}^\circ = 0$ that $\hat{w}_0^\circ = 0$ and $\hat{w}_1^\circ = -L_0 \hat{w}_0^\dagger$, where \hat{w}_0^\dagger is the first KF update, which is influenced only by the first example presented to the learning machine. Since $|\lambda|_{\max}(L_0) < 1$ (see formula 3.14), it is evident that the OLL estimate is less influenced by the presence of a possible outlier. Moreover, such an influence decreases by increasing the regularization parameter γ , since, by formula 3.14, the larger γ is, the smaller $|\lambda|_{\max}(L_0)$ is. Figure 6 confirms this advantage of the OLL estimates, showing that the l_2 -norms of the vectors of updates used to generate the OLL estimates are typically much smaller than the l_2 -norms of the vectors of updates used to generate the KF estimates. An additional significant advantage of the OLL estimates in the presence of time-varying parameter vectors is detailed in item e of section 8.

7 Nonlinear Models of Data Generation and Application of Kernel Methods

An interesting extension of the model investigated in the previous sections is obtained by mapping the input data x_k preliminarily to another Euclidean space, then applying the model in the new input space. More precisely, one introduces a (possibly nonlinear) mapping $\phi : \mathbb{R}^d \rightarrow E$, where E is a Euclidean space of dimension d_E , possibly larger than d (or even infinite). Then the equation 3.1 becomes

$$y_k = w' \phi(x_k) + \varepsilon_k, \quad (7.1)$$

where the parameter vector w now belongs to E . In this case, one can still apply all the techniques described in the letter taking E as the new input space. Of course, in doing this, the dimensions of some matrices would in general change (typically, they would increase); for instance, in case of a finite d_E , the matrices \bar{K} , L , and Σ_k would become $d_E \times d_E$ matrices, whereas the Kalman gain matrix H_k would become a $1 \times d_E$ matrix. In case of an infinite-dimensional Hilbert space E , they would be replaced by suitable infinite-dimensional linear operators.

Interestingly, as we show below, when doing such an extension, one can apply the so-called kernel trick of kernel machines (Cristianini & Shawe-Taylor, 2000). More precisely, we show some circumstances under which, for every (possibly unseen) input x , one can express both $(\hat{w}_k^+)'\phi(x)$ and $(\hat{w}_k^-)'\phi(x)$ in terms of inner products of the form $\phi(x_j)'\phi(x)$, where x_j is an input example already seen by the learning machine. Hence, if one is able to express $\phi(x_j)'\phi(x)$ in a simple way (e.g., through a symmetric kernel function $\mathcal{K} : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ such that $\phi(x_j)'\phi(x) = \mathcal{K}(x_j, x)$), one can compute $(\hat{w}_k^+)'\phi(x)$ and $(\hat{w}_k^-)'\phi(x)$ even without knowing explicitly the expression of the mapping ϕ .

Remark 11. As an example of a mapping ϕ and its associated kernel \mathcal{K} , we consider the case $d = 2$ and the feature mapping $\phi : \mathbb{R}^2 \rightarrow E = \mathbb{R}^6$, defined as

$$\phi(x) := \left(1, \sqrt{2}x_{(1)}, \sqrt{2}x_{(2)}, \sqrt{2}x_{(1)}x_{(2)}, x_{(1)}^2, x_{(2)}^2 \right)',$$

where $x_{(1)}$ and $x_{(2)}$ are the two components of the vector x . Then, given any two input vectors $x, z \in \mathbb{R}^2$, the inner product $\phi(x)'\phi(z)$ is expressed as

$$\begin{aligned} \phi(x)'\phi(z) &= 1 + 2x_{(1)}z_{(1)} + 2x_{(2)}z_{(2)} + 2x_{(1)}x_{(2)}z_{(1)}z_{(2)} + x_{(1)}^2z_{(1)}^2 + x_{(2)}^2z_{(2)}^2 \\ &= (1 + x'z)^2 \\ &:= \mathcal{K}(x, z), \end{aligned}$$

which is the so-called homogeneous polynomial kernel (Cristianini & Shawe-Taylor, 2000) of order 2, whose evaluation involves only computations to be performed in the original input space \mathbb{R}^2 .

We first consider how to compute $(\hat{w}_k^\dagger)' \phi(x)$ using kernels, then how to compute $(\hat{w}_k^\circ)' \phi(x)$, too. We make the following assumption.

Assumption 7 (covariance matrix of the measurement noise). Let

$$\Sigma_w = \nu I_{d_E}, \tag{7.2}$$

where $\nu > 0$ and I_{d_E} denotes the (matrix associated with the) identity operator on E .

The results presented in the next proposition for the kernel version of the KF estimate are essentially the same as the ones obtained in (Liu et al., 2009, theorems 2 and 3), which shows also how to express the linear combinations inside such equations, through an application of the matrix inversion lemma (see, e.g., Sayed, 2003). However, their extension to the kernel version of the OLL estimate, provided in proposition 10, is novel. In order to improve their readability, in the next formulas—equations 7.3–7.5 and 7.9–7.12—only the functional form of the right-hand side is provided.

Proposition 9. *Let assumptions 1–7 be satisfied for the kernel version of the KF estimate (i.e., with every x_k replaced by $\phi(x_k)$). Then, for every $k = 0, 1, \dots$,*

$$H_k = \Sigma_k (C_k^{(\phi)})' (\sigma_\varepsilon^2)^{-1} = \text{linear combination of } \phi(x_0), \dots, \phi(x_k), \tag{7.3}$$

$$\hat{w}_k^\dagger = \text{linear combination of } \phi(x_0), \dots, \phi(x_k), \tag{7.4}$$

and

$$(\hat{w}_k^\dagger)' \phi(x) = \text{linear combination of } \mathcal{K}(x_0, x), \dots, \mathcal{K}(x_k, x). \tag{7.5}$$

In case of a finite-dimensional space E , the convergence analysis is exactly the same as the one in proposition 5, and a similar (even though more technical) analysis is expected to hold for the infinite-dimensional case. Finally, in case the (matrix associated with the) covariance operator

$$\overline{Q}^{(\phi)} := \mathbb{E}_{\phi(x)} \{ \phi(x) \phi(x)' \} \tag{7.6}$$

is only positive-semidefinite but not positive-definite, one could still follow remark 9 to prove the convergence of the estimate on the subspace on which the input data lie with probability 1. Such a subspace could be estimated by, for example, an application of kernel principal component analysis (KPCA) (Schölkopf, Smola, & Müller, 1998). Moreover, one could even redefine the

problem taking that subspace as the new input space, making the operator $\overline{Q}^{(\phi)}$ be positive-definite when restricted on it.

After dealing with the kernel version of the KF estimate of w , we now investigate the kernel version of its OLL estimate. We make the following assumption:

Assumption 8 (covariance operator). Let one of the following hold:

- i. The covariance operator $\overline{Q}^{(\phi)}$ has the form

$$\overline{Q}^{(\phi)} = qI_{d_E}, \tag{7.7}$$

for some $q > 0$.

- ii.

$$\overline{Q}^{(\phi)} = \overline{Q}_{emp}^{(\phi)} := \frac{1}{l_U} \sum_{j=1}^{l_U} \phi(\tilde{x}_j)\phi(\tilde{x}_j)', \tag{7.8}$$

where l_U is a given positive integer, and $\{\tilde{x}_j, j = 1, \dots, l_U\}$ are some unsupervised examples (assumed here for simplicity to be available to the learning machine starting from the time $k = 0$).

Remark 12. Assumption 8.i refers to a particularly simple model for $\overline{Q}^{(\phi)}$ which is relaxed in assumption 8.ii, which refers to the case in which $\overline{Q}^{(\phi)}$ is modeled by an empirical estimate $\overline{Q}_{emp}^{(\phi)}$ obtained using the unsupervised examples.

Proposition 10. *i. Let assumptions 1–7, and 8.i be satisfied for the kernel version of the KF estimate (i.e., with every x_k replaced by $\phi(x_k)$). Then for every $k = 1, 2, \dots$,*

$$\hat{w}_k^\circ = \text{linear combination of } \phi(x_0), \dots, \phi(x_{k-1}) \tag{7.9}$$

and

$$(\hat{w}_k^\circ)' \phi(x) = \text{linear combination of } \mathcal{K}(x_0, x), \dots, \mathcal{K}(x_{k-1}, x). \tag{7.10}$$

- ii. *If, instead, assumption 8.ii is used, then for every $k = 1, 2, \dots$,*

$$\hat{w}_k^\circ = \text{linear combination of } \phi(\tilde{x}_1), \dots, \phi(\tilde{x}_{l_U}) \tag{7.11}$$

and

$$(\hat{w}_k^\circ)' \phi(x) = \text{linear combination of } \mathcal{K}(\tilde{x}_1, x), \dots, \mathcal{K}(\tilde{x}_{l_U}, x). \tag{7.12}$$

Remark 13. A significant advantage of the representation 7.11 over the representations 7.4 and 7.9 is that the vector \hat{w}_k° has dimension at most l_U .

Remark 14. More generally, $\tilde{x}_1, \dots, \tilde{x}_{l_U}$ in assumption 8.ii could be previously seen input data, preferably not used by the learning machine in combination with labels, to possibly avoid overtraining. So their number could grow as the learning machine acquires examples. Of course, after adding new empirical data in the estimate 7.8, one could also accordingly update the matrix L_k (or, in the infinite-horizon case, the stationary matrix L), in item d.1 of section 8.

We conclude this section by mentioning that in the nonlinear case, differently from techniques such as the extended KF (Grewal & Andrews, 2001), the kernel version of the OLL estimate has the advantage of solving an optimal control (or optimal estimation) problem. Other approaches to online learning with kernels are described, for example, in Diethel and Girolami (2013).

8 Extensions

In this section, we illustrate some other extensions of the proposed OLL estimation scheme investigated in the letter.

8.1 Extension a: Nonzero Mean of x_k . In this case, no significant change in the analysis is required. The only difference is that $\mathbb{E}\{Q_k\}$ and \bar{Q} are now correlation matrices instead of covariance matrices. Q_k

8.2 Extension b: Nonzero Mean of w . Propositions 5 and 6 still hold if $\mathbb{E}_w\{w\} \neq 0$, and the KF estimate and the OLL estimate are initialized, respectively, by

$$\hat{w}_{-1}^\dagger = \mathbb{E}_w\{w\}$$

and

$$\hat{w}_0^\circ = \hat{w}_0 = \mathbb{E}_w\{w\} \tag{8.1}$$

(notice that two different initialization indices have been used for the two estimates, where the subscript -1 has been used to denote the a priori KF estimate—the one obtained before the presentation of the first example, whereas the subscript 0 has been used for initializing of the OLL estimate).¹⁵

¹⁵Recall that \hat{w}_0° refers to the OLL estimate obtained before seeing the first example, whereas \hat{w}_1° refers to the OLL estimate obtained after seeing the first example but before seeing the second one, and so on. Instead, \hat{w}_{-1}^\dagger refers to the KF estimate obtained before seeing the first example, whereas \hat{w}_0^\dagger refers to the KF estimate obtained after seeing the first

Indeed, in such a case one obtains similar expressions as in the appendix for the matrices Σ_k in equation A.48 and for the matrices $\Sigma_{e_k^\dagger}$, $\Sigma_{e_k^\circ}$, $\Sigma_{e_k^\dagger, e_k^\circ}$, $\Sigma_{e_k^\circ, e_k^\dagger}$ in equations 5.4, 5.5, A.65 and A.66, respectively, and the same equation, A.73, which is used there to obtain the convergence result, equation 5.6, through an analysis of the convergence of $\text{Tr}\{\Sigma_{e_k^\circ}\}$ when k tends to $+\infty$.

Remark 15. The case $\mathbb{E}\{w\} \neq 0$ is important in practice, and, among other ones, it models the situation in which, after some number \bar{k} of measures, the time index k is shifted to the left (i.e., k is replaced by $k - \bar{k}$, or equivalently, one reformulates problem OLL using \bar{k} instead of 0 as the initial index in the summation of its objective, equation 3.5) and the knowledge derived by the previous estimates (i.e., the one up to the time $\bar{k} - 1$) is used to generate the term $\mathbb{E}\{w\}$ (this is actually a posteriori knowledge, since it summarizes the knowledge deriving from the previous estimates, but it becomes the new a priori knowledge for the new problem with modified starting index in the summation).

Similar convergence results are obtained even if one replaces equation 8.1 with

$$\hat{w}_0^\circ = \hat{w}_0 \neq \mathbb{E}\{w\}. \quad (8.2)$$

Indeed, in such a case, the convergence analysis of $\text{Tr}\{\Sigma_{e_k^\circ}\}$ made in the appendix is still valid. The only difference is that now one has $\mathbb{E}\{e_k^\circ\} \neq 0$, but in an infinite-horizon setting, $\mathbb{E}\{e_k^\circ\}$ also tends to 0 exponentially fast as k tends to $+\infty$, due to equation A.64 in the appendix with L_j replaced by L , since the matrix $I+L$ has spectral radius smaller than 1.

8.3 Extension c: Introduction of a Bias in the Model. Instead of the measurement equation 3.1, one could consider

$$y_k = w'x_k + \varepsilon_k + b, \quad (8.3)$$

where b is an additional parameter to be learned using the sequence of examples available online. This case can be reduced to equation 3.1 by replacing the input vector x_k by $(x_k', 1)'$, and the parameter vector w to be learned by $(w', 1)'$. As the last component of the new input vector $(x_k', 1)'$ has

example but before seeing the second one, and so on. Hence, according to the notation, there is a shift in the indices of the two estimates, the available information being the same. Of course, a more uniform notation could have been used instead at the expense of shifting and renaming the index for the KF estimate, but using a less common notation for it.

nonzero mean, one is also reduced to case a) above. Moreover, the assumption of positive-definiteness of the correlation matrix of $(x'_k, 1)'$ is satisfied automatically if it is satisfied by the covariance matrix of x_k .

8.4 Extension d: More Complex Models for the Measurement Errors.

The measurement errors ε_k could have nonzero means, nonidentical distributions, and/or be not mutually independent. The first two cases can be dealt with in a straightforward way. Indeed, in the first case, one has only to subtract the expectation of ε_k from the measure y_k before presenting it as an input to the KF,¹⁶ while in the second case, one has to insert an additional index k to σ_ε^2 , using terms of the form $\sigma_{\varepsilon_k}^2$ in the Kalman filter recursion scheme 3.25 and in the Kalman gain matrix 3.24. Finally, in the correlated case, one could model the measurement noise as the output of an auxiliary uncontrolled linear dynamical system, which receives mutually independent noises as inputs. In this case, when the horizon tends to $+\infty$, the convergence of the solution of the ARE to a stationary solution could be more difficult to prove (or such a convergence could even not hold at all), since the reachability condition needed for the application of Bertsekas (1995, proposition 4.1) would be violated.

8.5 Extension e: Time-Varying Models for Some Parameters. When solving the “shifted version” of problem OLL that uses \bar{k} as the initial index (see remark 15), one could exploit, for some of its parameters, time-varying models (which could be also estimated online), including the cases of:

- e.1: Slowly time-varying covariance matrices $\mathbb{E}\{Q_k\}$ of the input examples x_k
- e.2: Slowly time-varying variances $\sigma_{\varepsilon_k}^2$ of the measurement noises ε_k
- e.3: A slowly time-varying parameter vector in the data-generation model, equation 3.1.

About issue e.1, we notice that the covariance matrices $\mathbb{E}\{Q_k\}$ may be estimated online from the already-observed input data¹⁷ x_k ($k = 0, \dots, \bar{k} - 1$), when such covariance matrices do not depend on the time index (as assumed in the basic infinite-horizon version of the problem), but also if one

¹⁶By redefining \tilde{y}_{k+1}° in equation A.19 as

$$\tilde{y}_{k+1}^\circ = C_{k+1} \hat{w}_{k+1}^\circ - (y_{k+1} - \mathbb{E}\{\varepsilon_{k+1}\}).$$

¹⁷In practice, in order to possibly avoid overtraining, one could use a subset of the input data available online up to time \bar{k} to estimate the covariance matrices $\mathbb{E}\{Q_k\}$ and include only the data associated with the other subset in the definition of the learning functional of problem OLL. A similar remark holds for the online estimate of the variance of the measurement noise, which is described in item e.2.

makes the assumption instead that they are slowly time varying (in such a case, one could give more weight in such estimates to the last input data using, for example, a forgetting factor). Now let us suppose that one is solving the infinite-horizon version of problem OLL under the assumptions of section 4, replacing the “common”¹⁸ covariance matrix \bar{Q} of the future input data by its initial estimate, here denoted as $\bar{Q}(0)$. However, let us also suppose that the estimate at time \bar{k} of \bar{Q} derived from the previous input data (denoted in the following as $\bar{Q}(\bar{k})$) is significantly different from $\bar{Q}(0)$ (due, e.g., to slow changes with respect to time of the covariance matrices $\mathbb{E}\{Q_k\}$, or simply due to a possibly bad initial estimate $\bar{Q}(0)$). Then, using the updated estimate $\bar{Q}(\bar{k})$ as the model for the “common” covariance matrix of the future input data, one could also update the stationary matrix L of the proposed optimal controller, replacing \bar{Q} by $\bar{Q}(\bar{k})$ in the average Riccati equation, equation 3.17, and looking for its (new) “stationary” solution \bar{K} , hence deriving the (new) “stationary” matrix L from equation 4.2. Such new solutions will be “stationary” as long as the future estimates of the “common” covariance matrix of the future input data will not change significantly with respect to $\bar{Q}(\bar{k})$.

Concerning issue e.2, one may detect changes in $\sigma_{\varepsilon_k}^2$ from the already observed input/output data (x_k, y_k) ($k = 0, \dots, \bar{k} - 1$), if one assumes that $\sigma_{\varepsilon_k}^2$ is also slowly time varying. In particular, such changes would be easily detected if, inside that set, the learning machine has at its disposal some pairs (x_k, y_k) with similar values for x_k , or even several groups of such pairs, each characterized by similar values of x_k . In this way, one could generate, inside the i th such group,

$$G^{(i)} := \left\{ \left(x_{k_1}^{(i)}, y_{k_1}^{(i)} \right), \dots, \left(x_{k_{|G^{(i)}}}^{(i)}, y_{k_{|G^{(i)}}}^{(i)} \right) \right\},$$

the auxiliary $|G^{(i)}| \times |G^{(i)}|$ data matrix $Y^{\text{aux},(i)}$, whose element $Y_{(h,l)}^{\text{aux},(i)}$ is defined as

$$Y_{(h,l)}^{\text{aux},(i)} := y_{k_h}^{(i)} - y_{k_l}^{(i)},$$

and has a very small dependence on w , since

$$y_{k_h}^{(i)} - y_{k_l}^{(i)} = w' \left(x_{k_h}^{(i)} - x_{k_l}^{(i)} \right) + \varepsilon_{k_h}^{(i)} - \varepsilon_{k_l}^{(i)}$$

¹⁸ Although in this paragraph, such a covariance matrix is modeled as time varying, one could assume that on the basis of the knowledge available at time \bar{k} , the future input data will have a similar model as the present (time-varying) one.

and

$$x_{k_h}^{(i)} - x_{k_l}^{(i)} \simeq 0.$$

Then one could use such matrices $G^{(i)}$ to estimate the variance of the measurement noise, giving more importance or weight to the last among such measures. The obtained estimate would be then used as the common variance σ_ε^2 in the shifted version of problem OLL that uses \bar{k} as the initial index, presented in section 4.

Finally, about issue e.3, we will show that a minor modification of the proposed OLL model can learn even a time-varying parameter vector w . We first focus on the case in which, at some time $\hat{k} \leq \bar{k}$, the parameter vector w changes, then remains fixed. In this case, the mean square error of the estimate of the new parameter vector still converges to 0 when k tends to $+\infty$ (see formula 5.6, together with item b in this section). However, the trace of $\Sigma_{\hat{k}}$ may be extremely small (see equation 5.1, combined with Markov's inequality), making the trace of Σ_k be extremely small, for every $k \geq \hat{k}$. Since such matrices are used to construct the Kalman gain matrices H_k , the convergence to 0 of the error with respect to the new parameter vector may be extremely slow, for both the KF and the OLL estimates. This issue could be solved in both cases by a reinitialization at time \bar{k} of the covariance matrix $\Sigma_{\bar{k}}$ to the a priori covariance matrix Σ_w . More generally, a periodic reinitialization could be used to track a periodically (or continuously) changing parameter vector w . In this case, the OLL estimate would have the advantage, with respect to the KF estimate, to change more slowly in time, making it more suitable to a slowly time-varying parameter vector. The next figures provide more insight into this last issue.

