

Distributed Multi-Agent Gaussian Regression via Finite-Dimensional Approximations

Gianluigi Pillonetto¹, Luca Schenato², and Damiano Varagnolo³

Abstract—We consider the problem of distributedly estimating Gaussian processes in multi-agent frameworks. Each agent collects few measurements and aims to collaboratively reconstruct a common estimate based on all data. Agents are assumed with limited computational and communication capabilities and to gather M noisy measurements in total on input locations independently drawn from a known common probability density. The optimal solution would require agents to exchange all the M input locations and measurements and then invert an $M \times M$ matrix, a non-scalable task. Differently, we propose two suboptimal approaches using the first E orthonormal eigenfunctions obtained from the Karhunen-Loève (KL) expansion of the chosen kernel, where typically $E \ll M$. The benefits are that the computation and communication complexities scale with E and not with M , and computing the required statistics can be performed via standard average consensus algorithms. We obtain probabilistic non-asymptotic bounds that determine a priori the desired level of estimation accuracy, and new distributed strategies relying on Stein's unbiased risk estimate (SURE) paradigms for tuning the regularization parameters and applicable to generic basis functions (thus not necessarily kernel eigenfunctions) and that can again be implemented via average consensus. The proposed estimators and bounds are finally tested on both synthetic and real field data.

Index Terms—Gaussian processes, sensor networks, distributed estimation, kernel-based regularization, nonparametric estimation, average consensus

1 INTRODUCTION

MANY modern engineering problems involve networks containing a large number of agents which have to cooperate to obtain a common goal. Several of these tasks can be seen as problems of function estimation from sparse and noisy data, a central issue in the machine learning field [1], [2]. Examples include the determination of the wind speed and direction field in a wind farm from local measurements of the turbines [3], the reconstruction of the temperature field in a datacenter from local measurements at each server [4], and weather forecasts [5], [6]. Traditional centralized machine-learning estimation approaches are computationally non-scalable when the network is large. Moreover parallelization of computation using client-server architectures, which can alleviate this problem, might not be feasible. This happens, e.g., in applications where communication is peer-to-peer, as in wireless sensor networks or multi-agent robotics, and where each agent is expected to have a common copy of the global estimate. In these cases, fully distributed cooperation approaches are sought [7].

1.1 State-of-the-Art

This paper considers a distributed nonparametric Gaussian regression approach. In this context, the unknown map is modeled as a zero-mean Gaussian process whose covariance (also called kernel in the machine learning literature) has to embed expected properties like smoothness [8], [9]. Other approaches to function estimation could be also adopted, e.g., sparse regression based on the ℓ_1 norm, automatic relevance determination or the elastic net [10], [11], [12], [13], [14]. However, in our framework the implementation of these approaches is not trivial and would require sophisticated distributed optimization algorithms like ADMM [15]. In fact, we consider a scenario where N agents first collect a total of M direct and noisy measurements of the unknown map on input locations drawn from a common and known probability density. The aim is then to obtain a shared function estimate. To simplify the exposition, we assume w.l.o.g. $N = M$, i.e., each agent collects a single measurement. We also assume that computational and data storage capabilities are limited, and that the communication network is peer-to-peer, i.e., agents are able only to communicate with a restricted number of neighbors. As described below, this makes the problem difficult also under Gaussian process assumptions, but we will see that function estimation can be performed using simple average operations.

Assuming that f and the measurements noise are jointly Gaussian, achieving the minimum variance estimate requires knowing all the M measurements and related input locations, plus invert an $M \times M$ matrix with $O(M^3)$ operations, a difficult task in a distributed fashion. When the data set size M is large, the complexity is high also in centralized contexts. Therefore, many alternative approaches have been

- G. Pillonetto and L. Schenato are with the Department of Information Engineering, University of Padova, Padova 35122, Italy. E-mail: {giapi, schenato}@dei.unipd.it.
- D. Varagnolo is with the Department of Computer Science, Electrical and Space Engineering, Luleå University of Technology, Luleå 971 87, Sweden. E-mail: damvar@ltu.se.

Manuscript received 24 July 2017; revised 8 May 2018; accepted 9 May 2018. Date of publication 17 June 2018; date of current version 13 Aug. 2019.

(Corresponding author: Damiano Varagnolo.)

Recommended for acceptance by R. Adams.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TPAMI.2018.2836422

developed relying, e.g., on the notion of pseudo input locations [16], [17], [18], the use of matrix factorizations [19] and approximations of the kernel function [20], [21] through the Nyström method or greedy techniques [22], [23], [24]. Along this way, KL expansions [25] have been also used to decompose the kernel in terms of eigenfunctions that are orthogonal w.r.t. the input locations probability density. One can then approximate the Gaussian process via the E kernel eigenfunctions associated to the largest eigenvalues, an approximation that corresponds to perform the best process approximation before seeing the data [25] (see Section 3.1 for more details). A posteriori, i.e., after seeing the measurements and their input locations, the situation is instead more subtle since there exist E -dimensional subspaces that allow to come closer to the minimum variance estimator [26]. However, the a priori basis given by the KL expansion has important advantages. In fact, as proved in [27], the first E kernel eigenfunctions are asymptotically optimal, i.e., they provide the best E -dimensional approximation of the minimum variance estimator as the data set size M grows to infinity. In addition, differently from the a posteriori basis described in [26], the a priori basis can be computed off-line. Moreover, as detailed in Section 3.2, computing the final estimates requires computing sufficient statistics that have the structure of averages of M local matrices and local vectors of dimension respectively $E \times E$ and E . This implies that the basic building block of the estimators involves computing averages over networks which can be more efficient from a memory, computation and communication perspective when $E \ll M$. Such averages can be computed via the so called *average consensus algorithms* [28], [29] which require only mild assumption on network connectivity and communication. In particular, these algorithms require no global topological information, only minimal local coordination and can be implemented also in the context of asynchronous updates and lossy communication [30].

1.2 Contribution

Our stream of research pairs the ones of other authors focusing on distributed kernel regression. An example is [31], that proposes a distributed regularized kernel Least Squares (LS) regression algorithm that exploits successive orthogonal projections, or [32] that extends [31] by designing strategies to reduce the communication and synchronization needs. Estimators with reduced order model complexity have been proposed in [33], while nonparametric schemes using Nearest-Neighbors interpolation strategies have been studied also in [34]. Another Gaussian estimation approach is considered in [35], with focus on the problem of sequentially predicting the most informative future input locations to minimize simultaneously the prediction error and the uncertainty in the regularization parameters. Other distributed regression algorithms are proposed in [36] with the aim of estimating a dynamic Gaussian process and its gradient, while in [37] authors develop a distributed learning and cooperative control algorithm where agents estimate a static field modeled as a network of radial basis functions whose centers locations are known in advance.

Despite the many research efforts, none of the aforementioned works on distributed regression have addressed the following fundamental issue: *assigned a Gaussian prior (the*

kernel) and the input locations distribution, how much information does the network need to exchange to obtain, with a probability $1 - \alpha$, the desired level of estimation accuracy? In this paper we will answer this question adopting KL-based strategies which exploit E kernel eigenfunctions. In particular, we will study two different estimators denoted by \hat{f}_A and \hat{f}_B which have computational and communication complexities of order $O(E^2)$ and $O(E)$, respectively, originally proposed in [38]. Differently from [38] which focused on finding Monte Carlo based strategies for assessing the a posteriori statistical performance of the estimators, in this work the focus is on characterizing their a priori prediction capability on future data by first assigning the kernel and the input locations statistics, and then deriving non-asymptotic error bounds that are functions of E , M and α . This analysis can be also seen as the extension to the Bayesian context of the concept of *effective dimension* developed in deterministic frameworks, e.g., in [39]. There it has been shown that, in the worst case, subspaces of dimension \sqrt{M} , i.e., sub-polynomial in the data set size, capture the estimate. Parallel to this, our bound returns information on the *Bayesian effective dimension* revealing which subspace can be really influenced by the measurements.

Another major contribution provided in this work is to show that both \hat{f}_A and \hat{f}_B are asymptotically optimal, i.e., for fixed E , as M grows to infinity there is no other estimator which can perform better in the mean squared error sense. We will also see that, while \hat{f}_A is always consistent, i.e., convergent in probability to the true function as $E, M \rightarrow \infty$, consistency of \hat{f}_B requires E to grow slower than M . In some sense, such result clarifies the price to pay when adopting a estimator parsimonious in the information exchange.

Finally, in many applications the kernel scale factor is unknown and its tuning is critical since it strongly affects the performance of the Bayesian estimator. In addition, the kernel expansion could be hard to be obtained and one would rather use a different set of basis functions. In the terminal part of the paper, we address these problems by proposing a novel distributed tuning strategy based on the SURE criterion [40]. Standard approaches proposed in the literature in the context of a centralized framework (like cross-validation and maximum likelihood [41], [42], [43], [44], MAP estimation [45], expected improvement [46] and Markov chain Monte Carlo [47], [48]) require high computation and communication overhead, and are therefore not suited for distributed implementations. Instead, our strategy allows for simultaneous hyperparameter tuning and function estimation via a single average consensus algorithm over a vector of size $O(E^2)$ when \hat{f}_A is employed, and via only two averages of size $O(E)$ when \hat{f}_B is employed. Very importantly, the SURE criterion be used also for generic basis functions, such as kernel sections or Nyström bases, thus not necessarily restricted to be kernel eigenfunctions and defined.

1.3 Paper Outline

The paper is organized as follows. Section 2 formulates the Bayesian estimation problem while Section 3 describes the KL expansion of the Gaussian process and the distributed estimators. Section 4 provides the statistical characterization of our distributed estimators, also deriving error bounds which are then tested via some numerical experiments. Section 5 proposes distributed strategies to tune the possibly

unknown regularization parameter entering our estimators for generic basis functions and discusses practical implementation issues. These strategies are also tested on both synthetic and real data. Section 6 collects conclusions and future research directions while proofs are collected in the Appendix, which can be found on the Computer Society Digital Library at <http://doi.ieeecomputersociety.org/10.1109/TPAMI.2018.2836422>.

2 BAYESIAN ESTIMATION

2.1 The Measurements Model

We consider the measurements model

$$y_m = f(x_m) + v_m, \quad m = 1, \dots, M \quad (1)$$

with the input locations x_m following the stochastic generation scheme

$$x_m \sim \mu(\mathcal{X}) \text{ i.i.d.}, \quad m = 1, \dots, M, \quad (2)$$

with μ a non-degenerate probability measure on the compact \mathcal{X} . The unknown function $f: \mathcal{X} \rightarrow \mathbb{R}$ is a zero-mean Gaussian process with continuous covariance $K: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, i.e.,

$$f \sim \mathcal{N}(0, K). \quad (3)$$

The measurement noise is also Gaussian of known variance σ_v^2 :

$$v_m \sim \mathcal{N}(0, \sigma_v^2).$$

Finally, $\{v_m\}_{m=1}^M$, $\{x_m\}_{m=1}^M$ and f are all assumed mutually independent.

2.2 The Bayesian Estimator

The Gaussian assumptions of Section 2.1 imply that the posterior of f given the dataset $\{x_m, y_m\}_{m=1}^M$ is still Gaussian. Also, the Maximum A Posteriori (MAP) estimator coincides with the minimum variance estimator and is given by

$$\hat{f}_{\text{MAP}}(x) = [K(x, x_1) \ \dots \ K(x, x_M)] H_{\text{MAP}} \begin{bmatrix} y_1 \\ \vdots \\ y_M \end{bmatrix}$$

with

$$H_{\text{MAP}} := \left(\begin{bmatrix} K(x_1, x_1) & \dots & K(x_1, x_M) \\ \vdots & & \vdots \\ K(x_M, x_1) & \dots & K(x_M, x_M) \end{bmatrix} + \sigma_v^2 I \right)^{-1}.$$

The storage and computational requirements needed to compute \hat{f}_{MAP} are thus $O(M^2)$ and $O(M^3)$, respectively. The communication complexity is either $O(\dim(\mathcal{X})M)$ if agents share the input locations x_m or $O(M^2)$ if they share the covariances $K(x_m, x_{m'})$. Thus, storage, computational and communication complexities do not scale favorably with the dataset size M . Our aim is thus to find good approximators of \hat{f}_{MAP} that are suitable for distributed implementations.

3 FINITE-DIMENSIONAL APPROXIMATIONS OF THE BAYESIAN ESTIMATOR

3.1 KL Expansion: Kernel

The kernel (3) can be expanded in terms of eigenfunctions ϕ_e orthonormal w.r.t. the measure μ in (2) and related eigenfunctions λ_e [25]. They are defined by

$$\lambda_e \phi_e(x) = \int_{\mathcal{X}} K(x, x') \phi_e(x') d\mu(x'), \quad (4)$$

$$K(x, x') = \sum_{e=1}^{+\infty} \lambda_e \phi_e(x) \phi_e(x') \quad \lambda_1 \geq \lambda_2 \dots > 0, \quad (5)$$

and, using δ_{ij} for the Kronecker delta,

$$\int_{\mathcal{X}} \phi_i(x) \phi_j(x) d\mu(x) = \delta_{ij}. \quad (6)$$

Let E be a positive integer. Then (4), (5) and (6) allow us to reformulate the process f via the following KL expansion

$$f(x) = \underbrace{\sum_{e=1}^E a_e \phi_e(x)}_{=: f_a(x)} + \underbrace{\sum_{e=1}^{+\infty} b_e \phi_{E+e}(x)}_{=: f_b(x)}. \quad (7)$$

The expansion coefficients have been thus divided into two sets: a finite one composed by the E random variables a_e , and an infinite one given by the remaining variables b_e . The elements in these two sets are all mutually independent, and satisfy

$$a_e \sim \mathcal{N}(0, \lambda_e), \quad e = 1, \dots, E \quad (8a)$$

$$b_e \sim \mathcal{N}(0, \lambda_{E+e}), \quad e = 1, 2, \dots \quad (8b)$$

It is well known that

$$\mathcal{S} := \text{span}\langle \phi_1(\cdot), \dots, \phi_E(\cdot) \rangle \quad (9)$$

is that E -dimensional subspace that captures the biggest part of the statistical energy of f as measured by $\mathbb{E}[f^2 d\mu]$. In other words, f_a is the best E -dimensional approximation of f in the mean square sense [27].

In what follows, it is always assumed that all the kernel eigenfunctions are contained in a ball of finite radius in the space of continuous functions, i.e.,

Assumption 1. *There exists a $k < +\infty$ s.t.*

$$\sup_{x \in \mathcal{X}} |\phi_e(x)| \leq \sqrt{k} < +\infty \quad e = 1, 2, \dots \quad (10)$$

Assumption 1 is satisfied by all the finite-dimensional kernels and also by classical covariances like the spline kernels, e.g., see [49] for the case of uniform μ . In practice, if the KL expansion is not available in closed form, it can be obtained numerically with arbitrary accuracy, as for example described in [50], also permitting to compute the constant k .

3.2 KL Expansion: Measurement Model

Our next step is to search for finite-dimensional estimators of f suitable for distributed implementations. Below, we

introduce two different estimators, denoted by \hat{f}_A and \hat{f}_B , which assume values in the finite-dimensional subspace \mathcal{S} defined in (9). First, it is useful to rewrite model (1) in a more compact form.