Figures 7 and 8 refer to a parameter vector that changes periodically, the change in its components being small and random. In Figure 7, there is no reinitialization to Σ_w of any of the matrices Σ_k , and the convergence to the new parameter vector of both the KF estimate and the OLL estimate is slow. Figure 8 instead refers to the case in which—the change in the parameter vector being the same—there is a periodic reinitialization of the matrices Σ_k to Σ_w (this second period has been chosen as different from the first one according to which the parameter vector changes, just to avoid giving the learning machine the advantage of knowing when the parameter vector changes, in order to model the more realistic situation in which this knowledge is not available to the machine). In this case, both estimates are able to track the time-varying parameter vector w in a better way, but the OLL estimate is smoother than the KF estimate due to the presence of the regularization parameter $\gamma > 0$. So in this context, the OLL estimate is preferable to the KF estimate if the parameter vector changes slowly with time. Figures 9 and 10 show the reason that this happens: when there is no reinitialization

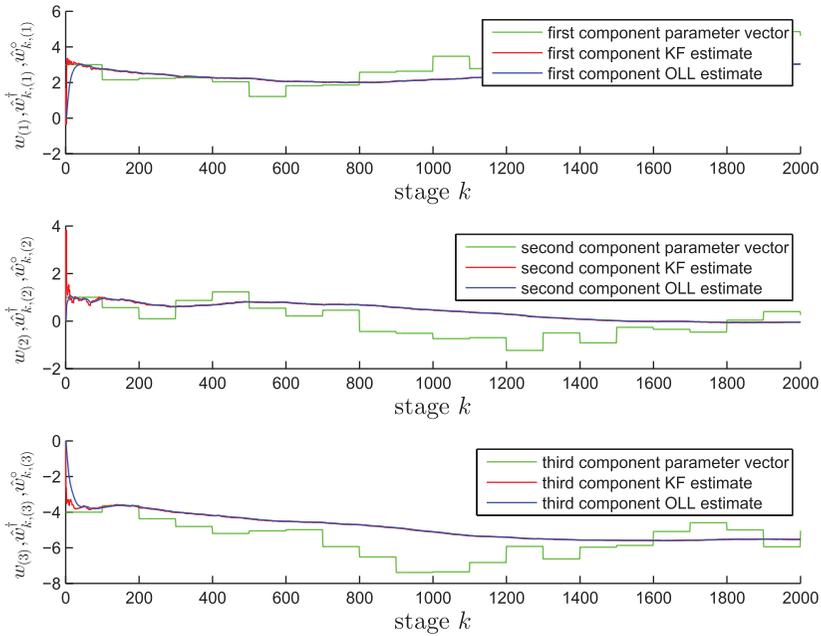


Figure 7: For the case of a time-varying parameter vector, a comparison between the components of the optimal estimate \hat{w}_k° at the time k of the parameter vector w , obtained by solving problem OLL modeling online learning, and the corresponding components of the estimate \hat{w}_k^\dagger at the time k , obtained by applying the Kalman filter. A setting similar to the one of Figure 2 has been considered, but with $N + 1 = 2001$ online examples, and the parameter vector w randomly changed of a small amount every 100 online examples. The covariance matrix Σ_w of the initial w has been chosen to be diagonal with diagonal entries equal to 64, the variance σ_ε^2 of the measurement noise has been chosen to be equal to 1, and the regularization parameter γ has been set to 30. No periodic reinitialization to Σ_w of any of the matrices Σ_k was performed.

to Σ_w of the matrices Σ_k , the Frobenius norm¹⁹ of the Kalman gain matrix H_k is expected to be small for k large (see formulas 3.24 and 5.1), which

¹⁹We recall that for a $n \times m$ real matrix S , its Frobenius norm is defined as

$$\|S\|_{\text{Frobenius}} := \sqrt{\sum_{i=1}^n \sum_{j=1}^m S_{(i,j)}^2}. \tag{8.4}$$

In this problem, as being H_k a column vector, its Frobenius norm coincides with its l_2 -norm.

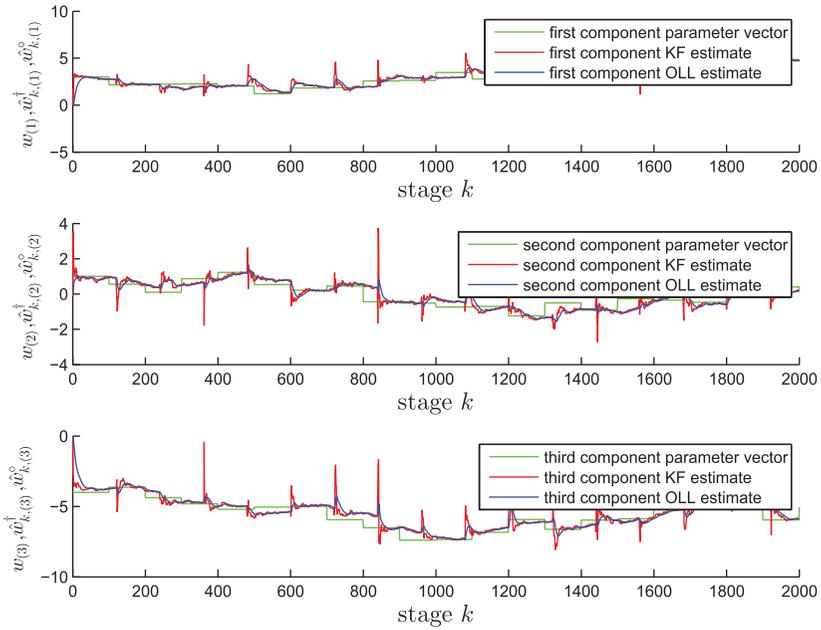


Figure 8: For the case of a time-varying parameter vector: a comparison similar to the one of Figure 7, with the same online examples and the same changes in w , but with a reinitialization to Σ_w of the matrices Σ_k performed every 120 online examples.

is confirmed by Figure 9. Hence, although the norm of the error $y_k - C_k \hat{w}_k^\dagger$ tends to increase when the parameter vector changes, the KF estimate of w at time k is not affected so much by this change (see formula 3.23), hence, also the OLL estimate does not change so much (see formula 3.22). Instead, a reinitialization to Σ_w tends to make the Frobenius norm of the Kalman gain matrix H_k much larger (which is confirmed by Figure 10), and this amplifies the effect of the larger norm of $y_k - C_k \hat{w}_k^\dagger$ (due to the change in the parameter vector w) on the KF estimate of w at time k . For the OLL estimate, the change in the estimate is expected to be smaller due to the smoothing effect of formula 3.22. So, like the KF estimate, the OLL estimate is also able to track the change in the parameter vector but with smoother behavior.

Figures 11 and 12 demonstrate also that a periodic reinitialization to Σ_w of the matrices Σ_k does not negatively affect so much the tracking of a time-invariant parameter vector w in the case of the OLL estimate, whereas the KF estimate is more negatively affected. Again, this is due to the smoothing effect of formula 3.22.

Finally, various approaches to deal with learning time-varying parameters online were presented in McGough (2003); Liu et al. (2009); Chang

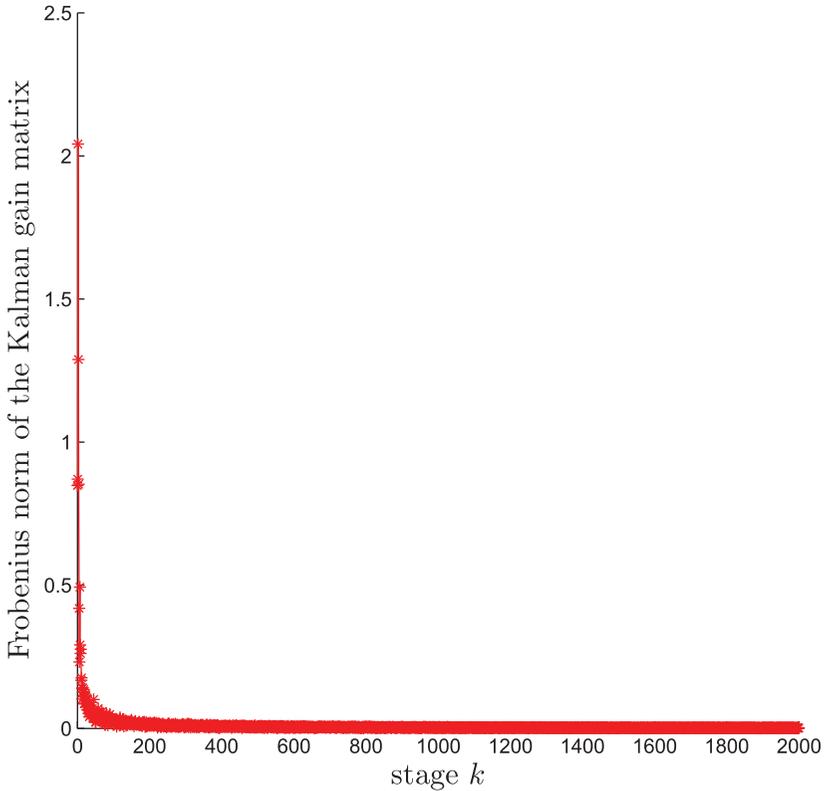


Figure 9: Frobenius norm of the Kalman gain matrix H_k for the experiment shown in Figure 7 (highly dense regions correspond to oscillations in the Frobenius norm with respect to the time index k).

(2014), in some cases in the context of state estimation of dynamical systems in the presence of outliers (Alessandri & Awawdeh, 2016; De Palma & Indiveri, 2016). As a possible extension, one could combine those approaches with the regularization of the updates included in our model, whose beneficial effects have also been demonstrated in this time-varying case.

8.6 Extension f: Insertion of a Discount Factor in the Problem. One may be interested in giving different weights to future expected errors, giving more importance to the present. This can be modeled by inserting a discount factor $\rho \in (0, 1)$, and modifying the learning functional, equation 3.9, as follows:

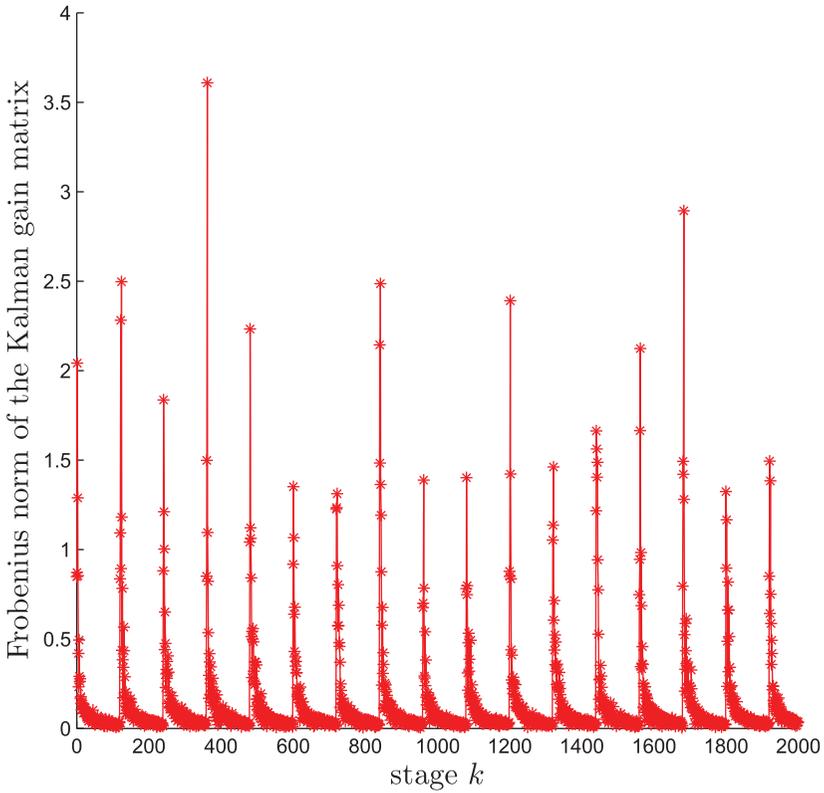


Figure 10: Frobenius norm of the Kalman gain matrix H_k for the experiment shown in Figure 8.

$$\begin{aligned}
 & J_{\gamma, \rho}^N \left(\{u_k(\tilde{I}_k)\}_{k=0}^{N-1} \right) \\
 & := \mathbb{E}_{e_0, \{x_k\}_{k=0}^N, \{\tilde{\epsilon}_k\}_{k=0}^{N-1}} \left\{ \sum_{k=0}^{N-1} \rho^k \left[(e'_k x_k)^2 + \gamma u'_k u_k \right] + \rho^N (e'_N x_N)^2 \right\} \\
 & = \mathbb{E}_{e_0, \{x_k\}_{k=0}^N, \{\tilde{\epsilon}_k\}_{k=0}^{N-1}} \left\{ \sum_{k=0}^{N-1} \rho^k \left[e'_k Q_k e_k + \gamma u'_k u_k \right] + \rho^N e'_N Q_N e_N \right\}. \tag{8.5}
 \end{aligned}$$

Then the resulting problem is just a variation, with random matrices, of the discounted LQ/LQG problem (see Davis & Vinter, 1985, sec. 6.3), for the version with deterministic matrices).

In practice, the modification 8.5 only slightly changes the Bellman equations for the cost-to-go functions, with the introduction of the discount

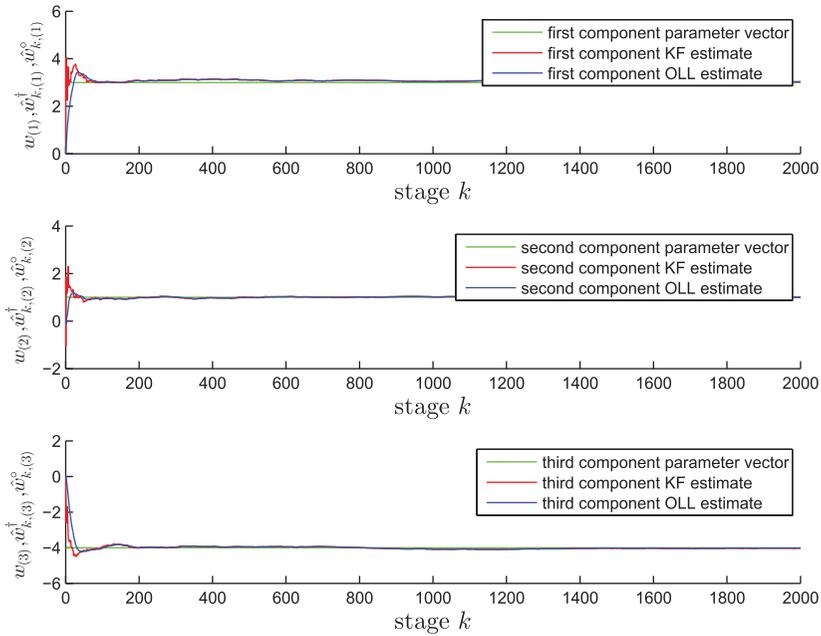


Figure 11: An experimental setting similar to the one of Figure 7 but with no change in the parameter vector.

factor ρ . Particularly, the updates 3.14 to 3.16, and the ARE, equation 3.17, become, respectively,

$$L_k := -(\rho \bar{K}_{k+1} + \gamma I)^{-1} \rho \bar{K}_{k+1}, \tag{8.6}$$

$$K_k := \rho \bar{K}_{k+1} - \rho \bar{K}_{k+1} (\rho \bar{K}_{k+1} + \gamma I)^{-1} \rho \bar{K}_{k+1} + Q_k, \tag{8.7}$$

$$F_k := \rho \bar{K}_{k+1} (\rho \bar{K}_{k+1} + \gamma I)^{-1} \rho \bar{K}_{k+1}, \tag{8.8}$$

and

$$\bar{K}_k := \mathbb{E}_{K_k} \{K_k\} = \rho \bar{K}_{k+1} - \rho \bar{K}_{k+1} (\rho \bar{K}_{k+1} + \gamma I)^{-1} \rho \bar{K}_{k+1} + \mathbb{E}_{Q_k} \{Q_k\}, \tag{8.9}$$

whereas the SRE, equation 3.25, is not changed at all. For the infinite-horizon case, the learning functional is the limit of equation 8.5 for $N \rightarrow +\infty$ (so, no “average” learning functional similar to equation 4.4 is needed). Moreover, the stationary matrix

$$L := -(\rho \bar{K} + \gamma I)^{-1} \rho \bar{K} \tag{8.10}$$

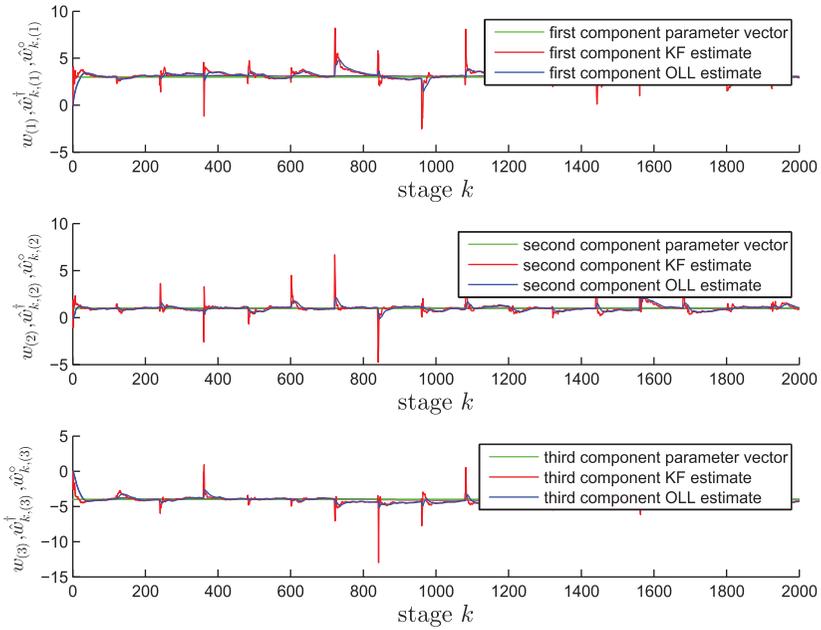


Figure 12: A comparison similar to the one of Figure 11 with the same online examples and with a reinitialization to Σ_w of the matrices Σ_k performed every 120 online examples.

associated with the stationary solution \bar{K} of the stationary version

$$\bar{K} = \rho\bar{K} - \rho\bar{K}(\rho\bar{K} + \gamma I)^{-1}\rho\bar{K} + \bar{Q} \tag{8.11}$$

of the ARE, equation 8.9, is symmetric and such that $(I + L)$ has all its eigenvalues inside the unit circle, as it follows from the positive-definiteness of \bar{K} , which is derived directly from equation 8.11, assuming that \bar{Q} is positive-definite, as in section 4. Concluding, these analogies with the undiscounted case would allow one, for example, to extend the properties stated in proposition 6 to the discounted case.

Regarding the modification 8.5, one can also observe that, like the basic version of problem OLL, the past updates of the estimate of the parameter vector w are not modified when a new example arrives and that, in any case, such updates have no influence on the cost-to-go functions at later stages. Moreover, the discount factor should be not confused with a forgetting

factor²⁰ (which gives less importance to the past errors, to determine the current estimate).

8.7 Extension g: Introduction of Additional Regularizations on the Estimates of w . Given a sequence of additional regularization parameters $\gamma_{\hat{w}_k} > 0$, one could insert terms of the form $\gamma_{\hat{w}_k} \hat{w}_k' \hat{w}_k$ in the per stage cost of the learning functional, equation 3.5, hence replacing it by

$$\begin{aligned}
 & J_{\gamma, \gamma_w}^N (\{u_k(I_k)\}_{k=0}^{N-1}) \\
 & := \mathbb{E}_{e_0, \{x_k\}_{k=0}^N, \{e_k\}_{k=0}^{N-1}} \left\{ \sum_{k=0}^{N-1} \left[((\hat{w}_k - w_k)' x_k)^2 + \gamma_{\hat{w}_k} \hat{w}_k' \hat{w}_k + \gamma u_k' u_k \right] \right. \\
 & \quad \left. + ((\hat{w}_N - w_N)' x_N)^2 \right\} \\
 & = \mathbb{E}_{e_0, \{x_k\}_{k=0}^N, \{e_k\}_{k=0}^{N-1}} \left\{ \sum_{k=0}^{N-1} \left[(\hat{w}_k - w_k)' Q_k (\hat{w}_k - w_k) + \gamma_{\hat{w}_k} \hat{w}_k' \hat{w}_k + \gamma u_k' u_k \right] \right. \\
 & \quad \left. + (\hat{w}_N - w_N)' Q_N (\hat{w}_N - w_N) \right\}. \tag{8.12}
 \end{aligned}$$

In this case, one would still obtain an ARE, but referred to the original dynamical system 3.3 instead of the reduced one, equation 3.7, due to the presence of the terms $\gamma_{\hat{w}_k} \hat{w}_k' \hat{w}_k$ in the objective functional, extension 8.12. So the matrix L_k would have in this case the size $2d \times 2d$. Instead, the SRE would be exactly the same as equation 3.25. A suitable choice for the sequence of the regularization parameters $\gamma_{\hat{w}_k}$ could be a sequence that decreases monotonically to 0. In this case, in the first stages, when the machine has seen a small number of examples, it would be preferable to make the a priori knowledge on w dominate the one associated with the small number of examples seen so far, whereas in the successive stages—having seen a much larger number of examples—it would be preferable to make the a-posteriori knowledge coming from the examples be predominant. This would also be in accordance with common choices of the regularization parameter (as a function of the number of available examples, plus other parameters) for batch learning with regularization (Cucker & Smale, 2002). For the infinite-horizon version of the problem, a constant γ_{w_k} (denoted by γ_w) could be used.

Another extension has to do with the insertion in the per stage cost of terms related to the previous updates. For instance, limiting for simplicity

²⁰ Also the case $\rho > 1$, not studied here because the resulting problem could be not well defined in the infinite-horizon case, does not correspond to a forgetting factor.

to the last previous update and adding another regularization parameter $\gamma_{-1} > 0$, one could extend the definition of the learning functional 8.12 to

$$\begin{aligned}
 & J_{\gamma, \gamma_{\hat{w}}, \gamma_{-1}}^N (\{u_k(I_k)\}_{k=0}^{N-1}) \\
 := & \mathbb{E}_{e_0, \{x_k\}_{k=0}^N, \{\varepsilon_k\}_{k=0}^{N-1}} \left\{ \sum_{k=0}^{N-1} [((\hat{w}_k - w_k)' x_k)^2 + \gamma_{\hat{w}_k} \hat{w}'_k \hat{w}_k \right. \\
 & \left. + \gamma_{-1} (u_k - (\hat{w}_k - \hat{w}_{k-1}))' (u_k - (\hat{w}_k - \hat{w}_{k-1})) + \gamma u'_k u_k \right] \\
 & \left. + ((\hat{w}_N - w_N)' x_N)^2 \right\} \\
 = & \mathbb{E}_{e_0, \{x_k\}_{k=0}^N, \{\varepsilon_k\}_{k=0}^{N-1}} \left\{ \sum_{k=0}^{N-1} [(\hat{w}_k - w_k)' Q_k (\hat{w}_k - w_k) + \gamma_{\hat{w}_k} \hat{w}'_k \hat{w}_k \right. \\
 & \left. + \gamma_{-1} (u_k - (\hat{w}_k - \hat{w}_{k-1}))' (u_k - (\hat{w}_k - \hat{w}_{k-1})) + \gamma u'_k u_k \right] \\
 & \left. + (\hat{w}_N - w_N)' Q_N (\hat{w}_N - w_N) \right\}, \tag{8.13}
 \end{aligned}$$

with $\hat{w}_{-1} := 0$. Here, each component $(u_k - (\hat{w}_k - \hat{w}_{k-1}))_{(j)}$ of the vector $u_k - (\hat{w}_k - \hat{w}_{k-1})$ can be interpreted as a central difference approximation (with discretization step $\Delta_k = 1$) of the second derivative of a continuous-time function (or, more precisely, of the second derivative of a stochastic process (Papoulis, 1991, appendix 10A) $u_{(j)}(t)$ such that $u_{(j)}(k) = u_{k,(j)}$). Indeed, one has

$$\frac{d}{dt} u_{(j)}(t) \Big|_{t=k} \simeq \frac{(\hat{w}_{k+1} - \hat{w}_k)_{(j)}}{\Delta_k} - \frac{(\hat{w}_k - \hat{w}_{k-1})_{(j)}}{\Delta_k} = (u_k - (\hat{w}_k - \hat{w}_{k-1}))_{(j)}.$$

In order to optimize the learning functional 8.13, one could at first extend the definition of the state vector of the dynamical system 3.3, including the previous estimate \hat{w}_{k-1} in the extended state at the time k . Then, tools similar to the ones used in this work could be still used for the optimization.

One can see that by choosing a suitable value for γ_{-1} , one could give different importance to the previous estimates of w in order to generate the optimal update u_k° at each time k . Indeed, when $\gamma_{-1} \simeq 0$, one would expect the estimate \hat{w}_{k-1} not to be practically taken into account to generate u_k° , whereas for γ_{-1} extremely large, one would expect the learning machine to penalize, for each j th component, a change in the “slope” of $u_{k,(j)}$ more than its absolute value $|u_{k,(j)}|$. In other words, to generate $u_{k,(j)}^\circ$, the difference $(\hat{w}_k - \hat{w}_{k-1})_{(j)}$ would be taken into account more than $|u_{k,(j)}|$. For intermediate values of γ_{-1} , both $(\hat{w}_k - \hat{w}_{k-1})_{(j)}$ and $|u_{k,(j)}|$ would be expected to be taken into account significantly to generate $u_{k,(j)}^\circ$. A more rigorous analysis

of these cases (outside the scope of this work) could be done by solving the resulting ARE for the specific problem.

Finally, a similar technique could be used to include terms approximating derivatives of a larger order of $u_{(j)}(t)$, extending the definition of the state vector of the dynamical system 3.3 in such a way as to include all the previous estimates of w that are used to approximate such derivatives.

8.8 Extension h: Extension of the Problem Formulation to the Continuous-Time Case. As already mentioned in section 8.7, some terms in the learning functional, equation 8.12, can be interpreted as approximations of terms arising in a continuous-time formulation of the problem. Also the basic version of problem OLL can be considered as a discrete-time version of a continuous-time problem whose learning functional is

$$\begin{aligned}
 & J_{\gamma,c}^T(u(\cdot, \cdot)) \\
 & := \mathbb{E}_{e_0, \{x(\cdot)\}, \{\varepsilon(\cdot)\}_{k=0}^{N-1}} \left\{ \int_0^T [(\hat{w}(t) - w(t))'x(t)]^2 + \gamma(u(t, I_t))'(u(t, I_t))]dt \right. \\
 & \qquad \qquad \qquad \left. + ((\hat{w}(T) - w(T))'x(T))^2 \right\} \\
 & = \mathbb{E}_{e_0, \{x(\cdot)\}, \{\varepsilon(\cdot)\}_{k=0}^{N-1}} \left\{ \int_0^T [(\hat{w}(t) - w(t))'Q(t)(\hat{w}(t) - w(t)) \right. \\
 & \qquad \qquad \qquad + \gamma(u(t, I_t))'(u(t, I_t))]dt \\
 & \qquad \qquad \qquad \left. + (\hat{w}(T) - w(T))'Q(T)(\hat{w}(T) - w(T)) \right\}, \tag{8.14}
 \end{aligned}$$

with $T = N/\Delta_k$ and²¹ $\Delta_k = 1$, where u, w, \hat{w}, x, Q are suitable stochastic processes, and the system dynamics are described by the stochastic differential equation

$$\begin{cases} dw = 0, \\ d\hat{w} = u dt. \end{cases} \tag{8.15}$$

Moreover, the measurement process would be modeled by the stochastic process

$$y = Cw + \varepsilon, \tag{8.16}$$

²¹ A discretization of equation 8.14, better than equation 3.5, would be obtained using a smaller value of Δ_k , but it would have exactly the same form as equation 3.5.

where ε is another stochastic process (white noise). Finally, the causality constraint that the current update depends only on the history of the measurement and decision processes up to the current time would be enforced, imposing that the stochastic process u is nonanticipative. As it would be rather technical, a rigorous analysis of this case is outside the scope of the work, but it could be done using classic tools, such as the ones used to solve the LQG optimal control problem in continuous time (Maybeck, 1982).

8.9 Extension i: Introduction of Constraints on the Updates. As a possible extension, one could insert constraints on the variable u_k , such as

$$\|u_k\|_2 \leq B_k$$

for some $B_k > 0$ (where $\|\cdot\|_2$ denotes the l_2 -norm) or constraints of the form

$$|u_{k,(j)}| \leq B_j$$

for some $B_j > 0$, $j = 1, \dots, d$. In particular, the latter are linear constraints, as they can be written as the union of the constraints

$$u_{k,(j)} \leq B_j$$

and

$$-u_{k,(j)} \leq B_j$$

(we refer to Gnecco, Gori, & Sanguineti, 2013, and Gnecco, Gori, Melacci, & Sanguineti, 2015a, 2015b, for other examples of constraints in machine learning problems, although presented in a batch framework there). From a theoretical point of view, one could still search for the optimal solution of the resulting constrained optimization problem by solving Bellman equations, provided that one is able to determine the conditional probability distribution of e_k given \tilde{I}_k . Under assumption 5, such a conditional probability distribution is still gaussian, so one needs to know only the conditional mean $\mathbb{E}_{e_k}\{e_k|\tilde{I}_k\}$ and the conditional covariance matrix $\Sigma_k = \mathbb{E}_{e_k}\{(e_k - \mathbb{E}_{e_k}\{e_k|\tilde{I}_k\})(e_k - \mathbb{E}_{e_k}\{e_k|\tilde{I}_k\})'\}$ (which are provided by the Kalman filter recursion scheme) to determine such a conditional probability distribution completely. However, in this case, solving Bellman equations would be not reduced to solving suitable AREs. In practice, being able to solve the problem optimally may be the exception, particularly in the case of a large finite horizon (or of an infinite horizon), since the complexity of the structure of the cost-to-go functions (e.g., the number of "pieces" in case of their possible piecewise quadraticity) may grow when performing the backward

phase of dynamic programming. In practice, one may be forced to give up searching for an optimal solution, and look for good suboptimal solutions instead.

8.10 Extension j: Modification of the Per Stage Cost. Besides the changes already discussed in section 8.7, one could also modify the per stage cost by inserting additive nonquadratic but still convex terms, for example, a LASSO (least absolute shrinkage and selection operator) term of the form $\gamma_L \|u_k\|_1$, where $\gamma_L > 0$ and $\|\cdot\|_1$ denotes the l_1 -norm. The goal of the LASSO is to enforce sparsity of the update u_k at optimality, which is possible due to geometrical properties of the l_1 -norm (Tibshirani, 1996). Here, we observe that as in the case discussed in section 8.9, at least from a theoretical point of view it is still possible to solve the problem by applying dynamic programming if one is able to express the conditional probability distribution of e_k given \tilde{I}_k . This is the case because under assumption 5, such a conditional probability distribution is still gaussian, and can be computed efficiently through the Kalman filter recursion scheme.

8.11 Extension k: Introduction of Constraints on the Parameter Vector w to be Learned. The parameter vector w could be subject to other constraints (e.g., the nonnegativity constraint $w \geq 0$, to be interpreted componentwise) modeling an additional prior knowledge on it. Again, dynamic programming can be still applied in such a variation of problem OLL if one is able to express the conditional probability distribution of e_k given both \tilde{I}_k and the constraint on w . In practice, this is the case since, under assumption 5, such a distribution is obtained from the gaussian conditional probability distribution of e_k given \tilde{I}_k only, by imposing the constraint $w \geq 0$, then doing a successive renormalization of that gaussian conditional probability distribution.

8.12 Extension l: Application of Moving Horizon Techniques. In practice, in certain situations, it could be very difficult to solve exactly the modified problems discussed in sections 8.9 to 8.11. In such cases, one could resort to variations of such problems, obtained following a moving horizon approach (Camacho & Bordons Alba, 2004), which uses a sliding optimization window of constant width and the current estimate as the initial estimate at the left extreme of the current optimization window. Such an approach would assign no importance to far-in-the-future expected errors outside the current optimization window. Moreover, as is typical of moving horizon approaches, it may allow one to find possibly good and stable suboptimal solutions to the original infinite-horizon problems; such solutions could be computed by solving, possibly in real time, simpler (and still convex) optimization problems, especially if the width of the optimization window is small.

8.13 Extension m: Reducing the Complexity of the Estimate and DOWDATING. A similar variation as the one discussed in section 8.12, which also uses a sliding optimization window—but with a reinitialization at 0 (or at a fixed vector) of the estimate of the parameter vector w at the left extreme of each window—has to do with the case in which one wants to completely forget the old examples, making the current estimate depend on only the examples contained in the current window. For a small width of the window, this would have the advantage of limiting the complexity of the current estimate of the parameter vector w .²² Moreover, when shifting the sliding window one unit to the right (hence, inserting a new example and removing the oldest one), a recursive approach could be used to generate the optimal solution of the resulting optimization problem, starting from the one of the previous problem. This approach, called *dowdating* in the literature (as opposed to *updating*), works, for example, for recursive least squares (see, e.g., Vu, 2007) and may be extended to the present problem.

8.14 Extension n: Active Online Learning. One could give to the learning machine the capability of influencing (at least partially) the choice of the sequence of input data. For instance, the learning machine could try to generate examples similar to the ones already seen, possibly making it easier to estimate the statistics of the measurement noises ε_k , or it could focus—at least in an initial learning stage—on certain components of the parameter vector w (e.g., components that seem to be easier to be learned), generating examples with 0 values for all of the other components. Focusing on such components in the initial stages could improve the convergence of the estimate to the true parameter vector w , making the machine learn the more difficult components of w in a second phase. However, in doing this, one should give to the learning machine not an excessive freedom to generate its input examples (e.g., giving it enough freedom only at certain time instants) in order to explore the state space enough, and to focus not only on tasks that appear easy to be learned. Finally, some of the input examples could be chosen deterministically and even presented periodically to the machine, with small changes in the analysis.

²² As shown in section 7, when the initial estimate is 0, for every $k = 1, 2, \dots$, the vector \hat{w}_k° belongs to a finite-dimensional subspace of dimension at most k . So considering only the examples contained in the current sliding window would make each estimate belong to a finite-dimensional subspace of dimension at most equal to the width of the window. However, it also has to be taken into account that in the basic version of problem OLL (the one without the nonlinear mapping ϕ), the maximal dimension of the subspace is equal to dimension d of the input space, whereas in its kernel version, it is equal to d_E . So, this sliding-window approach is expected to reduce significantly the complexity of the estimate only when the size of the window is small compared to d in the basic version of the problem and to d_E in its kernel version. Similar comments hold in the case of a different initialization of the estimate.

8.15 Extension o: Extension of the Problem Formulation through Techniques from Robust Estimation and Control. As a further extension, we discuss briefly an adversarial framework (e.g., both the input examples and the output disturbance noise—the latter not modeled anymore as gaussian random vectors—could be chosen in an adversarial way), in which one still wants to have the ability to learn the parameter vector even in a (suitably defined) worst-case setting. To study this possible extension, techniques from robust estimation or control could be used, particularly, the ones from \mathcal{H}_∞ -filtering or control (Başar & Bernhard, 2008), which, incidentally, is also based on suitable Riccati equations. Still, a direct application of such methods to online machine learning would be not trivial (e.g., in case not only the disturbance noise but also the input examples were chosen in an adversarial way). Compared with the LQ/LQG online learning framework investigated in this letter, a possible extension to online robust estimation or control could have the advantage of a further decrease in the sensitivity to outliers, since the worst case would be considered explicitly to generate estimates of the parameter vector. However, different from our setting, closed-form optimal solutions may not be available for such an extension.

8.16 Extension p: Possible Connections with Neuroscience. Finally, the proposed formulation of online learning from supervised examples with regularization of the updates also has a biological interpretation. Indeed, as evidenced by the analysis of cognitive behavior in children by the psychologist Jean Piaget (1961), human learning proceeds gradually, according to successive developmental stages. In other words, particularly in the earlier phases of life, children are not able to learn all the information presented as an input from the external world but just learn in a gradual way. It has recently been observed that this way of learning, successful in nature, may represent a key to inspire efficient machine learning algorithms (Gori, 2009). Piaget's theory has been related to current studies in neuroscience, such as those in structural and functional brain imaging (Crone & Ridderinkhof, 2011).

9 Discussion

We have proposed and investigated an optimal control approach to online learning from supervised examples, modeled as the online estimation of an unknown parameter vector relating the input examples x_k with their outputs y_k . We have shown the connections of the proposed problem with the classical LQ and LQG optimal control problems, of which the former is a nontrivial variation, as it involves random matrices. We have also compared the optimal solution with the KF estimate, showing cases in which the latter has advantages to it (e.g., more smoothness and less sensitivity to outliers). We have also described, and in some cases developed in detail, some extensions of the basic model, including these:

- a. The infinite-horizon case, with convergence results (in particular, convergence to 0 of the mean square estimation error of the OLL estimate, when the time index goes to infinity)
- b. Nonlinear models, exploiting kernel methods
- c. Nonzero-mean random variables
- d. More complex models for measurement errors
- e. Online estimates of some covariance matrices
- f. A slowly time-varying parameter vector to be learned from the sequence of supervised examples
- g. Discounted problems
- h. Higher-order regularizations of the estimates of w
- i. Continuous time
- j. Active online learning

Appendix: Proofs

The proof of proposition 1 is based on the following lemma:

Lemma 1. For $k = N - 2, \dots, 0$ one has

$$\begin{aligned} & \mathbb{E}_{e_{k+1}, \tilde{I}_{k+1}} \left\{ \left(e_{k+1} - \mathbb{E}_{e_{k+1}} \{ e_{k+1} | \tilde{I}_{k+1} \} \right)' F_{k+1} \left(e_{k+1} - \mathbb{E}_{e_{k+1}} \{ e_{k+1} | \tilde{I}_{k+1} \} \right) \middle| \tilde{I}_k, u_k \right\} \\ &= \text{Tr} \left\{ F_{k+1} \mathbb{E}_{C_{k+1}} \left\{ \Sigma_k - \Sigma_k C'_{k+1} (C_{k+1} \Sigma_k C'_{k+1} + \sigma_\varepsilon^2)^{-1} C_{k+1} \Sigma_k \middle| C_0, \dots, C_k \right\} \right\}. \end{aligned}$$

Proof. By the law of iterated expectations (Bierens, 2005), we get

$$\begin{aligned} & \mathbb{E}_{e_{k+1}, \tilde{I}_{k+1}} \left\{ \left(e_{k+1} - \mathbb{E}_{e_{k+1}} \{ e_{k+1} | \tilde{I}_{k+1} \} \right)' F_{k+1} \left(e_{k+1} - \mathbb{E}_{e_{k+1}} \{ e_{k+1} | \tilde{I}_{k+1} \} \right) \middle| \tilde{I}_k, u_k \right\} \\ &= \mathbb{E}_{\tilde{I}_{k+1}} \left\{ \mathbb{E}_{e_{k+1}} \left\{ \left(e_{k+1} - \mathbb{E}_{e_{k+1}} \{ e_{k+1} | \tilde{I}_{k+1} \} \right)' \right. \right. \\ & \quad \left. \left. \times F_{k+1} \left(e_{k+1} - \mathbb{E}_{e_{k+1}} \{ e_{k+1} | \tilde{I}_{k+1} \} \right) \middle| \tilde{I}_{k+1} \right\} \middle| \tilde{I}_k, u_k \right\}. \end{aligned} \tag{A.1}$$

Now, by properties of the trace and the linearity of the trace and of the expectation operator, one has

$$\begin{aligned} & \mathbb{E}_{e_{k+1}} \left\{ \left(e_{k+1} - \mathbb{E}_{e_{k+1}} \{ e_{k+1} | \tilde{I}_{k+1} \} \right)' F_{k+1} \left(e_{k+1} - \mathbb{E}_{e_{k+1}} \{ e_{k+1} | \tilde{I}_{k+1} \} \right) \middle| \tilde{I}_{k+1} \right\} \\ &= \mathbb{E}_{e_{k+1}} \left\{ \text{Tr} \left\{ \left(e_{k+1} - \mathbb{E}_{e_{k+1}} \{ e_{k+1} | \tilde{I}_{k+1} \} \right)' F_{k+1} \left(e_{k+1} - \mathbb{E}_{e_{k+1}} \{ e_{k+1} | \tilde{I}_{k+1} \} \right) \right\} \middle| \tilde{I}_{k+1} \right\} \end{aligned}$$

$$\begin{aligned}
 &= \mathbb{E}_{e_{k+1}} \left\{ \text{Tr} \left\{ F_{k+1} \left(e_{k+1} - \mathbb{E}_{e_{k+1}} \{ e_{k+1} | \tilde{I}_{k+1} \} \right) \left(e_{k+1} - \mathbb{E}_{e_{k+1}} \{ e_{k+1} | \tilde{I}_{k+1} \} \right)' \right\} \middle| \tilde{I}_{k+1} \right\} \\
 &= \text{Tr} \left\{ F_{k+1} \mathbb{E}_{e_{k+1}} \left\{ \left(e_{k+1} - \mathbb{E}_{e_{k+1}} \{ e_{k+1} | \tilde{I}_{k+1} \} \right) \left(e_{k+1} - \mathbb{E}_{e_{k+1}} \{ e_{k+1} | \tilde{I}_{k+1} \} \right)' \right\} \middle| \tilde{I}_{k+1} \right\} \\
 &= \text{Tr} \{ F_{k+1} \Sigma_{k+1} \}. \tag{A.2}
 \end{aligned}$$

Due to equations 3.15 to 3.17 and their initializations, 3.18 and 3.19, the last expression in equation A.2 is a function of the form $\tilde{f}_{k+1}(\{C_j\}_{j=0}^{k+1})$. Finally, by combining equations 3.25, A.1, and A.2, one gets

$$\begin{aligned}
 &\mathbb{E}_{\tilde{I}_{k+1}} \left\{ \mathbb{E}_{e_{k+1}} \left\{ \left(e_{k+1} - \mathbb{E}_{e_{k+1}} \{ e_{k+1} | \tilde{I}_{k+1} \} \right)' \right. \right. \\
 &\quad \times F_{k+1} \left(e_{k+1} - \mathbb{E}_{e_{k+1}} \{ e_{k+1} | \tilde{I}_{k+1} \} \right) \left. \middle| \tilde{I}_{k+1} \right\} \left. \middle| \tilde{I}_k, u_k \right\} \\
 &= \mathbb{E}_{\Sigma_{k+1}} \left\{ \text{Tr} \{ F_{k+1} \Sigma_{k+1} \} \middle| \tilde{I}_k, u_k \right\} \\
 &= \text{Tr} \left\{ F_{k+1} \mathbb{E}_{C_{k+1}} \left\{ \Sigma_k - \Sigma_k C'_{k+1} (C_{k+1} \Sigma_k C'_{k+1} + \sigma_\varepsilon^2)^{-1} C_{k+1} \Sigma_k \middle| C_0, \dots, C_k \right\} \right\}. \tag{A.3}
 \end{aligned}$$

□

According to lemma 1, the terms

$$\mathbb{E}_{e_{k+1}, \tilde{I}_{k+1}} \left\{ \left(e_{k+1} - \mathbb{E}_{e_{k+1}} \{ e_{k+1} | \tilde{I}_{k+1} \} \right)' F_{k+1} \left(e_{k+1} - \mathbb{E}_{e_{k+1}} \{ e_{k+1} | \tilde{I}_{k+1} \} \right) \middle| \tilde{I}_k, u_k \right\} \tag{A.4}$$

are functions of the form $f_k(\{C_j\}_{j=0}^k)$ of the random matrices C_j , for $j = 0, \dots, k$.