Let

$$\mathbf{x} := [x_1, \dots, x_M]^T$$

$$\mathbf{y} := [y_1, \dots, y_M]^T \quad \mathbf{v} := [v_1, \dots, v_M]^T \quad (11)$$

$$\mathbf{a} := [a_1, \dots, a_E]^T \quad \mathbf{b} := [b_1, b_2, \dots]^T$$

$$G := \begin{bmatrix} G_{11} & \dots & G_{1E} \\ \vdots & & \vdots \\ G_{M1} & \dots & G_{ME} \end{bmatrix} \quad Z := \begin{bmatrix} Z_{11} & Z_{12} & \dots \\ \vdots & \vdots & \vdots \\ Z_{M1} & Z_{M2} & \dots \end{bmatrix} \quad (12)$$

$$G_{me} := \phi_e(x_m), \quad m = 1, \dots, M, \quad e = 1, \dots, E,$$

$$Z_{me} := \phi_{E+e}(x_m), \quad m = 1, \dots, M, \quad e = 1, 2, \dots \quad (13)$$

Considering decomposition (7), definitions (11), (26), (27), (28), (29), (30), (13) and using classical algebraic notation to handle infinite-dimensional objects, the measurements model (1) becomes

$$\mathbf{y} = G\mathbf{a} + Z\mathbf{b} + \mathbf{v}.$$

With this novel notation $G\mathbf{a}$ accounts for the contribution from f_a while $Z\mathbf{b}$ accounts for the contribution from f_b .

3.3 The E -Dimensional Estimator \hat{f}_A

Let

$$\hat{f}_A(x) := [\phi_1(x) \cdots \phi_E(x)]H_A\mathbf{y}, \quad (14)$$

where

$$H_A := \left(\frac{G^T G}{M} + \frac{\sigma_v^2}{M} \Lambda_E^{-1} \right)^{-1} \frac{G^T}{M}$$

and $\Lambda_E := \text{diag}(\lambda_1, \dots, \lambda_E)$. The estimator \hat{f}_A is suitable for distributed computations. In fact, defining

$$G_m := [\phi_1(x_m), \dots, \phi_E(x_m)]$$

one has

$$\frac{G^T G}{M} = \frac{1}{M} \sum_{m=1}^M G_m^T G_m, \quad \frac{G^T \mathbf{y}}{M} = \frac{1}{M} \sum_{m=1}^M G_m^T y_m. \quad (15)$$

Since $G_m^T G_m \in \mathbb{R}^{E \times E}$ and $G_m^T y_m \in \mathbb{R}^E$ are local quantities, (15) points out that \hat{f}_A can be distributedly computed through the parallelization of two average consensus strategies: one on the $G_m^T G_m$'s and one on the $G_m^T y_m$'s, for a total of $E^2 + E$ scalars. This estimator would correspond to the Minimum Variance Unbiased Estimator (MVUE) estimator if the process f in (7) were truncated just to f_a .

3.4 The E -Dimensional Estimator \hat{f}_B

As stated in (6), one has

$$\mathbb{E} \left[\left[\frac{G^T G}{M} \right]_{e,e'} \right] = \int_{\mathcal{X}} \phi_e(x) \phi_{e'}(x) d\mu(x) = \delta_{e,e'}.$$

and, given the assumptions in Section 2.1 and Assumption 1, the following convergence in probability holds

$$\frac{G^T G}{M} = \frac{1}{M} \sum_{m=1}^M G_m^T G_m \xrightarrow{M \rightarrow +\infty} \mathbb{E} \left[\frac{G^T G}{M} \right] = I.$$

Thus, it is tempting to use the approximation

$$\frac{G^T G}{M} \approx I \quad (16)$$

and use, in place of H_A in (14), the matrix

$$H_B := \left(I + \frac{\sigma_v^2}{M} \Lambda_E^{-1} \right)^{-1} \frac{G^T}{M}.$$

In turn, this approach approximates \hat{f}_A with

$$\hat{f}_B(x) := [\phi_1(x) \cdots \phi_E(x)]H_B\mathbf{y}.$$

The estimator \hat{f}_B is more advantageous than \hat{f}_A for distributed computations. In fact, it requires an average consensus on just the column vectors $G_m^T y_m$'s, for a total of E scalars (differently from the $E^2 + E$ ones required by \hat{f}_A), and does not require any expensive matrix inversion since $I + \frac{\sigma_v^2}{M} \Lambda_E^{-1}$ is diagonal.

4 STATISTICAL ANALYSIS OF \hat{f}_A AND \hat{f}_B

Ideally one would like to compute $\mathbb{E}[\|f - \hat{f}_A\|^2]$ and $\mathbb{E}[\|f - \hat{f}_B\|^2]$, or at least some bounds that quantify the performance of the estimator for any specified E and M a priori. However, the computation of such quantities is intractable or, at least, requires an expensive Monte Carlo analysis, possibly to be repeated for many different design variables like, e.g., M, E, σ_v^2 . To circumvent this challenge, we will exploit the assumption that the input locations are randomly drawn from a known distribution μ and the orthonormality of the eigenfunctions to find bounds on $\mathbb{E}[\|f - \hat{f}_A\|^2]$ that hold with arbitrarily high probability. More specifically, the key idea is to find an event \mathcal{E} that occurs with arbitrarily high probability such that informative bounds on $\mathbb{E}[\|f - \hat{f}_A\|^2 | \mathcal{E}]$ can be computed. This is formally described in the next sections.

4.1 Performance Indexes and Lower Bound

Two important performance indexes we consider for \hat{f}_A and \hat{f}_B are the errors defined by the conditional expectations

$$\text{Err}_A(x) := \mathbb{E} \left[\|f - \hat{f}_A\|^2 \mid \mathbf{x} \right]$$

$$\text{Err}_B(x) := \mathbb{E} \left[\|f - \hat{f}_B\|^2 \mid \mathbf{x} \right],$$

where

$$\|g\|^2 := \int_{\mathcal{X}} g^2(x) d\mu(x).$$

The variables $\text{Err}_A(x)$ and $\text{Err}_B(x)$ are stochastic, since they are functions of the random input locations \mathbf{x} that in our settings are assumed random as described in (2). Hence, the crux of our analysis will be how to account for the

randomness coming from x . Note also that $\|\cdot\|$ depends on μ so that Err_A and Err_B quantify the prediction errors on future data independently drawn from the same training set distribution.

Exploiting the KL expansion introduced in Section 3.1 a lower bound on the errors $\text{Err}_A(x)$ and $\text{Err}_B(x)$ can be also easily obtained. More generally, the following result bounds the performance achievable by any generic E -dimensional estimator of f .

Theorem 1. *Let \hat{f}_* be any generic estimator of f , function of x and y and assuming values in any generic E -dimensional space fixed a priori. Then*

$$\min_{\hat{f}_*} \mathbb{E} \left[\|f - \hat{f}_*\|^2 \mid \mathbf{x} \right] \geq \sum_{e=E+1}^{+\infty} \lambda_e. \quad (17)$$

The following definition will be especially important for our future developments.

Definition 1. *We say that $\overline{\text{Err}}_A \leq q$ or $\overline{\text{Err}}_B \leq q$ with probability $1 - \alpha$ if there exists an event \mathcal{E} in the σ -algebra induced by \mathbf{x} of probability at least $1 - \alpha$ such that, respectively,*

$$\mathbb{E}[\text{Err}_A(\mathbf{x}) \mid \mathcal{E}] \leq q \quad (18)$$

or

$$\mathbb{E}[\text{Err}_B(\mathbf{x}) \mid \mathcal{E}] \leq q. \quad (19)$$

Thus, if α is close to zero saying that $\overline{\text{Err}}_A \leq q$ with probability $1 - \alpha$ is equivalent to saying that the average error associated to \hat{f}_A is smaller than q with high probability. Finally, note that setting \mathcal{E} to the entire sample space, the conditional expectations in the Left Hand Side (LHS) of (18) and (19) become unconditional ones, and actually correspond to the Mean Square Errors (MSEs) of \hat{f}_A and \hat{f}_B , i.e.,

$$\text{MSE}_{\hat{f}_A} = \int_{\mathcal{X}} \text{Err}_A(x) d\mu(x), \quad (20)$$

$$\text{MSE}_{\hat{f}_B} = \int_{\mathcal{X}} \text{Err}_B(x) d\mu(x). \quad (21)$$

4.2 Non Asymptotic Error Bounds

The key issue is to bound the performance indexes Err_A and Err_B for any finite number of measurements M and eigenfunctions E . The following theorem provides the desired bounds. It depends on the input locations distribution μ , the kernel eigenvalues λ_e and constant k defined in (4) and (10), the number of eigenfunctions E and measurements M . In addition the bound is also function of a parameter $\varepsilon \in (0, 1]$ connected to maximal and minimal (stochastic) eigenvalue of $\frac{G^T G}{M}$, as detailed in the proof contained in the Appendix, available in the online supplemental material.

Theorem 2. *Let the assumptions in Section 2.1 and Assumption 1 hold, $\alpha \in (0, 1)$ be a desired confidence level (e.g., 0.01 or 0.05), and $\varepsilon \in (0, 1]$ be given. If E, M and k satisfy*

$$1 - \varepsilon + \varepsilon \log(\varepsilon) \geq \frac{Ek}{M} \log\left(\frac{E}{\alpha}\right) \quad (22)$$

then with probability at least $1 - \alpha$ it holds that

$$\overline{\text{Err}}_A \leq \text{Bnd}_A$$

with

$$\begin{aligned} \text{Bnd}_A := & \frac{kM}{1 - \alpha} \left(\sum_{e=1}^E \frac{\lambda_e^2}{(\varepsilon M \lambda_e + \sigma_v^2)^2} \right) \left(\sum_{e=E+1}^{+\infty} \lambda_e \right) \\ & + \frac{\sigma_v^2}{1 - \alpha} \left(\sum_{e=1}^E \frac{\lambda_e}{\varepsilon M \lambda_e + \sigma_v^2} \right) + \left(\sum_{e=E+1}^{+\infty} \lambda_e \right). \end{aligned}$$

Under the same assumption but with E, M and k now satisfying

$$1 - \varepsilon + \varepsilon \log(\varepsilon) \geq \frac{Ek}{M} \log\left(\frac{2E}{\alpha}\right), \quad (23)$$

then with probability at least $1 - \alpha$ it holds that

$$\overline{\text{Err}}_B \leq \text{Bnd}_B$$

with

$$\begin{aligned} \text{Bnd}_B := & \frac{kM}{1 - \alpha} \left(\sum_{e=1}^E \frac{\lambda_e^2}{(M \lambda_e + \sigma_v^2)^2} \right) \left(\sum_{e=E+1}^{+\infty} \lambda_e \right) \\ & + \frac{\sigma_v^2}{1 - \alpha} \left(\sum_{e=1}^E \frac{\lambda_e}{\varepsilon M \lambda_e + \sigma_v^2} \right) + \left(\sum_{e=E+1}^{+\infty} \lambda_e \right) \\ & + \kappa \left(\frac{E}{M} \sigma_v^2 + \sum_{e=1}^E \lambda_e \right), \end{aligned}$$

where

$$\kappa = \frac{1}{1 - \alpha} \left(\varepsilon + \frac{\lambda_1^{-1} \sigma_v^2}{M} \right)^{-4} (1 - \varepsilon)^2 (2 - \varepsilon)^2.$$

The obtained bounds are now tested via a numerical example.

4.3 Numerical Study

Consider the first-order spline kernel [51] which corresponds to the Brownian motion covariance, i.e.,

$$K(x, x') = \min(x, x') = \sum_{e=1}^{\infty} \lambda_e \phi_e(x) \phi_e(x')$$

with the input locations probability measure μ in (2) set to the uniform distribution on $[0, 1]$. With these settings

$$\phi_e(x) = \sqrt{2} \sin(x(e\pi - \pi/2)), \quad \lambda_e = \frac{1}{(e\pi - \pi/2)^2}$$

and $k = 2$. To make the bounds only depend on E we set $M = 10000$, $1 - \alpha = 0.95$, the noise variance $\sigma_v^2 = 0.1^2$, and $\varepsilon \in (0, 1]$ that minimizes the bound while satisfying (22) or (23) accordingly.

The thick lines in the two top panels of Fig. 1 show how Bnd_A (left) and Bnd_B (right) vary with E (bounds are normalized using the prior process variance $\sum_{e=1}^{\infty} \lambda_e$). For the sake of comparison we also display the true (normalized) MSEs (dashed line) as defined in (20) and (21), calculated via a Monte Carlo of 1000 runs, and its lower bound (thin line), i.e., $\sum_{e=E+1}^{\infty} \lambda_e / \sum_{e=1}^{\infty} \lambda_e$ as illustrated in Theorem 1.

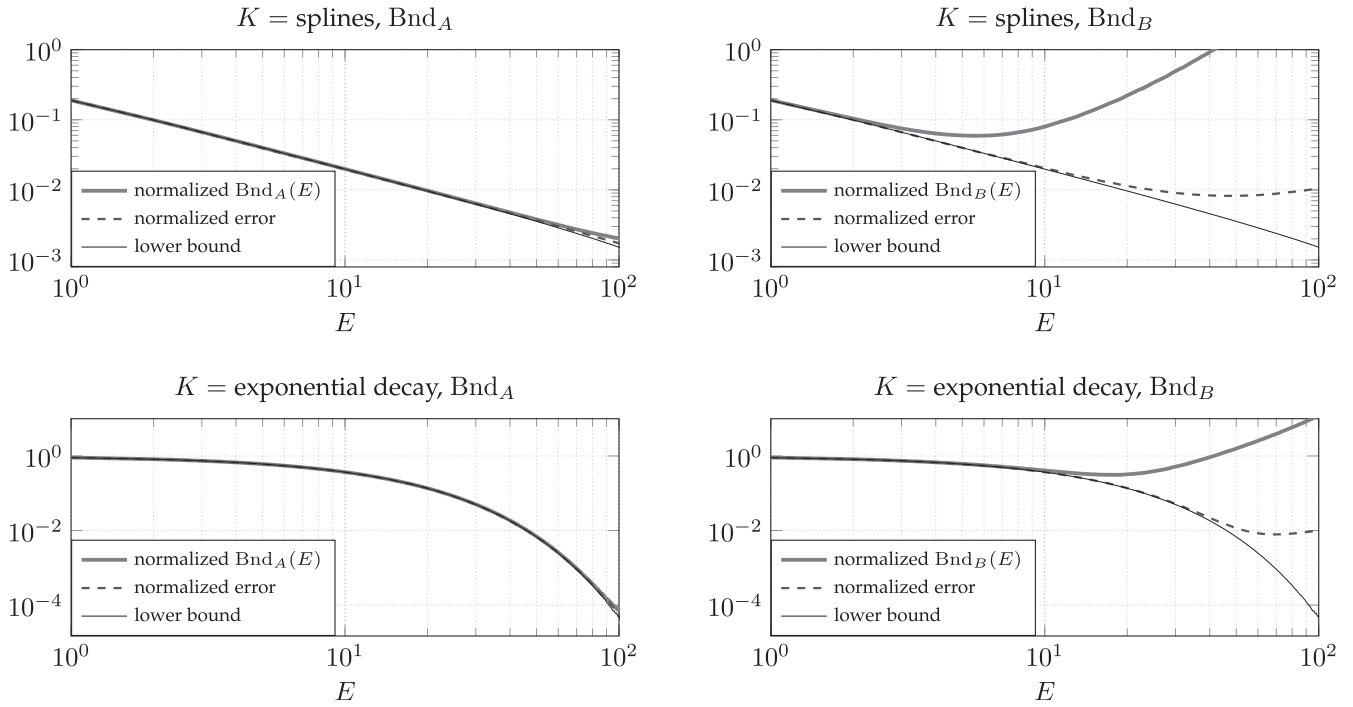


Fig. 1. Bnd_A and Bnd_B (normalized by the a priori function variance) as a function of E , with $\alpha = 0.05$, $M = 10000$ and for different eigenvalues decay rates.