Proof of Proposition 1. As in the classic derivations of the cost-to-go functions for the LQ problem shown in Bertsekas (1995, section 4.1), first we solve the Bellman equation 3.12, for $k = N - 1$ and $k = N - 2$. Then we infer the form of its solution for $k = N - 3, \dots, 0$. For $k = N - 1$, one obtains

$$\begin{aligned}
 &\tilde{f}_{N-1}^\circ(\tilde{I}_{N-1}) \\
 &= \inf_{u_{N-1} \in \mathbb{R}^d} \mathbb{E}_{e_{N-1}, \tilde{I}_N} \{ e'_{N-1} Q_{N-1} e_{N-1} + \gamma u'_{N-1} u_{N-1} + \tilde{f}_N^\circ(\tilde{I}_N) | \tilde{I}_{N-1}, u_{N-1} \} \\
 &= \inf_{u_{N-1} \in \mathbb{R}^d} \mathbb{E}_{e_{N-1}, Q_N} \{ e'_{N-1} Q_{N-1} e_{N-1} + \gamma u'_{N-1} u_{N-1} \\
 &\quad + (e_{N-1} + u_{N-1})' Q_N (e_{N-1} + u_{N-1}) | \tilde{I}_{N-1}, u_{N-1} \}
 \end{aligned}$$

$$\begin{aligned}
 &= \mathbb{E}_{e_{N-1}} \{e'_{N-1} Q_{N-1} e_{N-1} | \tilde{I}_{N-1}\} + \inf_{u_{N-1} \in \mathbb{R}^d} \mathbb{E}_{e_{N-1}, Q_N} \{\gamma u'_{N-1} u_{N-1} \\
 &\quad + (e_{N-1} + u_{N-1})' Q_N (e_{N-1} + u_{N-1}) | \tilde{I}_{N-1}, u_{N-1}\}. \tag{A.5}
 \end{aligned}$$

For uniformity of notation with some of the next equations, from now on we set $K_N := Q_N$. Now, we observe that K_N is conditionally mutually independent from e_{N-1} and u_{N-1} given \tilde{I}_{N-1} and u_{N-1} and is mutually independent from \tilde{I}_{N-1} and u_{N-1} . Hence, by setting

$$\bar{K}_N := \mathbb{E}_{K_N} \{K_N\}$$

(a symmetric and positive-semidefinite matrix), one gets

$$\begin{aligned}
 &\mathbb{E}_{e_{N-1}, K_N} \{(e_{N-1} + u_{N-1})' K_N (e_{N-1} + u_{N-1}) | \tilde{I}_{N-1}, u_{N-1}\} \\
 &= \mathbb{E}_{e_{N-1}} \{(e_{N-1} + u_{N-1})' \bar{K}_N (e_{N-1} + u_{N-1}) | \tilde{I}_{N-1}, u_{N-1}\}. \tag{A.6}
 \end{aligned}$$

Combining equations A.5 and A.6, one has

$$\begin{aligned}
 &\tilde{J}_{N-1}(\tilde{I}_{N-1}) \\
 &= \mathbb{E}_{e_{N-1}} \{e'_{N-1} (Q_{N-1} + \bar{K}_N) e_{N-1} | \tilde{I}_{N-1}\} \\
 &\quad + \inf_{u_{N-1} \in \mathbb{R}^d} \left[u'_{N-1} (\bar{K}_N + \gamma I) u_{N-1} + 2(\bar{K}_N \mathbb{E}_{e_{N-1}} \{e_{N-1} | \tilde{I}_{N-1}\})' u_{N-1} \right], \tag{A.7}
 \end{aligned}$$

where I denotes the $d \times d$ identity matrix. Now, the matrix $\bar{K}_N + \gamma I$ is symmetric and positive-definite; hence, by the first-order optimal condition, the optimal updating function $u_{N-1}^\circ(\tilde{I}_{N-1})$ in equation A.7 is given by

$$u_{N-1}^\circ(\tilde{I}_{N-1}) = L_{N-1} \mathbb{E}_{e_{N-1}} \{e_{N-1} | \tilde{I}_{N-1}\},$$

where

$$L_{N-1} := -(\bar{K}_N + \gamma I)^{-1} \bar{K}_N.$$

Moreover, by putting $u_{N-1} = u_{N-1}^\circ(\tilde{I}_{N-1})$ into equation A.7, one obtains

$$\begin{aligned}
 &\tilde{J}_{N-1}^\circ(\tilde{I}_{N-1}) \\
 &= \mathbb{E}_{e_{N-1}} \{e'_{N-1} K_{N-1} e_{N-1} | \tilde{I}_{N-1}\} + \mathbb{E}_{e_{N-1}} \left\{ \left(e_{N-1} - \mathbb{E}_{e_{N-1}} \{e_{N-1} | \tilde{I}_{N-1}\} \right)' F_{N-1} \right.
 \end{aligned}$$

$$\times \left(e_{N-1} - \mathbb{E}_{e_{N-1}} \{e_{N-1} | \tilde{I}_{N-1}\} \right) \Big| \tilde{I}_{N-1} \Big\}, \quad (\text{A.8})$$

where

$$K_{N-1} := \bar{K}_N - \bar{K}_N(\bar{K}_N + \gamma I)^{-1} \bar{K}_N + Q_{N-1}, \quad (\text{A.9})$$

and

$$F_{N-1} := \bar{K}_N(\bar{K}_N + \gamma I)^{-1} \bar{K}_N \quad (\text{A.10})$$

are symmetric and positive-semidefinite matrices. Similarly, for the stage $k = N - 2$, the Bellman equation 3.12 becomes

$$\begin{aligned} & \tilde{J}_{N-2}^{\circ}(\tilde{I}_{N-2}) \\ &= \inf_{u_{N-2} \in \mathbb{R}^d} \mathbb{E}_{e_{N-2}, \tilde{I}_{N-1}} \{e'_{N-2} Q_{N-2} e_{N-2} + \gamma u'_{N-2} u_{N-2} + \tilde{J}_{N-1}^{\circ}(\tilde{I}_{N-1}) | \tilde{I}_{N-2}, u_{N-2}\} \\ &= \mathbb{E}_{e_{N-2}} \left\{ e'_{N-2} Q_{N-2} e_{N-2} | \tilde{I}_{N-2} \right\} \\ &+ \inf_{u_{N-2} \in \mathbb{R}^d} \left[\mathbb{E}_{e_{N-1}, K_{N-1}} \left\{ \gamma u'_{N-2} u_{N-2} + e'_{N-1} K_{N-1} e_{N-1} | \tilde{I}_{N-2}, u_{N-2} \right\} \right. \\ &+ \mathbb{E}_{e_{N-1}, \tilde{I}_{N-1}} \left\{ \left(e_{N-1} - \mathbb{E}_{e_{N-1}} \{e_{N-1} | \tilde{I}_{N-1}\} \right)' F_{N-1} \right. \\ &\left. \left. \times \left(e_{N-1} - \mathbb{E}_{e_{N-1}} \{e_{N-1} | \tilde{I}_{N-1}\} \right) | \tilde{I}_{N-2}, u_{N-2} \right\} \right]. \end{aligned} \quad (\text{A.11})$$

Now, by Bertsekas (1995, section 5.2, lemma 2.1), the term

$$\begin{aligned} & \mathbb{E}_{e_{N-1}, \tilde{I}_{N-1}} \left\{ \left(e_{N-1} - \mathbb{E}_{e_{N-1}} \{e_{N-1} | \tilde{I}_{N-1}\} \right)' F_{N-1} \right. \\ & \left. \times \left(e_{N-1} - \mathbb{E}_{e_{N-1}} \{e_{N-1} | \tilde{I}_{N-1}\} \right) | \tilde{I}_{N-2}, u_{N-2} \right\} \end{aligned} \quad (\text{A.12})$$

does not depend on u_{N-2} or on the sequence of updates applied up to time $N - 2$. This is basically due to the linearity of the dynamical system and the measurement equation.²³ By lemma 1, equation A.12 is a function

²³To prove Bertsekas (1995, section 5.2, lemma 2.1), it is assumed there that the matrices C_k are deterministic. However, inspection of the proof shows that the result still holds when they are random matrices, generated according to the model described in this letter.

$f_{N-2}(\{C_j\}_{j=0}^{N-2})$ of the random matrices C_j , for $j = 0, \dots, N-2$, whose realizations can be derived directly from the information vector \tilde{I}_{N-2} . Hence, the term (A.12) does not influence the search for an optimal update at the time $N-2$, so one obtains

$$\begin{aligned} \tilde{J}_{N-2}^{\circ}(\tilde{I}_{N-2}) &= \inf_{u_{N-2} \in \mathbb{R}^d} \mathbb{E}_{e_{N-2}, K_{N-1}} \{ \gamma u'_{N-2} u_{N-2} + (e_{N-2} + u_{N-2})' K_{N-1} \\ &\quad \times (e_{N-2} + u_{N-2}) | \tilde{I}_{N-2}, u_{N-2} \} + f_{N-2}(\{C_j\}_{j=0}^{N-2}) \\ &= \inf_{u_{N-2} \in \mathbb{R}^d} \mathbb{E}_{e_{N-2}, K_{N-1}} \{ \gamma u'_{N-2} u_{N-2} + (e_{N-2} + u_{N-2})' K_{N-1} \\ &\quad \times (e_{N-2} + u_{N-2}) | \tilde{I}_{N-2}, u_{N-2} \} \\ &\quad + \text{a term that does not depend on } u_{N-2}. \end{aligned} \quad (\text{A.13})$$

Such an optimization problem has the same nature as the one in equation A.5. Hence, by setting

$$\bar{K}_{N-1} := \mathbb{E}_{K_{N-1}} \{ K_{N-1} \} = \bar{K}_N - \bar{K}_N (\bar{K}_N + \gamma I)^{-1} \bar{K}_N + \mathbb{E}_{Q_{N-1}} \{ Q_{N-1} \}, \quad (\text{A.14})$$

the optimal updating function at the time $k = N-2$ is

$$u_{N-2}^{\circ}(\tilde{I}_{N-2}) = L_{N-2} \mathbb{E}_{e_{N-2}} \{ e_{N-2} | \tilde{I}_{N-2} \},$$

where

$$L_{N-2} := -(\bar{K}_{N-1} + \gamma I)^{-1} \bar{K}_{N-1}.$$

Moreover, by putting $u_{N-2} = u_{N-2}^{\circ}(\tilde{I}_{N-2})$ into equation A.7, one obtains

$$\begin{aligned} &\tilde{J}_{N-2}^{\circ}(\tilde{I}_{N-2}) \\ &= \mathbb{E}_{e_{N-2}} \{ e'_{N-2} K_{N-2} e_{N-2} | \tilde{I}_{N-2} \} \\ &\quad + \mathbb{E}_{e_{N-2}} \left\{ \left(e_{N-2} - \mathbb{E}_{e_{N-2}} \{ e_{N-2} | \tilde{I}_{N-2} \} \right)' F_{N-2} \left(e_{N-2} - \mathbb{E}_{e_{N-2}} \{ e_{N-2} | \tilde{I}_{N-2} \} \right) \middle| \tilde{I}_{N-2} \right\} \\ &\quad + f_{N-2}(\{C_j\}_{j=0}^{N-2}), \end{aligned} \quad (\text{A.15})$$

where

$$K_{N-2} := \bar{K}_{N-1} - \bar{K}_{N-1} (\bar{K}_{N-1} + \gamma I)^{-1} \bar{K}_{N-1} + Q_{N-2} \quad (\text{A.16})$$

and

$$F_{N-2} := \bar{K}_{N-1}(\bar{K}_{N-1} + \gamma I)^{-1} \bar{K}_{N-1}$$

are symmetric and positive-semidefinite matrices. Finally, by equation A.16, the matrix

$$\bar{K}_{N-2} := \mathbb{E}_{K_{N-2}} \{K_{N-2}\} = \bar{K}_{N-1} - \bar{K}_{N-1}(\bar{K}_{N-1} + \gamma I)^{-1} \bar{K}_{N-1} + \mathbb{E}_{Q_{N-2}} \{Q_{N-2}\},$$

is symmetric and positive-semidefinite.

Remark 16. By lemma 1, for $k = N - 3, \dots, 0$, also the terms

$$\mathbb{E}_{e_{k+1}, \tilde{I}_{k+1}} \left\{ (e_{k+1} - \mathbb{E}_{e_{k+1}} \{e_{k+1} | \tilde{I}_{k+1}\})' F_{k+1} (e_{k+1} - \mathbb{E}_{e_{k+1}} \{e_{k+1} | \tilde{I}_{k+1}\}) \middle| \tilde{I}_k, u_k \right\} \quad (\text{A.17})$$

do not depend on u_k or on the sequence of updates applied up to the time k . Moreover, lemma 1 shows that they are functions $f_k(\{C_j\}_{j=0}^k)$ of the random matrices C_j , for $j = 0, \dots, k$. Again, their realizations can be derived directly from the information vector \tilde{I}_k .

The same arguments used for the stages $k = N - 1$ and $k = N - 2$ can be applied to $k = N - 3, \dots, 0$. Proceeding in such a way, one gets the following recursion for $k = N - 1, \dots, 0$,

$$u_k^\circ(\tilde{I}_k) = L_k \mathbb{E}_{e_k} \{e_k | \tilde{I}_k\},$$

where we recall that

$$\begin{aligned} L_k &:= -(\bar{K}_{k+1} + \gamma I)^{-1} \bar{K}_{k+1}, \\ \tilde{f}_k(\tilde{I}_k) &= \mathbb{E}_{e_k} \{e_k' K_k e_k | \tilde{I}_k\} + \mathbb{E}_{e_k} \left\{ \left(e_k - \mathbb{E}_{e_k} \{e_k | \tilde{I}_k\} \right)' F_k \left(e_k - \mathbb{E}_{e_k} \{e_k | \tilde{I}_k\} \right) \middle| \tilde{I}_k \right\} \\ &\quad + \mathbb{E}_{\{C_j\}_{j=k+1}^{N-2}} \left\{ \sum_{h=k}^{N-2} f_h(\{C_j\}_{j=0}^h) \middle| \tilde{I}_k \right\}, \end{aligned} \quad (\text{A.18})$$

and the matrices

$$K_k := \bar{K}_{k+1} - \bar{K}_{k+1}(\bar{K}_{k+1} + \gamma I)^{-1} \bar{K}_{k+1} + Q_k,$$

$$F_k := \bar{K}_{k+1}(\bar{K}_{k+1} + \gamma I)^{-1} \bar{K}_{k+1}$$

and

$$\bar{K}_k := \mathbb{E}_{K_k} \{K_k\} = \bar{K}_{k+1} - \bar{K}_{k+1}(\bar{K}_{k+1} + \gamma I)^{-1} \bar{K}_{k+1} + \mathbb{E}_{Q_k} \{Q_k\}$$

are symmetric and positive-semidefinite. □

Proof of Proposition 2. Let us first show that the recursion 3.25 holds. To this end, we exploit the classic Kalman filter recursion scheme for the specific problem (see, e.g., Bertsekas, 1995, appendix E.3). This can be done since, at time k , the realization of the random matrix C_k becomes known to the learning machine; hence, one can apply the Kalman recursion to compute $\mathbb{E}\{e_k | \tilde{I}_k\}$. Indeed, such a recursion requires the knowledge of such a matrix at time k , not before. Note that, different from the analysis in Bertsekas (1995, appendix E.3), the (conditional) covariance matrix Σ_k in equation 3.20 depends actually on \tilde{I}_k through the realizations of the random matrices C_j , for $j = 0, \dots, k$. Instead, in the deterministic case, there is no dependence of such a covariance matrix on the information vector, so it is just an unconditional covariance matrix.

Now let $\hat{e}_k^{\circ,\dagger} := \mathbb{E}\{e_k^{\circ} | \tilde{I}_k\}$. By Bertsekas (1995, appendix E.3), the KF estimate of e_k° at time k , based on information vector \tilde{I}_k , is given by

$$\hat{e}_{k+1}^{\circ,\dagger} = \hat{e}_k^{\circ,\dagger} + L_k \hat{e}_k^{\circ,\dagger} + H_{k+1} \left(\tilde{y}_{k+1}^{\circ} - C_{k+1} (\hat{e}_k^{\circ,\dagger} + L_k \hat{e}_k^{\circ,\dagger}) \right), \quad (\text{A.19})$$

which is initialized by $\hat{e}_{-1}^{\circ,\dagger} := 0$, where the Kalman gain matrix H_{k+1} is defined in equation 3.24. Moreover, since $e_k^{\circ} = \hat{w}_k^{\circ} - w$, $e_{k+1}^{\circ} = \hat{w}_{k+1}^{\circ} - w$, $\tilde{y}_{k+1}^{\circ} := C_{k+1} \hat{w}_{k+1}^{\circ} - y_{k+1}$, and \hat{w}_k° is known at time k , the KF estimate \hat{w}_k^{\dagger} of w at time k , based on information vector I_k , satisfies²⁴

$$\hat{e}_k^{\circ,\dagger} = \hat{w}_k^{\circ} - \hat{w}_k^{\dagger}. \quad (\text{A.20})$$

Similarly, since \hat{w}_{k+1}° , C_{k+1} and y_{k+1} are known at time $k+1$, the KF estimate \hat{w}_{k+1}^{\dagger} of w at time $k+1$, based on information vector I_{k+1} , satisfies

$$\hat{e}_{k+1}^{\circ,\dagger} = \hat{w}_{k+1}^{\circ} - \hat{w}_{k+1}^{\dagger}. \quad (\text{A.21})$$

So, \hat{w}_{k+1}° is derived from equation A.19 by replacement of equations A.20 and A.21, obtaining

$$\begin{aligned} & (\hat{w}_{k+1}^{\circ} - \hat{w}_{k+1}^{\dagger}) \\ &= (\hat{w}_k^{\circ} - \hat{w}_k^{\dagger}) + L_k (\hat{w}_k^{\circ} - \hat{w}_k^{\dagger}) \\ & \quad + H_{k+1} \left((C_{k+1} \hat{w}_{k+1}^{\circ} - y_{k+1}) - C_{k+1} ((\hat{w}_k^{\circ} - \hat{w}_k^{\dagger}) + L_k (\hat{w}_k^{\circ} - \hat{w}_k^{\dagger})) \right). \end{aligned} \quad (\text{A.22})$$

By equations 3.13 and A.20 and the definition of the error vector $e_k := \hat{w}_k - w$, one obtains

²⁴To make the notation uniform, let $\hat{w}_{-1}^{\circ} := 0$ in equation A.20.

$$\hat{w}_{k+1}^\circ = \hat{w}_k^\circ + L_k(\hat{w}_k^\circ - \hat{w}_k^\dagger) = \hat{w}_k^\circ + L_k(\hat{e}_k^\circ - \hat{e}_k^{\circ,\dagger}). \quad (\text{A.23})$$

This, combined with equation A.22, provides

$$\hat{w}_{k+1}^\dagger = \hat{w}_k^\dagger + H_{k+1}(y_{k+1} - C_{k+1}\hat{w}_k^\dagger),$$

which is initialized by $\hat{w}_{-1}^\dagger := 0$. Finally, L_{-1} can be chosen arbitrarily, since it multiplies vectors with all-zero components (e.g., one can choose $L_{-1} := -(\bar{K}_0 + \gamma I)^{-1} \bar{K}_0$, as in the statement of the proposition). \square

Remark 17. Note that the update expressed by equation 3.23 does not depend on the sequence of applied updates. This could have been obtained more directly by considering the evolution of the dynamical system

$$w_{k+1} = w_k \quad (\text{A.24})$$

only, together with the initial condition $w_0 := w$, and the measurement equation 3.4.