As for Bnd_A , it is interesting to notice that just 20 eigenfunctions are needed to obtain a high estimation accuracy in both the cases. In addition, the curve is very close to the true error profile (which in turn is close to the lower bound) and is monotonically decreasing. Indeed, as discussed in the proof of Theorem 5 contained in the next subsection, when one adopts \hat{f}_A one should set E as large as possible (compatibly with communication capabilities) since, at the limit, convergence to the minimum variance estimator holds.

The profile of Bnd_B is instead different and exhibit a clear minimum at $E = 7$. The reason is that \hat{f}_B relies on the asymptotic matrix approximation (16). The bound Bnd_B then points out that if E is too large then the quality of this approximation can worsen, hence leading to an increment of the corresponding MSE. One can see that also the true error profile is not monotonically decreasing (indeed, we will see in the next subsection that for M fixed and E going to infinity \hat{f}_B is not guaranteed to converge to the minimum variance estimator). Note that, in this case, Bnd_B is close to truth only for low values of E and that the Monte Carlo analysis suggests the best E to be around 50. Overall, this indicates that the eigenfunctions number has to be seen as an important design parameter for \hat{f}_B to optimize the performance. This point will be the focus of Section 5.

Finally, the two bottom panels of Fig. 1 display the same bounds except that the kernel eigenvalues now decay exponentially to zero as $\lambda_e = \exp(-0.1e)$. Exponentially decaying eigenvalues are typical for Gaussian kernels, and therefore of practical relevance. The shapes of the curves change but the same comments hold true.

4.4 Asymptotic Behaviors of the Estimators and of the Bounds

Now, we start investigating the asymptotic properties of our estimators considering a situation where their dimension E

is fixed while the number of measurements M grows to infinity. The next result then shows that \hat{f}_A and \hat{f}_B asymptotically reach the lower bound (17).

Theorem 3. *Given the assumptions in Section 2.1 and Assumption 1,*

$$\lim_{M \rightarrow +\infty} \text{Err}_A = \sum_{e=E+1}^{+\infty} \lambda_e \quad \text{in probability}$$

$$\lim_{M \rightarrow +\infty} \text{Err}_B = \sum_{e=E+1}^{+\infty} \lambda_e \quad \text{in probability.}$$

We now discuss the statistical consistency of our estimators. In this case, the conditions under which \hat{f}_A and \hat{f}_B converge to f as both E and M grow to infinity are different, as illustrated in the following two results.

Theorem 4. *Given the assumptions in Section 2.1 and Assumption 1,*

$$\lim_{M \rightarrow +\infty} \lim_{E \rightarrow +\infty} \text{Err}_A = 0 \quad \text{in probability.}$$

Theorem 5. *Let $E = E(M)$ such that*

$$E(M) \log E^\delta(M) \leq M^\delta, \quad \lim_{M \rightarrow +\infty} E(M) = +\infty$$

for some $\delta \in (0, 1)$. Given the assumptions in Section 2.1 and Assumption 1, then

$$\lim_{\substack{M \rightarrow +\infty \\ E = E(M)}} \text{Err}_B = 0 \quad \text{in probability}$$

Remark 1. The sufficient condition required in the theorem in terms of the growth rate of $E(M)$ as a function of M is

tight according to the Chernoff's bound. In fact, our requirement is that M grows up a bit more slowly w.r.t. the relationship $E \log E = M$. Now, assume instead that $E \log E = M$, i.e., $\delta = 1$ and fix any rule such that $\varepsilon \rightarrow 1$ and $\alpha \rightarrow 0$. Recall that

$$1 - \varepsilon + \varepsilon \log(\varepsilon) \geq \frac{Ek}{M} \log\left(\frac{2E}{\alpha}\right)$$

must be satisfied. Asymptotically, the lhs tends to 0^+ while the second term becomes $k - \frac{Ek}{M} \log(\alpha/2)$ and is larger than k when α is sufficiently close to zero. One would thus need $0 \geq k$ but this is not possible. Also note that the previous theorem implies that any sublinear power growth of $E(M) = M^a$, for any $a \in (0, 1)$, satisfies the consistency condition, which can be readily verified by choosing $\delta = \frac{1+a}{2}$.

The consistency properties of \hat{f}_A and \hat{f}_B are thus remarkably different. For what regards \hat{f}_A , as M goes to infinity its consistency is guaranteed without any control on the growth rate of the dimension E . Indeed, as E increases such estimator can approximate arbitrarily well the optimal \hat{f}_{MAP} . This agrees with what already discussed in the previous subsection: when using \hat{f}_A it is convenient for the network to use a dimension E as large as possible, just compatible with its communication constraints. Differently, the estimator \hat{f}_B is instead consistent only if M augments sufficiently faster than E .

5 DISTRIBUTED TUNING OF THE REGULARIZATION PARAMETER

The statistical bounds obtained in the previous section quantify the performance of \hat{f}_A and \hat{f}_B assuming that the prior function model is correct. Beyond their theoretical interest, in real applications these bounds can give useful guidelines to select the amount of information that agents need to exchange. However, the covariance K is often defined only except for a scalar factor γ . In addition, the prior is never perfect and the tuning of γ could also hinder possible undermodeling. So, in place of (3), in practical applications it is beneficial to consider

$$f \sim \mathcal{N}(0, \gamma^{-1}K)$$

with γ to be estimated from the observed noisy outputs and related input locations. Furthermore, when \hat{f}_B is considered, it has been shown that also the parameter E plays an important role since, for a fixed number of samples M , its performance degrades if E is too small or too large. Hence, it could be desirable to adjust also the number of eigenfunctions forming the estimate after seeing the data.

In the following we will follow the SURE approach for tuning the free parameters. Although alternative approaches are possible, such as cross validation and marginal likelihood optimization, we will see that SURE has the advantage to require less communication and computation processing, and also to be suitable for distributed implementations. We start by reporting a result obtained through a simple generalization of the arguments in [44] [Section 7.4].

Theorem 6. Let η be a deterministic unknown parameter vector. Assume that the measurements model is

$$z = \eta + e$$

and consider also future measurements

$$z^* = \eta + e^*,$$

where the noises e and e^* are uncorrelated, zero mean with covariance Σ . Then, given the linear estimator $\hat{z} = Sz$, an unbiased estimator of the risk $\mathbb{E}[\|z^* - \hat{z}\|^2]$ is given by:

$$\|z - \hat{z}\|^2 + 2\text{tr}(S\Sigma). \quad (24)$$

The quantity $\text{tr}(S\Sigma)$ entering the second part of the objective (24) is connected to the concept of *equivalent degrees of freedom* [52], [53].

In what follows, we assume that γ is unknown but belongs to the finite set Γ which is known in advance to the network. In addition, let us assume that the estimation step has been performed adopting a certain value E . Hence, if \hat{f}_A has been used, each agent has stored $\frac{G^T G}{M}$ and $\frac{G^T}{M} y$ so that, letting

$$H_A(\gamma) := \left(\frac{G^T G}{M} + \frac{\gamma \sigma_v^2}{M} \Lambda_E^{-1} \right)^{-1} \frac{G^T}{M},$$

it can compute $H_A y$ for any $\gamma \in \Gamma$.