Proof of Proposition 3. i. The first bound in equation 3.29 follows from the fact that Σ_{k+1} is a (conditional) covariance matrix. Let $\Sigma_k^{1/2}$ be the symmetric and positive-semidefinite square root of the matrix Σ_k and $M_{k+1} := C_{k+1}\Sigma_k^{1/2}$. The second bound follows by equation 3.25 and the fact that

$$\begin{aligned} \Sigma_k - \Sigma_{k+1} &= \Sigma_k^{1/2} \left(\Sigma_k^{1/2} C'_{k+1} (C_{k+1} \Sigma_k^{1/2} \Sigma_k^{1/2} C'_{k+1} + \sigma_\varepsilon^2)^{-1} C_{k+1} \Sigma_k^{1/2} \right) \Sigma_k^{1/2} \\ &= \Sigma_k^{1/2} M'_{k+1} (M_{k+1} M'_{k+1} + \sigma_\varepsilon^2)^{-1} M_{k+1} \Sigma_k^{1/2}. \end{aligned} \quad (\text{A.25})$$

Since $M'_{k+1} (M_{k+1} M'_{k+1} + \sigma_\varepsilon^2)^{-1} M_{k+1}$ is symmetric and positive-semidefinite, by equation A.25, $\Sigma_k - \Sigma_{k+1}$ is symmetric and positive-semidefinite too.

ii. Defining $N_{k+1} := \Sigma_k - \Sigma_{k+1}$, one has obviously

$$\Sigma_k = \Sigma_{k+1} + N_{k+1}, \quad (\text{A.26})$$

where, by equation 3.29, all the matrices involved in equation A.26 are symmetric and positive-semidefinite, so one gets

$$\text{Tr}\{\Sigma_k\} = \text{Tr}\{\Sigma_{k+1}\} + \text{Tr}\{N_{k+1}\} \geq \text{Tr}\{\Sigma_{k+1}\} \geq 0,$$

which proves equation 3.30. Moreover, by Weyl’s inequalities²⁵ of matrix perturbation theory, if one orders the eigenvalues of each of the three matrices Σ_k , Σ_{k+1} , and N_{k+1} in nondecreasing order taking into account their multiplicities, then for every $j = 1, \dots, d$ one gets

$$\lambda_j(\Sigma_{k+1}) + \lambda_1(\Sigma_{N_{k+1}}) \leq \lambda_j(\Sigma_k).$$

By the positive-semidefiniteness of N_{k+1} , for every $j = 1, \dots, d$ one has also

$$\lambda_j(\Sigma_{k+1}) \leq \lambda_j(\Sigma_k),$$

which implies (3.31).

iii. By Krall (2002, theorem II.5.4), every bounded and monotonic sequence of self-adjoint operators on a Hilbert space converges strongly to a self-adjoint operator. Then formula 3.32 is obtained as a finite-dimensional case of such a result.

iv. This part of the proof is based on the investigation of the limit behavior of equation A.25 for $k \rightarrow +\infty$, also using the expectation and trace operators. First, we exploit the assumption that the common probability distribution of the random vectors x_k has bounded support. This, together with the definition $M_{k+1} := C_{k+1} \Sigma_k^{1/2}$ and the bound 3.29, proves the existence of a positive constant c_1 such that

$$M_{k+1} M'_{k+1} = C_{k+1} \Sigma_k C'_{k+1} \leq C_{k+1} \Sigma_{-1} C'_{k+1} \leq c_1 \tag{A.27}$$

with a priori probability 1. Then one obtains

$$M'_{k+1} (M_{k+1} M'_{k+1} + \sigma_\varepsilon^2)^{-1} M_{k+1} \geq (c_1 + \sigma_\varepsilon^2)^{-1} M'_{k+1} M_{k+1}$$

with a priori probability 1. Moreover,

$$\Sigma_k^{1/2} M'_{k+1} (M_{k+1} M'_{k+1} + \sigma_\varepsilon^2)^{-1} M_{k+1} \Sigma_k^{1/2}$$

²⁵ Let $S_1, S_2 \in \mathbb{R}^{d \times d}$ and symmetric, and let their eigenvalues be ordered nondecreasingly with their multiplicities as

$$\begin{aligned} \lambda_1(S_1) &\leq \lambda_2(S_1) \leq \dots \leq \lambda_j(S_1) \leq \dots \leq \lambda_d(S_1), \\ \lambda_1(S_2) &\leq \lambda_2(S_2) \leq \dots \leq \lambda_j(S_2) \leq \dots \leq \lambda_d(S_2). \end{aligned}$$

Then, in their simplest form, Weyl’s inequalities (see, e.g., Bernstein, 2009, theorem 8.4.11) state that for every $j = 1, \dots, d$, one has

$$\lambda_1(S_1) + \lambda_j(S_2) \leq \lambda_j(S_1 + S_2) \leq \lambda_d(S_1) + \lambda_j(S_2).$$

$$\begin{aligned}
&\geq (c_1 + \sigma_\varepsilon^2)^{-1} \Sigma_k^{1/2} M'_{k+1} M_{k+1} \Sigma_k^{1/2} \\
&= (c_1 + \sigma_\varepsilon^2)^{-1} \Sigma_k C'_{k+1} C_{k+1} \Sigma_k,
\end{aligned} \tag{A.28}$$

where all the steps in equation A.28 hold with a priori probability 1.²⁶ Hence, exploiting properties of the trace operator and the independence between $C'_{k+1} C_{k+1}$ and Σ_k^2 , one obtains

$$\begin{aligned}
&\text{Tr} \left\{ \mathbb{E}_{\Sigma_k, M_{k+1}} \left\{ \Sigma_k^{1/2} M'_{k+1} (M_{k+1} M'_{k+1} + \sigma_\varepsilon^2)^{-1} M_{k+1} \Sigma_k^{1/2} \right\} \right\} \\
&\geq (c_1 + \sigma_\varepsilon^2)^{-1} \text{Tr} \left\{ \mathbb{E}_{\Sigma_k, C_{k+1}} \left\{ \Sigma_k C'_{k+1} C_{k+1} \Sigma_k \right\} \right\} \\
&= (c_1 + \sigma_\varepsilon^2)^{-1} \text{Tr} \left\{ \mathbb{E}_{\Sigma_k, C_{k+1}} \left\{ \Sigma_k^2 C'_{k+1} C_{k+1} \right\} \right\} \\
&= (c_1 + \sigma_\varepsilon^2)^{-1} \text{Tr} \left\{ \mathbb{E}_{\Sigma_k} \left\{ \Sigma_k^2 \right\} \mathbb{E}_{C_{k+1}} \left\{ C'_{k+1} C_{k+1} \right\} \right\} \\
&= (c_1 + \sigma_\varepsilon^2)^{-1} \text{Tr} \left\{ \mathbb{E}_{\Sigma_k} \left\{ \Sigma_k^2 \right\} \mathbb{E}_{Q_{k+1}} \left\{ Q_{k+1} \right\} \right\} \\
&= (c_1 + \sigma_\varepsilon^2)^{-1} \text{Tr} \left\{ \mathbb{E}_{\Sigma_k} \left\{ \Sigma_k^2 \right\} \bar{Q} \right\} \\
&\geq (c_1 + \sigma_\varepsilon^2)^{-1} \text{Tr} \left\{ \mathbb{E}_{\Sigma_k} \left\{ \Sigma_k^2 \right\} \right\} \lambda_{\min}(\bar{Q}),
\end{aligned} \tag{A.29}$$

where $\lambda_{\min}(\bar{Q})$ denotes the minimum eigenvalue of \bar{Q} , which is positive by the assumed positive-definiteness of \bar{Q} , and the last inequality in equation A.29 follows by Fang, Loparo, and Feng (1994, theorem 1).

²⁶To obtain the generalized inequality in equation A.28, we have exploited the fact that if S_1 and S_2 are symmetric and positive-semidefinite $d \times d$ matrices such that

$$S_1 \geq S_2,$$

then for every $d \times d$ matrix Σ , one also has

$$\Sigma' S_1 \Sigma \geq \Sigma' S_2 \Sigma.$$

This is proved by observing that for every $y \in \mathbb{R}^d$, one has

$$y' \Sigma' S_1 \Sigma y = (\Sigma y)' S_1 (\Sigma y) \geq (\Sigma y)' S_2 (\Sigma y) = y' \Sigma' S_2 \Sigma y.$$

At this point, we recall that for two symmetric and positive-semidefinite $d \times d$ matrices S_1 and S_2 , one has

$$|\text{Tr}\{S_1 S_2\}| \leq \sqrt{\text{Tr}\{S_1^2\}\text{Tr}\{S_2^2\}}, \tag{A.30}$$

which is the Cauchy-Schwarz inequality for the Hilbert-Schmidt norm $\sqrt{\text{Tr}\{S^2\}}$ (Conway, 1985). Hence, when $S_1 = \Sigma_k$ and $S_2 = I$, one obtains

$$|\text{Tr}\{\Sigma_k\}| = |\text{Tr}\{\Sigma_k I\}| \leq \sqrt{\text{Tr}\{\Sigma_k^2\}\text{Tr}\{I^2\}} = \sqrt{d \text{Tr}\{\Sigma_k^2\}}. \tag{A.31}$$

Then,

$$\begin{aligned} (\text{Tr}\{\bar{\Sigma}\})^2 &= \left(\lim_{k \rightarrow +\infty} \mathbb{E}_{\Sigma_k} \{\text{Tr}\{\Sigma_k\}\} \right)^2 \\ &= \lim_{k \rightarrow +\infty} \left(\mathbb{E}_{\Sigma_k} \{\text{Tr}\{\Sigma_k\}\} \right)^2 \\ &\leq \liminf_{k \rightarrow +\infty} \mathbb{E}_{\Sigma_k} \{(\text{Tr}\{\Sigma_k\})^2\} \\ &\leq \liminf_{k \rightarrow +\infty} \mathbb{E}_{\Sigma_k} \{d \text{Tr}\{\Sigma_k^2\}\}, \end{aligned} \tag{A.32}$$

where the last two inequalities derive, respectively, from the convexity of the function $\sqrt{\cdot}$ and Jensen’s inequality (Rudin, 1987, theorem 3.3), and from equation A.31.

Now, taking traces and expectations, making k tend to $+\infty$, and exploiting equations A.25, A.29, and A.32, we get

$$\begin{aligned} 0 &= \text{Tr}\{\bar{\Sigma}\} - \text{Tr}\{\bar{\Sigma}\} \\ &= \lim_{k \rightarrow +\infty} \text{Tr} \left\{ \mathbb{E}_{\Sigma_k} \{\Sigma_k\} \right\} - \lim_{k \rightarrow +\infty} \text{Tr} \left\{ \mathbb{E}_{\Sigma_{k+1}} \{\Sigma_{k+1}\} \right\} \\ &= \lim_{k \rightarrow +\infty} \text{Tr} \left\{ \mathbb{E}_{\Sigma_k, C_{k+1}} \left\{ \Sigma_k^{1/2} \left(\Sigma_k^{1/2} C'_{k+1} (C_{k+1} \Sigma_k^{1/2} \Sigma_k^{1/2} C'_{k+1} \right. \right. \right. \\ &\quad \left. \left. \left. + \sigma_\varepsilon^2 \right)^{-1} C_{k+1} \Sigma_k^{1/2} \right\} \Sigma_k^{1/2} \right\} \\ &\geq (c_1 + \sigma_\varepsilon^2)^{-1} \liminf_{k \rightarrow +\infty} \text{Tr} \{ \mathbb{E}_{\Sigma_k} \{\Sigma_k^2\} \} \lambda_{\min}(\bar{Q}) \\ &\geq (c_1 + \sigma_\varepsilon^2)^{-1} (\text{Tr} \{ \bar{\Sigma} \})^2 d^{-1} \lambda_{\min}(\bar{Q}). \end{aligned} \tag{A.33}$$

Hence, since $(c_1 + \sigma_\varepsilon^2)^{-1}$, d^{-1} , and $\lambda_{\min}(\bar{Q})$ are different from 0, equation A.33 implies

$$\text{Tr}\{\bar{\Sigma}\} = 0 \quad (\text{A.34})$$

and also

$$\bar{\Sigma} = 0,$$

since $\bar{\Sigma}$ is symmetric and positive-semidefinite. This concludes the proof of equation 3.34.

v. Let us denote by $F_{k+1}^{1/2}$ and $N_{k+1}^{1/2}$, respectively, a symmetric and positive-semidefinite square root of the symmetric and positive-semidefinite matrix F_{k+1} and a symmetric and positive-semidefinite square root of the symmetric and positive-semidefinite matrix $N_{k+1} := \Sigma_k - \Sigma_{k+1}$. Hence, one gets

$$\begin{aligned} \text{Tr}\{F_{k+1}(\Sigma_k - \Sigma_{k+1})\} &= \text{Tr}\{(F_{k+1}^{1/2})^2(N_{k+1}^{1/2})^2\} \\ &= \text{Tr}\{N_{k+1}^{1/2}(F_{k+1}^{1/2})^2N_{k+1}^{1/2}\} \\ &= \text{Tr}\{(N_{k+1}^{1/2}F_{k+1}^{1/2})(N_{k+1}^{1/2}F_{k+1}^{1/2})'\} \\ &\geq 0 \end{aligned} \quad (\text{A.35})$$

for all the realizations of the random matrices involved. Similarly, for all the realizations of the random matrices involved and for all $k = -1, 0, 1, 2, \dots$, one obtains

$$\text{Tr}\{F_{k+1}\Sigma_{k+1}\} \leq \text{Tr}\{F_{k+1}\Sigma_k\} \leq \dots \leq \text{Tr}\{F_{k+1}\Sigma_{-1}\} \quad (\text{A.36})$$

which is equation 3.36. \square

Proof of Proposition 4. In this proof, we use the superscript (N) to denote expressions obtained for the finite-horizon case with horizon N and $Q_N = 0$ with probability 1, and assuming that equation 3.33 holds for the other values of k . Due to equation A.18, for any sequence of feasible updates and any finite horizon N , one has

$$\begin{aligned} & \mathbb{E}_{e_0, \{x_k\}_{k=0}^{N-1}, \{\bar{\varepsilon}_k\}_{k=0}^{N-2}} \left\{ \sum_{k=0}^{N-1} [e'_k Q_k e_k + \gamma u'_k u_k] \right\} \\ & \geq \mathbb{E}_{\tilde{I}_0} \{ \tilde{f}_0^{c_0, (N)}(\tilde{I}_0) \} \\ & = \mathbb{E}_{\tilde{I}_0} \left\{ \mathbb{E}_{e_0} \left\{ e'_0 K_0^{(N)} e_0 \middle| \tilde{I}_0 \right\} + \mathbb{E}_{e_0} \left\{ \left(e_0 - \mathbb{E}_{e_0} \left\{ e_0 \middle| \tilde{I}_0 \right\} \right)' F_0^{(N)} \left(e_0 - \mathbb{E}_{e_0} \left\{ e_0 \middle| \tilde{I}_0 \right\} \right) \middle| \tilde{I}_0 \right\} \right\} \\ & \quad + \mathbb{E}_{\{C_j\}_{j=0}^{N-2}} \left\{ \sum_{h=0}^{N-2} f_h^{(N)}(\{C_j\}_{j=0}^h) \right\}. \end{aligned} \quad (\text{A.37})$$

Using steps similar to the ones made to obtain equations A.2 and A.3, and observing from equation A.3 that the functions $f_h^{(N)}$ can be written as

$$f_h^{(N)}((C_j)_{j=0}^h) = \mathbb{E}_{\Sigma_{h+1}} \left\{ \text{Tr} \left\{ F_{h+1}^{(N)} \Sigma_{h+1} \right\} \middle| C_0, \dots, C_h \right\},$$

one gets²⁷

$$\begin{aligned} & \mathbb{E}_{\tilde{I}_0} \left\{ \mathbb{E}_{e_0} \left\{ e_0' K_0^{(N)} e_0 \middle| \tilde{I}_0 \right\} + \mathbb{E}_{e_0} \left\{ \left(e_0 - \mathbb{E}_{e_0} \left\{ e_0 \middle| \tilde{I}_0 \right\} \right)' F_0^{(N)} \left(e_0 - \mathbb{E}_{e_0} \left\{ e_0 \middle| \tilde{I}_0 \right\} \right) \middle| \tilde{I}_0 \right\} \right\} \\ & + \mathbb{E}_{(C_j)_{j=0}^{N-2}} \left\{ \sum_{h=0}^{N-2} f_h^{(N)}((C_j)_{j=0}^h) \right\} \\ & = \text{Tr} \{ \bar{K}_0^{(N)} \Sigma_w \} + \mathbb{E}_{\Sigma_0} \left\{ \text{Tr} \{ F_0^{(N)} \Sigma_0 \} \right\} + \sum_{h=0}^{N-2} \mathbb{E}_{\Sigma_{h+1}} \left\{ \text{Tr} \{ F_{h+1}^{(N)} \Sigma_{h+1} \} \right\} \\ & = \text{Tr} \{ \bar{K}_0^{(N)} \Sigma_w \} + \sum_{h=0}^{N-1} \mathbb{E}_{\Sigma_h} \left\{ \text{Tr} \{ F_h^{(N)} \Sigma_h \} \right\} \\ & = \text{Tr} \{ \bar{K}_0^{(N)} \Sigma_w \} + \sum_{h=0}^{N-1} \text{Tr} \left\{ F_h^{(N)} \mathbb{E}_{\Sigma_h} \{ \Sigma_h \} \right\}. \end{aligned} \tag{A.38}$$

Hence, from equation A.37, one gets, for any feasible sequence of updates,

$$\begin{aligned} & \mathbb{E}_{e_0, \{x_k\}_{k=0}^{N-1}, \{\bar{\epsilon}_k\}_{k=0}^{N-2}} \left\{ \sum_{k=0}^{N-1} [e_k' Q_k e_k + \gamma u_k' u_k] \right\} \\ & \geq \text{Tr} \{ \bar{K}_0^{(N)} \Sigma_w \} + \sum_{h=0}^{N-1} \text{Tr} \left\{ F_h^{(N)} \mathbb{E}_{\Sigma_h} \{ \Sigma_h \} \right\}; \end{aligned} \tag{A.39}$$

then

$$\liminf_{N \rightarrow +\infty} \left(\frac{1}{N} \mathbb{E}_{e_0, \{x_k\}_{k=0}^{N-1}, \{\bar{\epsilon}_k\}_{k=0}^{N-2}} \left\{ \sum_{k=0}^{N-1} [e_k' Q_k e_k + \gamma u_k' u_k] \right\} \right)$$

²⁷ Here, one also applies the law of iterated expectations, together with the fact that the matrices $\bar{K}_0^{(N)}$ and $F_h^{(N)}$ are deterministic. Moreover, $e_0 e_0'$ and Q_0 are independent, which justifies the appearance of $\bar{K}_0^{(N)}$ and Σ_w in formula A.38.

$$\geq \liminf_{N \rightarrow +\infty} \left(\frac{1}{N} \left(\text{Tr}\{\bar{K}_0^{(N)} \Sigma_w\} + \sum_{h=0}^{N-1} \text{Tr} \left\{ F_h^{(N)} \mathbb{E}_{\Sigma_h} \{ \Sigma_h \} \right\} \right) \right). \tag{A.40}$$

Now, we show that the second lim inf in equation A.40 is actually a “lim,” and that

$$\lim_{N \rightarrow +\infty} \left(\frac{1}{N} \left(\text{Tr}\{\bar{K}_0^{(N)} \Sigma_w\} + \sum_{h=0}^{N-1} \text{Tr} \left\{ F_h^{(N)} \mathbb{E}_{\Sigma_h} \{ \Sigma_h \} \right\} \right) \right) = 0. \tag{A.41}$$

First, as N tends to $+\infty$, one gets

$$\lim_{N \rightarrow +\infty} \text{Tr}\{\bar{K}_0^{(N)} \Sigma_w\} = \text{Tr}\{\bar{K} \Sigma_w\} \tag{A.42}$$

by the convergence of $\bar{K}_0^{(N)}$ to the stationary solution \bar{K} of the ARE. Moreover, due to the definition of $F_h^{(N)}$, one gets

$$F_{h+1}^{(N+1)} = F_h^{(N)} = \dots = F_0^{(N-h)}, \tag{A.43}$$

for every N and $h = 0, \dots, N - 1$. From equation A.43 and the convergence of $F_0^{(N-h)}$ to \bar{F} as N tends to $+\infty$ (for fixed h), it follows that for every $\delta > 0$ there exists $t(\delta) \in \mathbb{N}_0$ such that, for every $N > t(\delta)$, one gets

$$-\delta I \preceq F_h^{(N)} - \bar{F} \preceq \delta I$$

for every $h = 0, \dots, N - t(\delta) - 1$, whereas

$$\begin{aligned} F_{N-t(\delta)}^{(N)} &= F_0^{(t(\delta))}, \\ F_{N-t(\delta)+1}^{(N)} &= F_0^{(t(\delta)-1)}, \\ &\dots, \\ F_{N-1}^{(N)} &= F_0^1 \end{aligned}$$

are a finite number $t(\delta)$ of fixed (i.e., independent from N) symmetric and positive-semidefinite matrices. Hence, recalling equations 3.21, 3.30, 3.36, and A.35 and the linearity of the trace operator, one obtains

$$0 \leq \sum_{h=0}^{N-1} \text{Tr} \left\{ F_h^{(N)} \mathbb{E}_{\Sigma_h} \{ \Sigma_h \} \right\}$$

$$\begin{aligned}
 &\leq \sum_{h=0}^{N-t(\delta)-1} \text{Tr} \left\{ (\bar{F} + \delta I) \mathbb{E}_{\Sigma_h} \{ \Sigma_h \} \right\} + \sum_{h=N-t(\delta)}^{N-1} \text{Tr} \left\{ F_h^{(N)} \mathbb{E}_{\Sigma_h} \{ \Sigma_h \} \right\} \\
 &\leq (N - t(\delta)) \delta \text{Tr} \left\{ \mathbb{E}_{\Sigma_{-1}} \{ \Sigma_{-1} \} \right\} + \sum_{h=0}^{N-t(\delta)-1} \text{Tr} \left\{ \bar{F} \mathbb{E}_{\Sigma_h} \{ \Sigma_h \} \right\} \\
 &\quad + \sum_{t=1}^{t(\delta)} \text{Tr} \left\{ F_0^{(t)} \mathbb{E}_{\Sigma_{-1}} \{ \Sigma_{-1} \} \right\} \\
 &= (N - t(\delta)) \delta \text{Tr} \{ \Sigma_w \} + \sum_{h=0}^{N-t(\delta)-1} \text{Tr} \left\{ \bar{F} \mathbb{E}_{\Sigma_h} \{ \Sigma_h \} \right\} + \sum_{t=1}^{t(\delta)} \text{Tr} \left\{ F_0^{(t)} \Sigma_w \right\}.
 \end{aligned}$$