If \hat{f}_B has been adopted, then also the optimal number of eigenbases E' has to be found within the set $E' \in \Omega$. In this case, each agent knows only $\frac{G^T}{M} y$ and, letting

$$H_B(\gamma, E') := \mathcal{I}_{E'} \left(I + \frac{\gamma \sigma_v^2}{M} \Lambda_{E'}^{-1} \right)^{-1} \frac{G^T}{M}, \quad (25)$$

where

$$\mathcal{I}_{E'} := \begin{bmatrix} I_{E'} & \\ & \mathbf{0}_{E-E'} \end{bmatrix}, \quad (26)$$

it can compute $H_B y$ for any $\gamma \in \Gamma$ and integer $E' \in \Omega$.

5.1 Distributed SURE for \hat{f}_A : Tuning of γ

The first strategy is suited for \hat{f}_A . Surprisingly, we will see that the tuning of γ can be performed by the network using only local operations, without the need of performing any additional consensus operation. Now, let us reconsider our measurements model

$$y = Ga + Zb + v,$$

where a is E -dimensional. Hereby, we break away from the assumptions on prior correctness by thinking of $Ga + Zb$ as a deterministic vector. It thus corresponds to the deterministic function f sampled on the realizations of the input locations.

We then create a (projected) measurement model via pre-multiplication by G^T/M , i.e.,

$$\underbrace{\frac{G^T y}{M}}_z = \underbrace{\frac{G^T G}{M} a + \frac{G^T Z}{M} b}_\eta + \underbrace{\frac{G^T v}{M}}_e, \quad (27)$$

where the correspondences with the key quantities defining the risk estimator (24) have been pointed out. From such definitions, we also obtain $\hat{z} = \frac{G^T G}{M} H_A y = Sz$ where

$$S := \frac{G^T G}{M} \left(\frac{G^T G}{M} + \frac{\gamma \sigma_v^2}{M} \Lambda_E^{-1} \right)^{-1}$$

and

$$\Sigma = \sigma_v^2 \frac{G^T G}{M^2}.$$

Recall that the matrix $V = \frac{G^T G}{M} = \frac{1}{M} \sum_{m=1}^M G_m^T G_m$ and the vector $z = \frac{1}{M} \sum_{m=1}^M G_m^T y_m$ have been already computed by each agent via a distributed consensus algorithm [29] to implement \hat{f}_A . Then, since the network cardinality M is known, each agent can tune γ by optimizing the SURE score (24) connected with the prediction risk on the future data $z^* = \frac{G^T y^*}{M}$, i.e.,

$$\hat{\gamma}_A = \arg \min_{\gamma \in \Gamma} J_A(\gamma) \quad (28)$$

with

$$\begin{aligned} J_A(\gamma) &:= \|(I - S)z\|^2 + 2\text{tr}(S\Sigma) \\ &= \left\| \frac{\gamma \sigma_v^2}{M} \left(V \Lambda_E + \frac{\gamma \sigma_v^2}{M} I \right)^{-1} z \right\|^2 + \\ &\quad + \frac{2\sigma_v^2}{M} \text{tr} \left(V^2 \left(V + \frac{\gamma \sigma_v^2}{M} \Lambda_E^{-1} \right)^{-1} \right). \end{aligned}$$

To understand the rationale underlying this strategy we have just to consider that the novel process (27) is formed by E measurements, each corresponding to the projection of the original ones on the space of the sampled eigenfunctions $[\phi_e(x_1) \cdots \phi_e(x_M)]$. For large M , the quantity $G^T Z b$ vanishes so that $\eta \approx a$. This means that the SURE score becomes an unbiased estimator of those signal components which are expected to capture the most part of the energy.

Remark 2. Based on the previous analysis, it is straightforward to observe that the SURE strategy described above is not suited for \hat{f}_B . In fact, it requires each agent to know $\frac{G^T G}{M}$. But if this quantity were known, each agent could implement \hat{f}_A , an estimator that has more favorable features than \hat{f}_B .

5.2 Distributed SURE for \hat{f}_B : Tuning of E and γ

The second strategy is designed for \hat{f}_B . It tunes $\gamma \in \Gamma$ and $E' \in \Omega$ just using an additional average consensus on a vector of size $E \dim(\Omega) \dim(\Gamma)$. Our starting point is still (27), i.e., the E -dimensional projected measurement space, where $z = \frac{G^T y}{M} = \frac{1}{M} \sum_{m=1}^M G_m^T y_m \in \mathbb{R}^E$ has been computed to implement \hat{f}_B via a standard distributed consensus algorithm and is therefore known to each agent. Let us define $\hat{a}(\gamma, E') = H_B(\gamma, E') y \in \mathbb{R}^{E'}$. Clearly $\hat{a}(\gamma, E')$ for $E' < E$ is simply the truncated version of $\hat{a}(\gamma, E)$ where the last $E - E'$ components are set to zero. Moreover, the vectors $\hat{a}(\gamma, E')$ can be independently computed by each agent for each value of $\gamma \in \Gamma, E' \in \Omega$ once z is available. The output prediction can be written as

$$\hat{z}(\gamma, E') = \frac{G^T G}{M} \hat{a}(\gamma, E') = \frac{1}{M} \sum_{m=1}^M G_m^T G_m \hat{a}(\gamma, E') \in \mathbb{R}^{E'}.$$

Hence, each agent can compute the vectors $\hat{z}(\gamma, E')$ for each $\gamma \in \Gamma$ and $E' \in \Omega$ by running an additional consensus of

size $O(\dim(\Omega) \dim(\Gamma) E)$. As so, the first part of the SURE score $\|z - \hat{z}(\gamma, E')\|^2$ can be readily computed by each agent.

As for the second part of SURE related to the equivalent degrees of freedom, we need to compute $\text{tr}(S(\gamma, E')\Sigma)$ where

$$S(\gamma, E') = \frac{G^T G}{M} \mathcal{I}_{E'} \left(I + \frac{\gamma \sigma_v^2}{M} \Lambda_E^{-1} \right)^{-1}, \quad \Sigma = \frac{\sigma_v^2}{M} \frac{G^T G}{M}.$$

Obviously, this would not make too much sense in the context of \hat{f}_B since the computation of $V = \frac{G^T G}{M}$ would allow us to compute \hat{f}_A which has better performance anyways. Therefore we will approximate such matrix V (similarly to what we did to obtain \hat{f}_B) by replacing it with an identity matrix. This corresponds to use a sort of *expected equivalent degrees of freedom*:

$$\begin{aligned} \text{tr}(S(\gamma, E')\Sigma) &\approx \frac{\sigma_v^2}{M} \text{tr} \left(\mathcal{I}_{E'} \left(I + \frac{\gamma \sigma_v^2}{M} \Lambda_E^{-1} \right)^{-1} \right) \\ &= \frac{\sigma_v^2}{M} \sum_{e=1}^{E'} \frac{\lambda_e}{\lambda_e + \gamma \sigma_v^2 / M}. \end{aligned}$$

The optimal tuning of the parameter is then obtained as

$$\left(\hat{\gamma}_B, \hat{E}_B \right) = \underset{\gamma \in \Gamma, E' \in \Omega}{\text{argmin}} J_B(\gamma, E') \quad (29)$$

with

$$J_B(\gamma, E') := \|z - \hat{z}(\gamma, E')\|^2 + 2 \frac{\sigma_v^2}{M} \sum_{e=1}^{E'} \frac{\lambda_e}{\lambda_e + \gamma \sigma_v^2 / M}.$$

Note that this strategy for tuning \hat{f}_B is more efficient from a communication and computational point of view than \hat{f}_A only if $\dim(\Omega) \dim(\Gamma) < E$.