This, combined with the finiteness of $\text{Tr} \left\{ \bar{F} \mathbb{E}_{\Sigma_0} \{ \Sigma_0 \} \right\}$ and

$$\text{Tr} \left\{ \bar{F} \mathbb{E}_{\Sigma_{h+1}} \{ \Sigma_{h+1} \} \right\} \leq \text{Tr} \left\{ \bar{F} \mathbb{E}_{\Sigma_h} \{ \Sigma_h \} \right\}$$

for every $h = 0, 1, \dots$, shows that

$$\begin{aligned}
 0 &\leq \liminf_{N \rightarrow +\infty} \left(\frac{1}{N} \sum_{h=0}^{N-1} \text{Tr} \left\{ F_h^{(N)} \mathbb{E}_{\Sigma_h} \{ \Sigma_h \} \right\} \right) \\
 &\leq \liminf_{N \rightarrow +\infty} \left(\text{Tr} \left\{ \bar{F} \mathbb{E}_{\Sigma_N} \{ \Sigma_N \} \right\} \right) + \delta \text{Tr} \{ \Sigma_w \}
 \end{aligned}$$

(see also the derivation of equation A.35 for a similar proof). Moreover, since this holds for every $\delta > 0$, one obtains

$$\begin{aligned}
 0 &\leq \liminf_{N \rightarrow +\infty} \left(\frac{1}{N} \sum_{h=0}^{N-1} \text{Tr} \left\{ F_h^{(N)} \mathbb{E}_{\Sigma_h} \{ \Sigma_h \} \right\} \right) \\
 &\leq \liminf_{N \rightarrow +\infty} \left(\text{Tr} \left\{ \bar{F} \mathbb{E}_{\Sigma_N} \{ \Sigma_N \} \right\} \right) \\
 &= \text{Tr} \left\{ \bar{F} \liminf_{N \rightarrow +\infty} \mathbb{E}_{\Sigma_N} \{ \Sigma_N \} \right\}, \\
 &= \text{Tr} \left\{ \bar{F} \lim_{N \rightarrow +\infty} \mathbb{E}_{\Sigma_N} \{ \Sigma_N \} \right\}, \\
 &= 0,
 \end{aligned} \tag{A.44}$$

where the last three steps are due to the linearity of the trace operator and to equations 3.32 and 3.34. This, combined with equation A.42, proves equation A.41. This means that the infimum of the average learning functional to be optimized, equation 4.4, is 0.

Now, we show that the sequence of updating functions 4.5 minimizes the average learning functional 4.4 with respect to all feasible sequences of updates. To see this, we use the superscript (N, \bar{K}) to denote expressions obtained for the finite-horizon case with horizon N and with $\bar{K}_N = \bar{K}$ (i.e., assuming that \bar{K}_N is equal to its stationary value \bar{K}), and assuming that equation 3.33 holds for the other values of k . Then, due to equation A.18, when the updates are generated by the sequence of updating functions 4.5, one has, for any finite horizon N ,

$$\begin{aligned}
 & \mathbb{E}_{e_0, \{x_k\}_{k=0}^{N-1}, \{\bar{\varepsilon}_k\}_{k=0}^{N-2}} \left\{ \sum_{k=0}^{N-1} [e'_k Q_k e_k + \gamma u'_k u_k] \right\} \\
 & \leq \mathbb{E}_{e_0, \{x_k\}_{k=0}^N, \{\bar{\varepsilon}_k\}_{k=0}^{N-2}} \left\{ \sum_{k=0}^{N-1} [e'_k Q_k e_k + \gamma u'_k u_k] + e'_N \bar{K} e_N \right\} \\
 & = \mathbb{E}_{\tilde{I}_0} \{ \tilde{J}_0^{\infty, (N, \bar{K})}(\tilde{I}_0) \} \\
 & = \mathbb{E}_{\tilde{I}_0} \left\{ \mathbb{E}_{e_0} \left\{ e'_0 \bar{K} e_0 \mid \tilde{I}_0 \right\} + \mathbb{E}_{e_0} \left\{ \left(e_0 - \mathbb{E}_{e_0} \left\{ e_0 \mid \tilde{I}_0 \right\} \right)' \bar{F} \left(e_0 - \mathbb{E}_{e_0} \left\{ e_0 \mid \tilde{I}_0 \right\} \right) \mid \tilde{I}_0 \right\} \right\} \\
 & \quad + \mathbb{E}_{\{C_j\}_{j=0}^{N-2}} \left\{ \sum_{h=0}^{N-2} f_h^{(N, \bar{K})}(\{C_j\}_{j=0}^h) \right\}. \tag{A.45}
 \end{aligned}$$

Moreover, likewise in the derivations of equations A.38 to A.44, one gets

$$\begin{aligned}
 & \mathbb{E}_{\tilde{I}_0} \left\{ \mathbb{E}_{e_0} \left\{ e'_0 \bar{K} e_0 \mid \tilde{I}_0 \right\} + \mathbb{E}_{e_0} \left\{ \left(e_0 - \mathbb{E}_{e_0} \left\{ e_0 \mid \tilde{I}_0 \right\} \right)' \bar{F} \left(e_0 - \mathbb{E}_{e_0} \left\{ e_0 \mid \tilde{I}_0 \right\} \right) \mid \tilde{I}_0 \right\} \right\} \\
 & \quad + \mathbb{E}_{\{C_j\}_{j=0}^{N-2}} \left\{ \sum_{h=0}^{N-2} f_h^{(N, \bar{K})}(\{C_j\}_{j=0}^h) \right\} \\
 & = \text{Tr} \{ \bar{K} \Sigma_w \} + \mathbb{E}_{\Sigma_0} \left\{ \text{Tr} \{ \bar{F} \Sigma_0 \} \right\} + \sum_{h=0}^{N-2} \mathbb{E}_{\Sigma_{h+1}} \left\{ \text{Tr} \{ \bar{F} \Sigma_{h+1} \} \right\} \\
 & = \text{Tr} \{ \bar{K} \Sigma_w \} + \sum_{h=0}^{N-1} \mathbb{E}_{\Sigma_h} \left\{ \text{Tr} \{ \bar{F} \Sigma_h \} \right\}
 \end{aligned}$$

$$= \text{Tr}\{\bar{K}\Sigma_w\} + \sum_{h=0}^{N-1} \text{Tr} \left\{ \bar{F}_{\Sigma_h} \mathbb{E} \{ \Sigma_h \} \right\}, \tag{A.46}$$

and

$$\begin{aligned} 0 &\leq \liminf_{N \rightarrow +\infty} \left(\frac{1}{N} \left(\text{Tr}\{\bar{K}\Sigma_w\} + \sum_{h=0}^{N-1} \text{Tr} \left\{ \bar{F}_{\Sigma_h} \mathbb{E} \{ \Sigma_h \} \right\} \right) \right) \\ &= \lim_{N \rightarrow +\infty} \left(\frac{1}{N} \left(\text{Tr}\{\bar{K}\Sigma_w\} + \sum_{h=0}^{N-1} \text{Tr} \left\{ \bar{F}_{\Sigma_h} \mathbb{E} \{ \Sigma_h \} \right\} \right) \right) \\ &= \lim_{N \rightarrow +\infty} \left(\frac{1}{N} \sum_{h=0}^{N-1} \text{Tr} \left\{ \bar{F}_{\Sigma_h} \mathbb{E} \{ \Sigma_h \} \right\} \right) \\ &\leq \lim_{N \rightarrow +\infty} \text{Tr} \left\{ \bar{F}_{\Sigma_N} \mathbb{E} \{ \Sigma_N \} \right\} \\ &= 0. \end{aligned} \tag{A.47}$$

This, combined with equations A.41, A.45, and A.46, proves the optimality of the sequence of updating functions 4.5 with respect to the average learning functional, 4.4, and all sequences of feasible updating functions. \square

Proof of Proposition 5. i. One has

$$\begin{aligned} \text{MSE}_k^\dagger &= \text{Tr} \left\{ \mathbb{E}_{w, \hat{w}_k^\dagger} \left\{ (w - \hat{w}_k^\dagger)' (w - \hat{w}_k^\dagger) \right\} \right\} \\ &= \text{Tr} \left\{ \mathbb{E}_{w, \hat{w}_k^\dagger} \left\{ (w - \hat{w}_k^\dagger) (w - \hat{w}_k^\dagger)'\right\} \right\} \\ &= \text{Tr} \left\{ \mathbb{E}_{w_k, \tilde{I}_k} \left\{ \left(w_k - \mathbb{E}_{w_k} \left\{ w_k \mid I_k \right\} \right) \left(w_k - \mathbb{E}_{w_k} \left\{ w_k \mid I_k \right\} \right)'\right\} \right\} \\ &= \text{Tr} \left\{ \mathbb{E}_{\tilde{I}_k} \left\{ \mathbb{E}_{w_k} \left\{ \left(w_k - \mathbb{E}_{w_k} \left\{ w_k \mid I_k \right\} \right) \left(w_k - \mathbb{E}_{w_k} \left\{ w_k \mid I_k \right\} \right)'\right\} \mid \tilde{I}_k \right\} \right\} \\ &= \text{Tr} \left\{ \mathbb{E}_{\tilde{I}_k} \left\{ \mathbb{E}_{e_k} \left\{ \left(- \left(e_k - \mathbb{E}_{e_k} \left\{ e_k \mid \tilde{I}_k \right\} \right) \right) \left(- \left(e_k - \mathbb{E}_{e_k} \left\{ e_k \mid \tilde{I}_k \right\} \right) \right)'\right\} \mid \tilde{I}_k \right\} \right\} \\ &= \text{Tr} \left\{ \mathbb{E}_{\Sigma_k} \{ \Sigma_k \} \right\}. \end{aligned} \tag{A.48}$$

ii. To obtain the rate of convergence to 0 of the mean square error of the KF estimate of w at the time k , we observe that, using similar steps as in the derivation of equations A.29 and A.33, one gets

$$\text{Tr}\{\mathbb{E}_{\Sigma_k}\{\Sigma_k\}\} - \text{Tr}\{\mathbb{E}_{\Sigma_{k+1}}\{\Sigma_{k+1}\}\} \geq (c_1 + \sigma_\varepsilon^2)^{-1} \text{Tr}\{\mathbb{E}_{\Sigma_k}\{\Sigma_k^2\}\} \lambda_{\min}(\bar{Q}). \quad (\text{A.49})$$

By iterating equation A.49, we obtain

$$\text{Tr}\{\mathbb{E}_{\Sigma_0}\{\Sigma_0\}\} - \text{Tr}\{\mathbb{E}_{\Sigma_{k+1}}\{\Sigma_{k+1}\}\} \geq (c_1 + \sigma_\varepsilon^2)^{-1} \left(\sum_{j=0}^k \text{Tr}\{\mathbb{E}_{\Sigma_j}\{\Sigma_j^2\}\} \right) \lambda_{\min}(\bar{Q}). \quad (\text{A.50})$$

This, combined with equation 3.31, provides

$$\begin{aligned} \text{Tr}\{\mathbb{E}_{\Sigma_0}\{\Sigma_0\}\} - \text{Tr}\{\mathbb{E}_{\Sigma_{k+1}}\{\Sigma_{k+1}\}\} &\geq (c_1 + \sigma_\varepsilon^2)^{-1} (k+1) \text{Tr}\{\mathbb{E}_{\Sigma_k}\{\Sigma_k^2\}\} \lambda_{\min}(\bar{Q}) \\ &\geq (c_1 + \sigma_\varepsilon^2)^{-1} (k+1) \text{Tr}\{\mathbb{E}_{\Sigma_{k+1}}\{\Sigma_{k+1}^2\}\} \lambda_{\min}(\bar{Q}). \end{aligned} \quad (\text{A.51})$$

Now, we apply the property A.30 with $S_1 = \Sigma_{k+1}$ and $S_2 = I$, obtaining

$$|\text{Tr}\{\Sigma_{k+1}\}| \leq \sqrt{\text{Tr}\{\Sigma_{k+1}^2\}} d;$$

hence

$$(\text{Tr}\{\Sigma_{k+1}\})^2 \leq \text{Tr}\{\Sigma_{k+1}^2\} d$$

and

$$\mathbb{E}_{\Sigma_{k+1}} \left\{ (\text{Tr}\{\Sigma_{k+1}\})^2 \right\} \leq \mathbb{E}_{\Sigma_{k+1}} \left\{ \text{Tr}\{\Sigma_{k+1}^2\} \right\} d = \text{Tr} \left\{ \mathbb{E}_{\Sigma_{k+1}} \{\Sigma_{k+1}^2\} \right\} d. \quad (\text{A.52})$$

Moreover, by the convexity of the square function $(\cdot)^2$ and Jensen's inequality, one gets

$$\mathbb{E}_{\Sigma_{k+1}} \left\{ (\text{Tr}\{\Sigma_{k+1}\})^2 \right\} \geq \left(\mathbb{E}_{\Sigma_{k+1}} \{ \text{Tr}\{\Sigma_{k+1}\} \} \right)^2 = \left(\text{Tr} \left\{ \mathbb{E}_{\Sigma_{k+1}} \{ \Sigma_{k+1} \} \right\} \right)^2. \quad (\text{A.53})$$

Then, combining equations A.51 to A.53, one obtains

$$\begin{aligned} & \text{Tr}\{\mathbb{E}_{\Sigma_0}\{\Sigma_0\}\} - \text{Tr}\{\mathbb{E}_{\Sigma_{k+1}}\{\Sigma_{k+1}\}\} \\ & \geq (c_1 + \sigma_\varepsilon^2)^{-1}(k+1) \left(\text{Tr} \left\{ \mathbb{E}_{\Sigma_{k+1}} \{ \Sigma_{k+1} \} \right\} \right)^2 d^{-1} \lambda_{\min}(\bar{Q}). \end{aligned} \quad (\text{A.54})$$

Now, for a given tolerance $\eta > 0$, we use equation A.54 to find an upper bound, as a function of η , on the maximal value $(k+1)(\eta)$ for $k+1$ for which $\text{Tr}\{\mathbb{E}_{\Sigma_{k+1}}\{\Sigma_{k+1}\}\} \geq \eta$. Since, of course, one has

$$\text{Tr}\{\mathbb{E}_{\Sigma_0}\{\Sigma_0\}\} \geq \text{Tr}\{\mathbb{E}_{\Sigma_0}\{\Sigma_0\}\} - \text{Tr}\{\mathbb{E}_{\Sigma_{k+1}}\{\Sigma_{k+1}\}\}, \quad (\text{A.55})$$

by combining the inequalities A.54 and A.55, and the definition of $(k+1)(\eta)$, one obtains

$$\begin{aligned} \text{Tr}\{\mathbb{E}_{\Sigma_0}\{\Sigma_0\}\} & \geq (c_1 + \sigma_\varepsilon^2)^{-1}(k+1)(\eta) \left(\text{Tr} \left\{ \mathbb{E}_{\Sigma_{(k+1)(\eta)}} \{ \Sigma_{(k+1)(\eta)} \} \right\} \right)^2 d^{-1} \lambda_{\min}(\bar{Q}) \\ & \geq (c_1 + \sigma_\varepsilon^2)^{-1}(k+1)(\eta) \eta^2 d^{-1} \lambda_{\min}(\bar{Q}) \\ & \geq 0. \end{aligned} \quad (\text{A.56})$$

Now, equation A.56 can hold only if

$$(k+1)(\eta) \leq \frac{(c_1 + \sigma_\varepsilon^2) d \text{Tr}\{\mathbb{E}_{\Sigma_0}\{\Sigma_0\}\}}{\eta^2 \lambda_{\min}(\bar{Q})}. \quad (\text{A.57})$$

Hence, renaming the index $k+1$ still by k ($k = 1, 2, \dots$), one obtains

$$k(\eta) \leq \frac{(c_1 + \sigma_\varepsilon^2) d \text{Tr}\{\mathbb{E}_{\Sigma_0}\{\Sigma_0\}\}}{\eta^2 \lambda_{\min}(\bar{Q})}. \quad (\text{A.58})$$

Similarly, for $k = 1, 2, \dots$, denoting by $\eta(k)$ the maximal value of $\eta > 0$ for which $\text{Tr}\{\mathbb{E}\{\Sigma_k\}\} \geq \eta$, one obtains

$$\eta(k) \leq \sqrt{\frac{(c_1 + \sigma_\varepsilon^2) d \text{Tr}\{\mathbb{E}\{\Sigma_0\}\}}{k \lambda_{\min}(\bar{Q})}}, \quad (\text{A.59})$$

which provides the desired rate of convergence, equation 5.1. Finally, the last part of ii follows from such a rate of convergence, from the definition of $\text{MSE}_{k'}^+$ and from part i.

iii. By equations 3.24, 3.32, and 3.34, and the assumed (uniform on k) almost-sure boundedness of $x_k = C'_k$, we get

$$\begin{aligned} \lim_{k \rightarrow +\infty} \mathbb{E}\{H_k\} &= \lim_{(k+1) \rightarrow +\infty} \mathbb{E}\{H_{k+1}\} \\ &= \lim_{k \rightarrow +\infty} \mathbb{E}\{\Sigma_{k+1} C'_{k+1} (\sigma_\varepsilon^2)^{-1}\} \\ &= \lim_{k \rightarrow +\infty} \mathbb{E}\{\Sigma_{k+1} C'_{k+1}\} (\sigma_\varepsilon^2)^{-1} \\ &= 0 \cdot (\sigma_\varepsilon^2)^{-1} \\ &= 0. \end{aligned}$$

iv. An application of equation 5.1, of the assumed (uniform on k) almost-sure boundedness of x_k , of Fang et al. (1994, theorem 1), and of the matrix version of Markov's inequality,²⁸ shows that for every $\delta > 0$, one has

$$\lim_{k \rightarrow +\infty} \Pr\{|H_{k,(h,l)}| > \delta\} = 0.$$

□

²⁸The matrix version of Markov's inequality (Recht, 2011, theorem A.1) states that for any symmetric and positive-semidefinite random matrix X and any fixed symmetric and positive-definite matrix M of the same dimension, one has

$$\Pr\{X \not\leq M\} \leq \text{Tr}\{\mathbb{E}\{X\}M^{-1}\}. \quad (\text{A.60})$$

Moreover, by Fang et al. (1994, theorem 2), one also has

$$\text{Tr}\{\mathbb{E}\{X\}M^{-1}\} \leq \text{Tr}\{\mathbb{E}\{X\}\} \frac{1}{\lambda_{\max}(M)}, \quad (\text{A.61})$$

where $\lambda_{\max}(M)$ is the largest eigenvalue of M . Then, one applies equations A.60 and A.61 with X replaced by Σ_k , and M by a sequence of positive-definite matrices M_j of the same dimension, with $\lambda_{\max}(M_j)$ decreasing and tending to 0 when j tends to $+\infty$.