5.3 Practical Implementation Issues

We now illustrate how to implement the proposed distributed estimators, also in connection with the properties of the SURE tuning strategies described above. We discuss first the use and the derivation of the KL expansion and then how f can be estimated in a distributed way also adopting basis functions different from the kernel eigenfunctions. All the code developed for implementing the algorithms below is publicly available in the repository github.com/damianovar/Gaussian-regression-via-finite-dimensional-approximations.

5.3.1 Computing the KL Expansions

Assume that the prior on f is correct and that the input locations distribution μ is known. Then, according to Theorem 5, at least for large data set size M , the use of the eigenfunctions in (5) is statistically optimal. Obtaining the kernel expansion in closed form is in general difficult but important exceptions are the popular spline and Gaussian kernel. In particular, for uniform μ the expansion of the linear and cubic smoothing spline kernel is reported in [49]. For Gaussian μ on the real line, the Gaussian kernel expansion is given via Hermite polynomials, as reported in [27] [Section 4]. Such result then immediately generalizes to multi-dimensional domains: if $\mu(\cdot)$ and $K(\cdot, \cdot)$ are tensor products of one-dimensional distributions and

kernels, respectively, the expansion involves tensor products of the one-dimensional eigenfunctions.

Assume then that the kernel expansion is not available in closed form. It is worth pointing out that in many relevant distributed problems the dimension of the function domain \mathcal{X} is limited to 2 or 3, and this makes the numerical determination of the eigenfunctions and eigenvalues viable. More specifically, let $\{\tilde{x}_e\}_{e=1}^q$ be independent samples from μ , and let \mathbf{K} be the $q \times q$ kernel matrix whose (i, j) -entry is

$$[\mathbf{K}]_{ij} = K(\tilde{x}_i, \tilde{x}_j), \quad i = 1, \dots, q, \quad j = 1, \dots, q. \quad (30)$$

Then, according to [50] [Lemma 9 and Corollary 10], the eigenvalues and (normalized) eigenvectors from the Singular Values Decomposition (SVD) of \mathbf{K} converge to the eigenvalues and eigenfunctions of $K(\cdot, \cdot)$ as $q \rightarrow +\infty$. Hence, the agents can be equipped with arbitrarily accurate approximations of the KL expansion.

5.3.2 Generic Basis Functions: Kernel Sections

As discussed above, in some circumstances the kernel eigenfunctions could be not available in closed form, or have a complex functional form that makes storing them in the agents' memory unpractical. In such cases, one would rather use basis functions which admit simple closed-form expressions, possibly also non orthonormal. Even if the bounds developed in Section 4 cannot be used anymore, we will see that the SURE strategies for hyperparameters tuning generalize well also to this situation.

We limit our discussion to the use of the kernel sections as basis (an important case also in view of their connections with the representer theorem [54], [55]). This basis is associated to a set $\{\tilde{x}_e\}_{e=1}^E$ of input locations¹ which could be drawn from μ or selected in a deterministic way to cover sufficiently well \mathcal{X} . We then define our E basis functions as

$$\phi_1(\cdot) = K(\tilde{x}_1, \cdot) \quad \dots \quad \phi_E(\cdot) = K(\tilde{x}_E, \cdot).$$

Using the kernel sections in the decomposition (7), we can think of f_a as

$$f_a(x) = \sum_{e=1}^E a_e K(\tilde{x}_e, x),$$

where the vector $\mathbf{a} := [a_1, \dots, a_E]^T$ is now zero-mean Gaussian with covariance proportional to the inverse of the kernel matrix

$$[\mathbf{K}]_{ij} = K(\tilde{x}_i, \tilde{x}_j), \quad i = 1, \dots, E, \quad j = 1, \dots, E$$

i.e.,

$$\mathbf{a} \sim \mathcal{N}(0, \gamma^{-1} \mathbf{K}^{-1}).$$

In fact, if the prior were correct, this would indeed correspond to see f_a sampled on $\{\tilde{x}_e\}_{e=1}^E$ as zero-mean Gaussian with covariance $\gamma^{-1} \mathbf{K}^{-1}$.

Since the kernel sections are generally not orthonormal w.r.t. μ , even if $M \rightarrow \infty$ the projected measurements $\frac{G^T \mathbf{y}}{M}$ do not converge to the expansion coefficients a_e . However, these

1. As in (30), the set $\{\tilde{x}_e\}_{e=1}^q$ is available a priori and has not to be confused with the input locations $\{x_m\}_{m=1}^M$ then visited by the agents.

can be still used to tune the regularization parameters. In particular, for what concerns \hat{f}_A , the distributed SURE estimator introduced in Section 5.1 can estimate f and γ with a single consensus just replacing Λ_E^{-1} with \mathbf{K} . Thus, this estimator does not even need the knowledge of μ and the agents can implement it once they know the function $K(\cdot, \cdot)$ and the expansion grid $\{\tilde{x}_e\}_{e=1}^E$. The estimator \hat{f}_A is thus given by:

$$\hat{\mathbf{a}}(\gamma) := \left(\frac{G^T G}{M} + \frac{\sigma_v^2}{M} \mathbf{K} \right)^{-1} \frac{G^T \mathbf{y}}{M}$$

$$S(\gamma) := \frac{G^T G}{M} \left(\frac{G^T G}{M} + \frac{\gamma \sigma_v^2}{M} \mathbf{K} \right)^{-1}.$$

Consider now the implementation of \hat{f}_B through the kernel sections with the estimator defined by the set of potential $E' \in \Omega$. In particular, for the sake of simplicity, assume that each E' is associated to the kernel sections induced by the first E' input locations in the (ordered) set $\{\tilde{x}_e\}_{e=1}^E$. Given a generic matrix A , the submatrix obtained by retaining its first E' rows and columns is denoted by $[A]_{E'}$. Assume moreover that the same notation applies to vectors to retain only their first E' elements. Then, the same SURE strategy developed in Section 5.2 can be adopted by setting

$$\hat{\mathbf{a}}(\gamma, E') = \begin{bmatrix} I_{E'} \\ \mathbf{0} \end{bmatrix} \left(\left[\mathbb{E} \frac{G^T G}{M} \right]_{E'} + \frac{\gamma \sigma_v^2}{M} [\mathbf{K}]_{E'} \right)^{-1} \begin{bmatrix} G^T \mathbf{y} \\ M \end{bmatrix}_{E'}$$

and

$$S(\gamma, E') = \left[\mathbb{E} \frac{G^T G}{M} \right]_{E'} \begin{bmatrix} I_{E'} \\ \mathbf{0} \end{bmatrix} \left(\left[\mathbb{E} \frac{G^T G}{M} \right]_{E'} + \frac{\gamma \sigma_v^2}{M} [\mathbf{K}]_{E'} \right)^{-1},$$

where, instead of using $\mathcal{I}_{E'}$ defined in (26), we use $I_{E'} \in \mathbb{R}^{(E-E') \times E'}$ to account for the non-diagonal nature of the matrices now at stake, with the trace of S given by the sum of the its (i, i) entries with $i = 1, \dots, E'$, and where $\mathbb{E} \left[\frac{G^T G}{M} \right]$ substitutes I in (25), since the kernel sections are generally not orthonormal. The exact expectation of $\mathbb{E} \left[\frac{G^T G}{M} \right]$ can be explicitly computed in some special cases as in the example below, or can be approximated via its sampled version, i.e., $\mathbb{E} \left[\frac{G^T G}{M} \right] = \mathbb{E} \left[\sum_{m=1}^M \frac{G_m^T G_m}{M} \right] \approx \frac{1}{E} \sum_{e=1}^E \frac{\tilde{G}_e^T \tilde{G}_e}{E}$ where \tilde{G}_e are computed on the input locations $\{\tilde{x}_e\}_{e=1}^E$ that shall be used for computing this empirical expectation, not to be confused with the set of a-posteriori input locations $\{x_m\}_{m=1}^M$ used in the actual coefficients estimation step.