Proof of Proposition 6. i. First, we observe that

$$\begin{aligned} \mathbb{E}\{e_k^\dagger\} &= \mathbb{E}_{w_k^\dagger, w} \{\hat{w}_k^\dagger - w\} = \mathbb{E}_{w, I_k} \left\{ \mathbb{E}_{w_k} \left\{ w_k \middle| I_k \right\} - w \right\} = \mathbb{E}_{w_k} \{w_k\} - \mathbb{E}\{w\} \\ &= \mathbb{E}_w \{w\} - \mathbb{E}_w \{w\} = 0. \end{aligned} \quad (\text{A.62})$$

Proceeding as in the proof of equation A.48, we get

$$\Sigma_{e_k^\dagger} = \mathbb{E}_{\Sigma_k} \{\Sigma_k\}. \quad (\text{A.63})$$

Similarly, recalling that

$$e_k^\circ := \hat{w}_k^\circ - w,$$

and using equations 3.22 and A.62, we get

$$\begin{aligned} \mathbb{E}_{e_{k+1}^\circ} \{e_{k+1}^\circ\} &= \mathbb{E}_{\hat{w}_{k+1}^\circ, w} \{\hat{w}_{k+1}^\circ - w\} \\ &= \mathbb{E}_{\hat{w}_k^\circ} \{\hat{w}_k^\circ\} + L_k \mathbb{E}_{\hat{w}_k^\circ, \hat{w}_k^\dagger} \{\hat{w}_k^\circ - \hat{w}_k^\dagger\} - \mathbb{E}_w \{w\} \\ &= \mathbb{E}_{\hat{w}_k^\circ} \{\hat{w}_k^\circ\} + L_k \mathbb{E}_{\hat{w}_k^\circ, w} \{\hat{w}_k^\circ - w\} - \mathbb{E}_w \{w\} \\ &= (I + L_k) \mathbb{E}_{\hat{w}_k^\circ, w} \{\hat{w}_k^\circ - w\} \\ &= (I + L_k) \mathbb{E}_{\hat{e}_k^\circ} \{\hat{e}_k^\circ\} \\ &= \Pi_{j=0}^k (I + L_j) \mathbb{E}_{\hat{w}_0^\circ, w} \{\hat{w}_0^\circ - w\} \\ &= \Pi_{j=0}^k (I + L_j) \mathbb{E}_{\hat{w}_0^\circ, w} \{\hat{w}_0 - w\} \\ &= 0. \end{aligned} \quad (\text{A.64})$$

Let

$$\begin{aligned} \Sigma_{e_k^\circ, e_k^\dagger} &:= \mathbb{E}_{e_k^\circ, e_k^\dagger} \left\{ \begin{pmatrix} e_k^\circ - \mathbb{E}\{e_k^\circ\} \\ e_k^\dagger - \mathbb{E}\{e_k^\dagger\} \end{pmatrix} \begin{pmatrix} e_k^\circ - \mathbb{E}\{e_k^\circ\} \\ e_k^\dagger - \mathbb{E}\{e_k^\dagger\} \end{pmatrix}' \right\} \\ &= \mathbb{E}_{e_k^\circ, e_k^\dagger} \left\{ (e_k^\circ) (e_k^\dagger)' \right\}, \end{aligned} \quad (\text{A.65})$$

and denote by

$$\begin{aligned} \Sigma_{e_k^\dagger, e_k^\circ} &:= \mathbb{E}_{e_k^\dagger, e_k^\circ} \left\{ \left(e_k^\dagger - \mathbb{E}_{e_k^\dagger} \{ e_k^\dagger \} \right) \left(e_k^\circ - \mathbb{E}_{e_k^\circ} \{ e_k^\circ \} \right)' \right\} \\ &= \mathbb{E}_{e_k^\dagger, e_k^\circ} \{ (e_k^\dagger) (e_k^\circ)' \} \end{aligned} \tag{A.66}$$

and

$$\Sigma_{e_k^\circ, e_k^\dagger} = \Sigma_{e_k^\dagger, e_k^\circ}' \tag{A.67}$$

the two (unconditional) cross-covariance matrices of e_k° and e_k^\dagger . Finally, we consider the vector

$$E_k := ((e_k^\circ)' , (e_k^\dagger)')'$$

and we denote by

$$\Sigma_{E_k} = \begin{pmatrix} \Sigma_{e_k^\circ} & \Sigma_{e_k^\circ, e_k^\dagger} \\ \Sigma_{e_k^\dagger, e_k^\circ} & \Sigma_{e_k^\dagger} \end{pmatrix}$$

its (unconditional) covariance matrix. As they are needed in the following analysis, we also provide the following upper bounds on the absolute values of the traces of the two matrices $(I + L_k) \Sigma_{e_k^\circ, e_k^\dagger} L_k'$ and $L_k \Sigma_{e_k^\dagger, e_k^\circ} (I + L_k)'$:

$$\begin{aligned} & \left| \text{Tr} \{ (I + L_k) \Sigma_{e_k^\circ, e_k^\dagger} L_k' \} \right| \\ &= \left| \text{Tr} \left\{ \mathbb{E}_{e_k^\circ, e_k^\dagger} \left\{ ((I + L_k) e_k^\circ) (L_k e_k^\dagger)' \right\} \right\} \right| \\ &= \left| \text{Tr} \left\{ \mathbb{E}_{e_k^\dagger, e_k^\circ} \left\{ (L_k e_k^\dagger)' ((I + L_k) e_k^\circ) \right\} \right\} \right| \\ &= \left| \mathbb{E}_{e_k^\dagger, e_k^\circ} \left\{ (L_k e_k^\dagger)' ((I + L_k) e_k^\circ) \right\} \right| \\ &\leq \sqrt{\mathbb{E}_{e_k^\dagger} \left\{ (L_k e_k^\dagger)' (L_k e_k^\dagger) \right\} \mathbb{E}_{e_k^\circ} \left\{ ((I + L_k) e_k^\circ)' ((I + L_k) e_k^\circ) \right\}} \\ &= \sqrt{\text{Tr} \left\{ \mathbb{E}_{e_k^\dagger} \left\{ (L_k e_k^\dagger)' (L_k e_k^\dagger) \right\} \right\} \text{Tr} \left\{ \mathbb{E}_{e_k^\circ} \left\{ ((I + L_k) e_k^\circ)' ((I + L_k) e_k^\circ) \right\} \right\}} \end{aligned}$$

$$\begin{aligned}
 &= \sqrt{\text{Tr} \left\{ \mathbb{E}_{e_k^\dagger} \left\{ (L_k e_k^\dagger) (L_k e_k^\dagger)' \right\} \right\} \text{Tr} \left\{ \mathbb{E}_{e_k^\circ} \left\{ ((I + L_k) e_k^\circ) ((I + L_k) e_k^\circ)' \right\} \right\}} \\
 &= \sqrt{\text{Tr} \{ L_k \Sigma_{e_k^\dagger} L_k' \} \text{Tr} \{ (I + L_k) \Sigma_{e_k^\circ} (I + L_k)' \}}, \tag{A.68}
 \end{aligned}$$

and, similarly,

$$\begin{aligned}
 \left| \text{Tr} \{ L_k \Sigma_{e_k^\dagger, e_k^\circ} (I + L_k)' \} \right| &= \left| \text{Tr} \left\{ \left((I + L_k) \Sigma_{e_k^\circ, e_k^\dagger} L_k' \right)' \right\} \right| \\
 &\leq \sqrt{\text{Tr} \{ L_k \Sigma_{e_k^\dagger} L_k' \} \text{Tr} \{ (I + L_k) \Sigma_{e_k^\circ} (I + L_k)' \}}. \tag{A.69}
 \end{aligned}$$

The mean square error of the OLL estimate \hat{w}_k° of w at time k is given by

$$\begin{aligned}
 \text{MSE}_k^\circ &= \mathbb{E}_{w, \hat{w}_k^\circ} \{ (w - \hat{w}_k^\circ)' (w - \hat{w}_k^\circ) \} \\
 &= \mathbb{E}_{e_k^\circ} \{ (-e_k^\circ)' (-e_k^\circ) \} = \mathbb{E}_{e_k^\circ} \{ (e_k^\circ)' (e_k^\circ) \} = \text{Tr} \left\{ \mathbb{E}_{e_k^\circ} \{ (e_k^\circ)' (e_k^\circ) \} \right\} \\
 &= \text{Tr} \left\{ \mathbb{E}_{e_k^\circ} \{ (e_k^\circ) (e_k^\circ)' \} \right\} \\
 &= \text{Tr} \{ \Sigma_{e_k^\circ} \}, \tag{A.70}
 \end{aligned}$$

which proves part i.

ii. By equation 3.22, one has the following recursion:

$$\begin{aligned}
 e_{k+1}^\circ &= e_k^\circ + L_k (e_k^\circ - e_k^\dagger), \\
 &= \begin{pmatrix} I + L_k \\ -L_k \end{pmatrix}' E_k. \tag{A.71}
 \end{aligned}$$

Hence, one gets

$$\begin{aligned}
 &\Sigma_{e_{k+1}^\circ} \\
 &= \begin{pmatrix} I + L_k \\ -L_k \end{pmatrix}' \Sigma_{E_k} \begin{pmatrix} I + L_k \\ -L_k \end{pmatrix} \\
 &= \begin{pmatrix} I + L_k \\ -L_k \end{pmatrix}' \begin{pmatrix} \Sigma_{e_k^\circ} & \Sigma_{e_k^\circ, e_k^\dagger} \\ \Sigma_{e_k^\dagger, e_k^\circ} & \Sigma_{e_k^\dagger} \end{pmatrix} \begin{pmatrix} I + L_k \\ -L_k \end{pmatrix} \\
 &= \begin{pmatrix} I + L_k \\ -L_k \end{pmatrix}' \begin{pmatrix} \Sigma_{e_k^\circ} (I + L_k)' - \Sigma_{e_k^\circ, e_k^\dagger} L_k' \\ \Sigma_{e_k^\dagger, e_k^\circ} (I + L_k)' - \Sigma_{e_k^\dagger} L_k' \end{pmatrix}
 \end{aligned}$$

$$= (I + L_k)\Sigma_{e_k^c}(I + L_k)' - (I + L_k)\Sigma_{e_k^c, e_k^+}L_k' - L_k\Sigma_{e_k^+, e_k^c}(I + L_k)' + L_k\Sigma_{e_k^+}L_k'.$$

Then, by using equations A.68 and A.69, we get

$$\begin{aligned} \text{Tr} \left\{ \Sigma_{e_{k+1}^c} \right\} &\leq \text{Tr} \left\{ (I + L_k)\Sigma_{e_k^c}(I + L_k)' \right\} \\ &\quad + \text{Tr} \left\{ L_k\Sigma_{e_k^+}L_k' \right\} \\ &\quad + 2\sqrt{\text{Tr} \left\{ (I + L_k)\Sigma_{e_k^c}(I + L_k)' \right\} \text{Tr} \left\{ L_k\Sigma_{e_k^+}L_k' \right\}}, \end{aligned} \tag{A.72}$$

which proves part ii.

iii. Let us now consider the infinite-horizon case. Then, L_k is replaced by L , and equation A.72 becomes

$$\begin{aligned} \text{Tr} \left\{ \Sigma_{e_{k+1}^c} \right\} &\leq \text{Tr} \left\{ (I + L)\Sigma_{e_k^c}(I + L)' \right\} \\ &\quad + \text{Tr} \left\{ L\Sigma_{e_k^+}L' \right\} \\ &\quad + 2\sqrt{\text{Tr} \left\{ (I + L)\Sigma_{e_k^c}(I + L)' \right\} \text{Tr} \left\{ L\Sigma_{e_k^+}L' \right\}}, \end{aligned} \tag{A.73}$$

where the matrix $(I + L)$ has all its eigenvalues inside the unit circle by Bertsekas (1995, section 4.1, proposition 4.1), and $\Sigma_{e_k^+}$ tends to the 0 matrix when k tends to $+\infty$, due to equations 3.32, 3.34, and A.63. We first show that the nonnegative sequence

$$\left\{ \text{Tr} \left\{ \Sigma_{e_k^c} \right\}, k = 0, 1, \dots \right\} \tag{A.74}$$

is bounded. Indeed, let M_0 be a symmetric and positive-semidefinite $d \times d$ matrix such that

$$\text{Tr} \{M_0\} \geq \text{Tr} \left\{ \Sigma_{e_0^c} \right\}. \tag{A.75}$$

We first note that the matrix L is symmetric by its definition, equation 4.2, since the symmetric matrices $(\bar{K} + I)^{-1}$ and \bar{K} commute, being associated with the same basis of eigenvectors. Then, by Fang et al. (1994, theorem 1), one has

$$\text{Tr} \left\{ (I + L)\Sigma_{e_0^c}(I + L)' \right\} \leq (|\lambda|_{\max}(I + L))^2 \text{Tr} \left\{ \Sigma_{e_0^c} \right\} < \text{Tr} \left\{ \Sigma_{e_0^c} \right\}, \tag{A.76}$$

since the spectral radius $|\lambda|_{\max}(I + L) < 1$, and

$$\text{Tr} \left\{ L\Sigma_{e_0^+}L' \right\} \leq (|\lambda|_{\max}(L))^2 \text{Tr} \left\{ \Sigma_{e_0^+} \right\}. \tag{A.77}$$

Combining inequalities A.76 and A.77 with equation A.73, we obtain

$$\begin{aligned}
 \text{Tr} \{ \Sigma_{e_1^e} \} &\leq (|\lambda|_{\max}(I + L))^2 \text{Tr} \{ \Sigma_{e_0^e} \} \\
 &\quad + |\lambda|_{\max}(L)^2 \text{Tr} \{ \Sigma_{e_0^t} \} \\
 &\quad + 2\sqrt{(|\lambda|_{\max}(I + L))^2 \text{Tr} \{ \Sigma_{e_0^e} \} (|\lambda|_{\max}(L))^2 \text{Tr} \{ \Sigma_{e_0^t} \}} \\
 &\leq (|\lambda|_{\max}(I + L))^2 \text{Tr} \{ M_0 \} \\
 &\quad + (|\lambda|_{\max}(L))^2 \text{Tr} \{ \Sigma_{e_0^t} \} \\
 &\quad + 2\sqrt{(|\lambda|_{\max}(I + L))^2 \text{Tr} \{ M_0 \} (|\lambda|_{\max}(L))^2 \text{Tr} \{ \Sigma_{e_0^t} \}}. \tag{A.78}
 \end{aligned}$$

Moreover, if $\text{Tr} \{ M_0 \}$ is sufficiently large, one gets²⁹

$$\begin{aligned}
 &(|\lambda|_{\max}(I + L))^2 \text{Tr} \{ M_0 \} + (|\lambda|_{\max}(L))^2 \text{Tr} \{ \Sigma_{e_0^t} \} \\
 &+ 2\sqrt{(|\lambda|_{\max}(I + L))^2 \text{Tr} \{ M_0 \} (|\lambda|_{\max}(L))^2 \text{Tr} \{ \Sigma_{e_0^t} \}} \\
 &\leq \text{Tr} \{ M_0 \}. \tag{A.79}
 \end{aligned}$$

Hence, one concludes

$$\text{Tr} \{ \Sigma_{e_1^e} \} \leq \text{Tr} \{ M_0 \}. \tag{A.80}$$

²⁹Inequality A.79 is of the form

$$a_1x + a_2 + a_3\sqrt{x} \leq x,$$

where $0 \leq a_1 < 1$, and $a_2 > 0, a_3 > 0$, and $x := \text{Tr} \{ M_0 \} > 0$, and it holds for

$$x \geq \left(\frac{a_3 + \sqrt{a_3^2 + 4a_2(1 - a_1)}}{2(1 - a_1)} \right)^2.$$

Since M_0 also has to satisfy equation A.75, one finally chooses

$$x \geq \max \left\{ \text{Tr} \{ \Sigma_{e_0^e} \}, \left(\frac{a_3 + \sqrt{a_3^2 + 4a_2(1 - a_1)}}{2(1 - a_1)} \right)^2 \right\}.$$

By similar reasoning, also for any $k = 2, 3, \dots$, one obtains

$$\mathrm{Tr} \left\{ \Sigma_{e_k^c} \right\} \leq \mathrm{Tr} \{M_0\}, \quad (\text{A.81})$$

since the sequence $\{\mathrm{Tr}\{\Sigma_{e_k^+}\}, k = 0, 1, \dots\}$ is nonincreasing (see equations 3.30 and A.63). This shows that the sequence A.74 is bounded.

Now, we investigate the convergence of $\mathrm{Tr} \left\{ \Sigma_{e_k^c} \right\}$ as k tends to $+\infty$. Let

$$\alpha := (|\lambda|_{\max}(I + L))^2 < 1,$$

and choose any $\beta > 0$ such that $\alpha + \beta < 1$. Moreover, let M_k denote any symmetric and positive-semidefinite $d \times d$ matrix such that

$$\mathrm{Tr} \{M_k\} \geq \mathrm{Tr} \left\{ \Sigma_{e_k^c} \right\}. \quad (\text{A.82})$$

Then, proceeding in the same way in the proof of equation A.78 and exploiting the fact that

$$|\lambda|_{\max}(L) \leq 2$$

(which follows from equation 4.9 and $|\lambda|_{\max}(I + L) \leq 1$), one obtains

$$\begin{aligned} & \mathrm{Tr} \left\{ \Sigma_{e_{k+1}^c} \right\} \\ & \leq (|\lambda|_{\max}(I + L))^2 \mathrm{Tr} \{M_k\} \\ & \quad + (|\lambda|_{\max}(L))^2 \mathrm{Tr} \left\{ \Sigma_{e_k^+} \right\} \\ & \quad + 2\sqrt{(|\lambda|_{\max}(I + L))^2 \mathrm{Tr} \{M_k\} (|\lambda|_{\max}(L))^2 \mathrm{Tr} \left\{ \Sigma_{e_k^+} \right\}} \\ & \leq \alpha \mathrm{Tr} \{M_k\} + 4 \mathrm{Tr} \left\{ \Sigma_{e_k^+} \right\} + 4\sqrt{\alpha \mathrm{Tr} \{M_k\} \mathrm{Tr} \left\{ \Sigma_{e_k^+} \right\}} \\ & \leq (\alpha + \beta) \mathrm{Tr} \{M_k\} \\ & \leq \mathrm{Tr} \{M_k\} \end{aligned} \quad (\text{A.83})$$

for any symmetric and positive-semidefinite $d \times d$ matrix M_k such that³⁰

$$\mathrm{Tr} \{M_k\} \geq \max \left\{ \mathrm{Tr} \left\{ \Sigma_{e_k^c} \right\}, \frac{4 \left(\sqrt{\alpha} + \sqrt{\alpha + \beta} \right)^2 \mathrm{Tr} \left\{ \Sigma_k^+ \right\}}{\beta^2} \right\}. \quad (\text{A.84})$$

³⁰As in footnote 29, formula A.84 is obtained by reducing equation A.83 to a quadratic inequality, expressing its solution through the two roots of the associated quadratic equality, and using equation A.82.

Moreover, if equation A.84 is satisfied, one gets

$$\text{Tr} \left\{ \Sigma_{e_{k+j}}^\circ \right\} \leq \text{Tr} \{M_k\} \tag{A.85}$$

also for any $j = 2, 3, \dots$, since the sequence $\{\text{Tr}\{\Sigma_{e_{k+j}}^\dagger\}, j = 0, 1, \dots\}$ is non-increasing. Finally, for a given value of β , we generate a sequence of symmetric and positive-semidefinite $d \times d$ matrices $M_0, M_1 \dots$ as follows:

- M_0 is chosen such that equation A.75 and A.79 are satisfied.
- For $k = 0, 1, \dots$:

$$M_{k+1} := \begin{cases} (\alpha + \beta)M_k & \text{if equation A.84 is satisfied;} \\ M_k & \text{otherwise.} \end{cases} \tag{A.86}$$

By construction, one has

$$\text{Tr}\{\Sigma_{e_k}^\circ\} \leq \text{Tr}\{M_k\} \tag{A.87}$$

for every $k = 0, 1, \dots$. Moreover, since, as already shown,

$$\lim_{k \rightarrow +\infty} \text{Tr}\{\Sigma_k^\dagger\} = 0,$$

the first condition in equation A.86 is satisfied an infinite number of times, which, combined with $0 < \alpha + \beta < 1$, shows that

$$\lim_{k \rightarrow +\infty} \text{Tr}\{M_k\} = 0;$$

hence, by equation A.87, also

$$\lim_{k \rightarrow +\infty} \text{Tr}\{\Sigma_{e_k}^\circ\} = 0,$$

which, combined with equation A.70, proves equation 5.6. □

Remark 18. An estimate of the rate of convergence to 0 of the MSE of the OLL estimate may be derived by combining the upper bound, equation 5.1, on the rate of convergence to 0 of the MSE of the KF estimate with some results contained in the proof of proposition 6—particularly, equation A.73 and the procedure used to generate the matrices M_k described in equation A.86.

Proof of Proposition 7. By proposition 6.iii and the Cauchy-Schwarz inequality, one gets

$$\lim_{k \rightarrow +\infty} \mathbb{E} \left\{ (u_k^\circ)' (u_k^\circ) \right\}$$

$$\begin{aligned}
 &= \lim_{k \rightarrow +\infty} \mathbb{E}_{\hat{w}_{k+1}^\circ, \hat{w}_k^\circ} \left\{ (\hat{w}_{k+1}^\circ - \hat{w}_k^\circ)' (\hat{w}_{k+1}^\circ - \hat{w}_k^\circ) \right\} \\
 &= \lim_{k \rightarrow +\infty} \mathbb{E}_{\hat{w}_{k+1}^\circ, \hat{w}_k^\circ, w} \left\{ ((\hat{w}_{k+1}^\circ - w) + (w - \hat{w}_k^\circ))' ((\hat{w}_{k+1}^\circ - w) + (w - \hat{w}_k^\circ)) \right\} \\
 &= \lim_{k \rightarrow +\infty} \mathbb{E}_{\hat{w}_{k+1}^\circ, w} \left\{ ((\hat{w}_{k+1}^\circ - w))' ((\hat{w}_{k+1}^\circ - w)) \right\} \\
 &\quad + \lim_{k \rightarrow +\infty} \mathbb{E}_{\hat{w}_k^\circ, w} \left\{ ((\hat{w}_k^\circ - w))' ((\hat{w}_k^\circ - w)) \right\} \\
 &\quad + 2 \lim_{k \rightarrow +\infty} \mathbb{E}_{\hat{w}_{k+1}^\circ, \hat{w}_k^\circ, w} \left\{ ((\hat{w}_{k+1}^\circ - w))' ((\hat{w}_k^\circ - w)) \right\} \\
 &\leq \lim_{k \rightarrow +\infty} \mathbb{E}_{\hat{w}_{k+1}^\circ, w} \left\{ ((\hat{w}_{k+1}^\circ - w))' ((\hat{w}_{k+1}^\circ - w)) \right\} \\
 &\quad + \lim_{k \rightarrow +\infty} \mathbb{E}_{\hat{w}_k^\circ, w} \left\{ ((\hat{w}_k^\circ - w))' ((\hat{w}_k^\circ - w)) \right\} \\
 &\quad + 2 \lim_{k \rightarrow +\infty} \sqrt{\mathbb{E}_{\hat{w}_{k+1}^\circ, w} \left\{ ((\hat{w}_{k+1}^\circ - w))' ((\hat{w}_{k+1}^\circ - w)) \right\}} \\
 &\quad \times \sqrt{\mathbb{E}_{\hat{w}_k^\circ, w} \left\{ ((\hat{w}_k^\circ - w))' ((\hat{w}_k^\circ - w)) \right\}} \\
 &= 0. \tag{A.88}
 \end{aligned}$$

Hence, the proof is concluded likewise for the KF estimates (see formulas 6.5 and 6.6 in section 6.1), exploiting again the fact that, by proposition 6.iii, one has $\lim_{k \rightarrow +\infty} \text{MSE}_k^\circ = 0$. \square

Proof of Proposition 8. By the optimality for the problem with $\gamma > 0$ and the limit problem, respectively, one has (with $e_0^\circ = e_0^\dagger = e_0$)

$$\begin{aligned}
 &\mathbb{E}_{\{e_k^\circ\}_{k=0}^N, \{u_k^\circ\}_{k=0}^{N-1}, \{Q_k\}_{k=0}^N} \left\{ \sum_{k=0}^{N-1} [(e_k^\circ)' Q_k (e_k^\circ) + \gamma (u_k^\circ)' (u_k^\circ)] + (e_N^\circ)' Q_N (e_N^\circ) \right\} \\
 &\leq \mathbb{E}_{\{e_k^\dagger\}_{k=0}^N, \{u_k^\dagger\}_{k=0}^{N-1}, \{Q_k\}_{k=0}^N} \left\{ \sum_{k=0}^{N-1} [(e_k^\dagger)' Q_k (e_k^\dagger) + \gamma (u_k^\dagger)' (u_k^\dagger)] + (e_N^\dagger)' Q_N (e_N^\dagger) \right\} \\
 &\tag{A.89}
 \end{aligned}$$

and

$$\mathbb{E}_{\{e_k^\dagger\}_{k=0}^N, \{Q_k\}_{k=0}^N} \left\{ \sum_{k=0}^{N-1} [(e_k^\dagger)' Q_k (e_k^\dagger)] + (e_N^\dagger)' Q_N (e_N^\dagger) \right\}$$

$$\leq \mathbb{E}_{\{e_k^\circ\}_{k=0}^N, \{Q_k\}_{k=0}^N} \left\{ \sum_{k=0}^{N-1} [(e_k^\circ)' Q_k (e_k^\circ)] + (e_N^\circ)' Q_N (e_N^\circ) \right\}. \tag{A.90}$$

Hence, combining equations A.89 and A.90, one obtains

$$\begin{aligned} & \mathbb{E}_{\{u_k^\circ\}_{k=0}^{N-1}} \left\{ \sum_{k=0}^{N-1} [\gamma (u_k^\circ)' (u_k^\circ)] \right\} \\ \leq & \mathbb{E}_{\{u_k^\circ\}_{k=0}^{N-1}} \left\{ \sum_{k=0}^{N-1} [\gamma (u_k^\circ)' (u_k^\circ)] \right\} \\ & + \mathbb{E}_{\{e_k^\circ\}_{k=0}^N, \{Q_k\}_{k=0}^N} \left\{ \sum_{k=0}^{N-1} [(e_k^\circ)' Q_k (e_k^\circ)] + (e_N^\circ)' Q_N (e_N^\circ) \right\} \\ & - \mathbb{E}_{\{e_k^\dagger\}_{k=0}^N, \{Q_k\}_{k=0}^N} \left\{ \sum_{k=0}^{N-1} [(e_k^\dagger)' Q_k (e_k^\dagger)] + (e_N^\dagger)' Q_N (e_N^\dagger) \right\} \\ \leq & \mathbb{E}_{\{u_k^\dagger\}_{k=0}^{N-1}} \left\{ \sum_{k=0}^{N-1} [\gamma (u_k^\dagger)' (u_k^\dagger)] \right\}. \end{aligned} \tag{A.91}$$

Concluding, as $\gamma > 0$, one gets the result. □

Proof of Proposition 9. We recall here equations 3.25, 3.24, and 3.23, which are needed to compute the KF estimate \hat{w}_k^\dagger :

$$\Sigma_{k+1} = \Sigma_k - \Sigma_k (C_{k+1}^{(\phi)})' ((C_{k+1}^{(\phi)}) \Sigma_k (C_{k+1}^{(\phi)})' + \sigma_\varepsilon^2)^{-1} (C_{k+1}^{(\phi)}) \Sigma_k, \tag{A.92}$$

$$H_{k+1} := \Sigma_{k+1} (C_{k+1}^{(\phi)})' (\sigma_\varepsilon^2)^{-1}, \tag{A.93}$$

$$\hat{w}_{k+1}^\dagger = \hat{w}_k^\dagger + H_{k+1} (y_{k+1} - (C_{k+1}^{(\phi)}) \hat{w}_k^\dagger), \tag{A.94}$$

where we have used the notation $C_k^{(\phi)}$ to recall that they are now defined using $\phi(x_k)$ instead than x_k ; that is, one has

$$C_k^{(\phi)} := (\phi(x_k))'.$$

We also recall the initializations

$$\hat{w}_{-1}^\dagger = 0$$

and

$$\Sigma_{-1} = \Sigma_w.$$

Then, for $k = -1$, using equations A.92 to A.94, we obtain

$$\begin{aligned} \Sigma_0 &= \nu I_{d_E} - \nu I_{d_E} \phi(x_0) (\nu (\phi(x_0))' \phi(x_0) + \sigma_\varepsilon^2)^{-1} (\phi(x_0))' \nu I_{d_E} \\ &= \nu I_{d_E} - \nu^2 \phi(x_0) (\nu \mathcal{K}(x_0, x_0) + \sigma_\varepsilon^2)^{-1} (\phi(x_0))', \\ H_0 &= (\nu I_{d_E} - \nu^2 \phi(x_0) (\nu \mathcal{K}(x_0, x_0) + \sigma_\varepsilon^2)^{-1} (\phi(x_0))') \phi(x_0) (\sigma_\varepsilon^2)^{-1} \\ &= (\nu - \nu^2 (\nu \mathcal{K}(x_0, x_0) + \sigma_\varepsilon^2)^{-1} \mathcal{K}(x_0, x_0)) (\sigma_\varepsilon^2)^{-1} \phi(x_0), \\ \hat{w}_0^\dagger &= (\nu - \nu^2 (\nu \mathcal{K}(x_0, x_0) + \sigma_\varepsilon^2)^{-1} \mathcal{K}(x_0, x_0)) (\sigma_\varepsilon^2)^{-1} y_0 \phi(x_0), \end{aligned}$$

and

$$\begin{aligned} (\hat{w}_0^\dagger)' \phi(x) &= (\nu - \nu^2 (\nu \mathcal{K}(x_0, x_0) + \sigma_\varepsilon^2)^{-1} \mathcal{K}(x_0, x_0)) (\sigma_\varepsilon^2)^{-1} y_0 (\phi(x_0))' \phi(x), \\ &= (\nu - \nu^2 (\nu \mathcal{K}(x_0, x_0) + \sigma_\varepsilon^2)^{-1} \mathcal{K}(x_0, x_0)) (\sigma_\varepsilon^2)^{-1} y_0 \mathcal{K}(x_0, x), \end{aligned}$$

where the last expression does not involve an explicit computation of $\phi(x_0)$ and $\phi(x)$. Similarly, by iterating the procedure above, one obtains, for every k , expressions of the form (see Liu et al., 2009, theorem 2)

$$\Sigma_k = \nu_k I_{d_E} - \Phi_k \Psi_k \Phi_k', \tag{A.95}$$

where ν_k and Ψ_k are, respectively, suitable scalars and suitable $(k + 1) \times (k + 1)$ matrices (both of which can be computed recursively), and

$$\Phi_k := (\phi(x_0), \dots, \phi(x_k)).$$

This, combined with equations A.93 and A.94, respectively, provides equations 7.3 to 7.5. □

Proof of Proposition 10. i. Following remark 10, the matrix (operator) L_k is proportional to I_{d_E} , say,

$$L_k = \alpha_{L_k} I_{d_E}$$

for some $\alpha_{L_k} > 0$. Then the update equation 3.22 of the OLL estimate of w becomes

$$\begin{aligned} \hat{w}_{k+1}^\circ &= \hat{w}_k^\circ + \alpha_{L_k} I_{d_E} (\hat{w}_k^\circ - \hat{w}_k^\dagger) \\ &= \hat{w}_k^\circ + \alpha_{L_k} (\hat{w}_k^\circ - \hat{w}_k^\dagger). \end{aligned} \tag{A.96}$$

This, combined with the initialization

$$\hat{w}_0^\circ = 0$$

and with equation 7.4 shows that, for $k = 1, 2, \dots$, one gets equations 7.9 and 7.10.

ii. Applying KPCA, one can show that the eigenvectors associated with the positive eigenvalues of $\overline{Q}_{\text{emp}}^{(\phi)}$ are linear combinations of $\phi(\tilde{x}_1), \dots, \phi(\tilde{x}_{l_U})$. Hence, following again remark 10, one obtains

$$L_k = \sum_{j=1}^{l_U} (\text{jth lin. comb. of } \phi(\tilde{x}_1), \dots, \phi(\tilde{x}_{l_U})) (\Lambda_{L_k})_{(j,j)} \\ \times (\text{same jth lin. comb. of } \phi(\tilde{x}_1), \dots, \phi(\tilde{x}_{l_U}))'$$

Then, reasoning as in the proof of part i, one obtains equations 7.11 and 7.12. \square

Acknowledgments

G. G. and M. S. are members of GNAMPA-INdAM (Gruppo Nazionale per l'Analisi Matematica, la Probabilità e le loro Applicazioni—Istituto Nazionale di Alta Matematica). The work by A. B. and G. G. was partially supported by the European Commission under Project H2020-SPIRE-636834, DISIRE (Distributed In-Situ Sensors Integrated into Raw Material and Energy Feedstock, <http://spire2030eu/disire/>).

References

- Alessandri, A., & Awawdeh, M. (2016). Moving-horizon estimation with guaranteed robustness for discrete-time linear systems and measurements subject to outliers. *Automatica*, 67, 85–93.
- Alessandri, A., Sanguineti, M., & Maggiore, M. (2002). Optimization-based learning with bounded error for feedforward neural networks. *IEEE Transactions on Neural Networks*, 15, 261–273.
- Antsaklis, P. J., & Michel, A. N. (2007). *A linear systems primer*. Cambridge, MA: Birkhäuser.
- Başar, T., & Bernhard, P. (2008). *\mathcal{H}_∞ -Optimal control and related minimax design problems: A dynamic game approach*. Cambridge, MA: Birkhäuser.
- Belkin, M., & Niyogi, P. (2006). Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *Journal of Machine Learning Research*, 7, 2299–2434.
- Bernstein, D. S. (2009). *Matrix mathematics: Theory, facts, and formulas*. Princeton, NJ: Princeton University Press.
- Bertsekas, D. P. (1995). *Dynamic programming and optimal control*, vol. 1. Belmont, MA: Athena Scientific.
- Bertsekas, D. P. (1996a). Incremental least squares methods and the extended Kalman filter. *SIAM Journal on Optimization*, 6, 807–822.

- Bertsekas, D. P. (1996b). *Neuro-dynamic programming*. Belmont, MA: Athena Scientific.
- Bertsekas, D. P., & Shrieve, S. E. (1978). Dynamic programming in Borel spaces. In M. Puterman (Ed.), *Dynamic programming and its applications* (pp. 115–130). Orlando, FL: Academic Press.
- Bierens, H. J. (2005). *Introduction to the mathematical and statistical foundations of econometrics*. Cambridge: Cambridge University Press.
- Camacho, E. F., & Bordons Alba, C. (2004). *Model predictive control*. New York: Springer.
- Chakrabarti, D., Wang, Y., Wang, D., Leskovec, J., & Faloutsos, C. (2008). Epidemic spreading in real networks. *ACM Transactions on Information and System Security*, 10, 1–26.
- Chang, G. (2014). Kalman filter with both adaptivity and robustness. *Journal of Process Control*, 24, 81–87.
- Conway, J. B. (1985). *A course in functional analysis*. New York: Springer.
- Cristianini, N., & Shawe-Taylor, J. (2000). *An introduction to support vector machines and other kernel-based learning methods*. Cambridge: Cambridge University Press.
- Crone, E. A., & Ridderinkhof, K. E. (2011). The developing brain: From theory to neuroimaging and back. *Developmental Cognitive Neuroscience*, 1, 101–109.
- Cucker, F., & Smale, S. (2002). Best choices for regularization parameters in learning theory: On the bias-variance problem. *Foundations of Computational Mathematics*, 2, 413–428.
- Davis, M. H. A., & Vinter, R. B. (1985). *Stochastic modelling and control*. London: Chapman and Hall.
- De Palma, D., & Indiveri, G. (2016). Output outlier robust state estimation. *International Journal of Adaptive Control and Signal Processing*. doi:10.1002/acs.2673
- Diethé, T., & Girolami, M. (2013). Online learning with (multiple) kernels: A review. *Neural Computation*, 25, 567–625.
- Fang, Y., Loparo, K. A., & Feng, X. (1994). Inequalities for the trace of matrix product. *IEEE Transactions on Automatic Control*, 39, 2489–2490.
- Gaggero, M., Gnecco, G., & Sanguineti, M. (2013). Dynamic programming and value-function approximation in sequential decision problems: Error analysis and numerical results. *Journal of Optimization Theory and Applications*, 156, 380–416.
- Gaggero, M., Gnecco, G., & Sanguineti, M. (2014). Approximate dynamic programming for stochastic n-stage optimization with application to optimal consumption under uncertainty. *Computational Optimization and Applications*, 58, 31–85.
- Gallieri, M., & Maciejowski, J. M. (2012). LASSO MPC: Smart regulation of over-actuated systems. In *Proceedings of the American Control Conference* (pp. 1217–1222). Piscataway, NJ: IEEE.
- Gnecco, G., Bemporad, A., Gori, M., Morisi, R., & Sanguineti, M. (2015). Online learning as an LQG optimal control problem with random matrices. In *Proceedings of the 14th IEEE European Control Conference* (pp. 2487–2494). Piscataway, NJ: IEEE.
- Gnecco, G., Gori, M., Melacci, S., & Sanguineti, M. (2015a). Foundations of support constraint machines. *Neural Computation*, 27, 388–480.
- Gnecco, G., Gori, M., Melacci, S., & Sanguineti, M. (2015b). Learning with mixed hard/soft pointwise constraints. *IEEE Transactions on Neural Networks and Learning Systems*, 26, 2019–2032.

- Gnecco, G., Gori, M., & Sanguineti, M. (2013). Learning with boundary conditions. *Neural Computation*, 25, 1029–1106.
- Gnecco, G., Morisi, R., & Bemporad, A. (2015). Sparse solutions to the average consensus problem via various regularizations of the fastest mixing Markov-chain problem. *IEEE Transactions on Network Science and Engineering*, 2, 97–111.
- Gnecco, G., & Sanguineti, M. (2010). Suboptimal solutions to dynamic optimization problems via approximations of the policy functions. *Journal of Optimization Theory and Applications*, 46, 746–794.
- Gori, M. (2009). Semantic-based regularization and Piaget's cognitive stages. *Neural Networks*, 22, 1036–1036.
- Grewal, M., & Andrews, A. (2001). *Kalman filtering: Theory and practice using Matlab*. New York: Wiley.
- Jay, E., Duvaut, P., Darolles, S., & Gouriéroux, C. (2011). l^q -regularization of the Kalman filter for exogenous outlier removal: Application to hedge funds analysis. In *Proceedings of the 4th IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing* (pp. 29–32). Piscataway, NJ: IEEE.
- Khan, J., Bouaynaya, N., & Fathallah-Shaykh, H. (2014). Tracking of time-varying genomic regulatory networks with a LASSO-Kalman smoother. *EURASIP Journal on Bioinformatics and Systems Biology*, 3.
- Krall, A. M. (2002). *Hilbert space, boundary value problems, and orthogonal polynomials*. Cambridge, MA: Birkhäuser.
- Lataire, J., Piga, D., & Tóth, R. (2014). Frequency-domain least-squares support vector machines to deal with correlated errors when identifying linear time-varying systems. In *Proceedings of the 19th IFAC World Congress* (pp. 10024–10029). Piscataway, NJ: IEEE.
- Li, Y., Gui, Q., Gu, Y., Han, S., & Du, K. (2014). Ridge-type Kalman filter and its algorithm. *WSEAS Transactions on Mathematics*, 13, 852–862.
- Liu, M., & Chang, G. (2016). Numerically and statistically stable Kalman filter for INS/GNSS integration. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, 230, 321–332.
- Liu, W., Park, I. M., Wang, Y., & Príncipe, J. C. (2009). Extended kernel recursive least squares algorithm. *IEEE Transactions on Signal Processing*, 57, 3801–3814.
- Maybeck, P. S. (1982). *Stochastic models, estimation, and control*, vol. 3. Orlando, FL: Academic Press, Inc.
- McGough, B. (2003). Statistical learning with time-varying parameters. *Macroeconomic Dynamics*, 7, 119–139.
- Mesbahi, M., & Egerstedt, M. (2010). *Graph theoretic methods in multiagent networks*. Princeton, NJ: Princeton University Press.
- Papoulis, A. (1991). *Probability, random variables, and stochastic processes*. New York: McGraw-Hill.
- Piaget, J. (1961). *La psychologie de l'intelligence*. Paris: Armand Colin.
- Ralaivola, L., & d'Alché Buc, F. (2005). Time series filtering, smoothing and learning using the kernel Kalman filter. In *Proceedings of the International Joint Conference on Neural Networks* (pp. 1449–1454). Piscataway, NJ: IEEE.
- Recht, B. (2011). A simpler approach to matrix completion. *Journal of Machine Learning Research*, 12, 3413–3430.
- Rudin, W. (1987). *Real and complex analysis*. New York: McGraw-Hill.

- Sayed, A. (2003). *Fundamentals of adaptive filtering*. New York: Wiley-Interscience.
- Schölkopf, B., Smola, A., & Müller, K.-R. (1998). Nonlinear component analysis as a kernel eigenvalue problem. *Neural Computation*, *10*, 1299–1319.
- Shalev-Shwartz, S. (2012). Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, *4*, 107–194.
- Smale, S., & Tao, Y. (2006). Online learning algorithms. *Foundations of Computational Mathematics*, *6*, 145–170.
- Söderström, T. (2002). *Discrete-time stochastic systems: Estimation and control*. New York: Springer-Verlag.
- Spall, J. C. (2003). *Introduction to stochastic search and optimization: Estimation, simulation, and control*. New York: Wiley-Interscience.
- Sutton, R. (1992). Gain adaptation beats least squares? In *Proceedings of the 7th Yale Workshop on Adaptive and Learning Systems* (pp. 161–166). New Haven, CT: Yale University.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning*. Cambridge, MA: MIT Press.
- Suykens, J., Vandewalle, J., & De Moor, B. (2001). Optimal control by least squares support vector machines. *Neural Networks*, *14*, 23–35.
- Tibshirani, R. (1996). Regression shrinkage and selection via the LASSO. *Journal of the Royal Statistical Society, Series B*, *58*, 267–288.
- Vu, K. M. (2007). *Optimal discrete control theory: The rational function structure model*. Ottawa: AuLac Technologies.
- Wang, K., Li, Y., & Rigos, C. (2012). Practical approaches to Kalman filtering with time-correlated measurement errors. *IEEE Transactions on Aerospace and Electronic Systems*, *48*, 1669–1681.
- Ying, Y., & Pontil, M. (2008). Online gradient descent algorithms. *Foundations of Computational Mathematics*, *8*, 561–596.
- Zinkevich, M. (2003). Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning* (pp. 928–936). New York: ACM.