Example 1. An interesting case, relevant for many applications, arises when one wants to use the Gaussian kernel and its kernel sections as basis with μ a mixture of Gaussians. In this case $\mathbb{E} \left[\frac{G^T G}{M} \right]$ can be obtained in closed form. In fact, consider first a scalar scenario, i.e., $x \in \mathbb{R}$, with a mixture of Gaussians made of a single component:

$$K(x, x') = \exp\left(-\frac{(x - x')^2}{\eta}\right), \quad \mu \sim \mathcal{N}(\mu_0, a^2).$$

After simple computations, one finds that the (e, e') -entry of $\mathbb{E} \left[\frac{G^T G}{M} \right]$ is

$$\int_{-\infty}^{+\infty} \frac{\exp\left(-\frac{(x-x_e)^2}{\eta} - \frac{(x-x_{e'})^2}{\eta} - \frac{(x-\mu_0)^2}{2a^2}\right)}{\sqrt{2\pi a}} dx = \frac{\sqrt{\eta}}{\sqrt{\eta + 4a^2}} \exp\left(-\frac{\star}{\eta^2 + 4\eta a^2}\right)$$

with $\star = \eta(x_e^2 - 2\mu_0 x_e + x_{e'}^2 - 2\mu_0 x_{e'} + 2\mu^2) + 2a^2(x_e - x_{e'})^2$.

In the multivariate case, assume that $K(x, x') = e^{-\frac{\|x-x'\|^2}{\eta}}$ while μ is given by tensor products of one-dimensional Gaussian densities. Then, the result is still available in closed form: $\mathbb{E}\left[\frac{G^T G}{M}\right]$ corresponds to convex combinations of Hadamard products of the matrices obtained in the scalar case.

5.3.3 Generic Basis Functions: Nyström Method

The analysis provided in the previous section can also be applied to the popular Nyström method [56], [57]. The idea is to find a basis for f of dimension E which has almost the same performance of the basis composed of $q \gg E$ kernel sections with q an arbitrary number as in Section 5.3.1. More specifically, let $\{\tilde{x}_n\}_{n=1}^q$ be q input location defined a-priori² from μ , and consider both the corresponding kernel matrix $K \in \mathbb{R}^{q \times q}$ defined in (30) and its SVD decomposition $K = VD_q V^T$ with $V := [v_1, \dots, v_q]$ the orthonormal eigenvectors of K and D_q the diagonal matrix formed by the corresponding eigenvalues of K sorted in non-increasing order. If $V_E := [v_1, \dots, v_E]$ and $D_E \in \mathbb{R}^{E \times E}$ is the diagonal matrix with the first E sorted eigenvalues of K , then $K_E = V_E D_E V_E^T$ is the best rank- E approximation of K . The a priori basis

$$\phi_e(x) := \sum_{n=1}^q v_e(n) K(\tilde{x}_n, x), \quad e = 1, \dots, E$$

with $v_e(n)$ the n -th element of the vector v_e , can then be used to define

$$f_a(x) = \sum_{e=1}^E a_e \phi_e(x)$$

where $\mathbf{a} := [a_1, \dots, a_E]^T$ is a zero-mean Gaussian vector with

$$\mathbf{a} \sim \mathcal{N}\left(0, \gamma^{-1} (V_E^T K V_E)^{-1}\right) = \mathcal{N}\left(0, \gamma^{-1} D_E^{-1}\right).$$

We can then use once again (12) to build G using the ϕ_e 's above, and exploit the same strategies considered in Section 5.3.2 just replacing K with D_E .

5.4 Numerical Study on Synthetic Data

Let us consider the same data generators based on the spline and the exponentially decaying kernels described in Section 4.3. The unknown function has to be reconstructed from $M = 10000$ measurements by \hat{f}_A and \hat{f}_B . The errors are still the MSEs defined in (20) and (21) normalized by the prior variance (the same definition was used to build Fig. 1). The

2. The Nyström method closely resembles the eigenfunctions/eigenvalues numerical computation method presented in Section 5.3.1, the difference being that in Nyström the parameter q is in general not extremely large and the input locations $\{\tilde{x}_n\}_{n=1}^q$ are not generated by μ but are randomly extracted from the training set.

difference however is that our estimators now depend on unknown hyperparameters that need to be inferred from data. More specifically, when \hat{f}_A is adopted we fix $E = 400$ and the regularization parameter is searched over a grid Γ containing 50 logarithmically spaced values between 10^{-3} and 10^3 . When using \hat{f}_B the grid Γ contains only the three values $\{10^{-3}, 0, 10^3\}$ while E is estimated from data over $\Omega = \{1, 5, 10, 20, 50, 100, 200, 300, 400\}$. We still consider a Monte Carlo study of 1000 runs where at any run independent realizations of f , of the M input locations and of the measurement noises are generated. Hyperparameters tuning is then performed by:

- “ \hat{f}_A + oracle” and “ \hat{f}_B + oracle”, where “oracle” indicates that these approaches know at any run the realization of f (which is the object to estimate) and select exactly those hyperparameters that minimize the MSE achievable by those two estimators. For instance, assume that $f = \sum_{e=1}^{\infty} a_e \phi_e$ is the realization of the function at a certain run. Let also $\hat{a}(\gamma, E')$ denote the vector with the estimates of the first E' coefficients a_e returned by \hat{f}_B . Then \hat{f}_B + oracle determines the hyperparameters as

$$\left(\hat{\gamma}, \hat{E}\right) := \arg \min_{\gamma \in \Gamma, E' \in \Omega} \sum_{e=1}^{\infty} (a_e - \hat{a}_e(\gamma, E'))^2,$$

where $\hat{a}_e(\gamma, E') := 0$ for $e > E'$. Thus, this estimator is not implementable in practice and provides the lower bound on the MSEs (20) and (21) achievable by the two estimators;

- “ \hat{f}_A + SURE” and “ \hat{f}_B + SURE”, where the hyperparameters tuning step is performed following the SURE approaches described in the previous subsection. Recall that “ \hat{f}_A + SURE” requires only a single consensus on a vector of size $O(E^2)$ to obtain simultaneously both the hyperparameters and function estimates, while “ \hat{f}_B + SURE” requires two consensus operations of size $O(E)$.

Fig. 2 compares with a scatter-plot the various (normalized) MSEs obtained by the oracle- and SURE-based approaches as a function of the Monte Carlo run. Remarkably, SURE's performance (dashed lines) is very close to that of the oracle (solid lines). When using “ \hat{f}_B + SURE” (right panels) the curves are in practice indistinguishable.

The set of four Monte Carlo experiments have been also repeated adopting different data set sizes M . To synthesize SURE's performance with an index function only of M , let $S_p \in [0, 1]$ denote the ratio between the mean of the 400 errors obtained by the oracle and the SURE strategies respectively for a certain value of M . Note that a value of $S_p = 1$ indicates that SURE is performing as well as the oracle and that, for $M = 10000$, S_p becomes the distillate of Fig. 2. Table 1 reports S_p for $M = 100, 1000, 10000$: one can see that the proposed hyperparameter estimation procedure behaves very nicely.

5.5 Numerical Study on Field Data – Colorado Rain

Let us now consider the reconstruction of monthly precipitations using data collected in Colorado in the years 1995-1997 [58]. Many alternative solutions are available in the context of weather forecasts, but they are limited to centralized

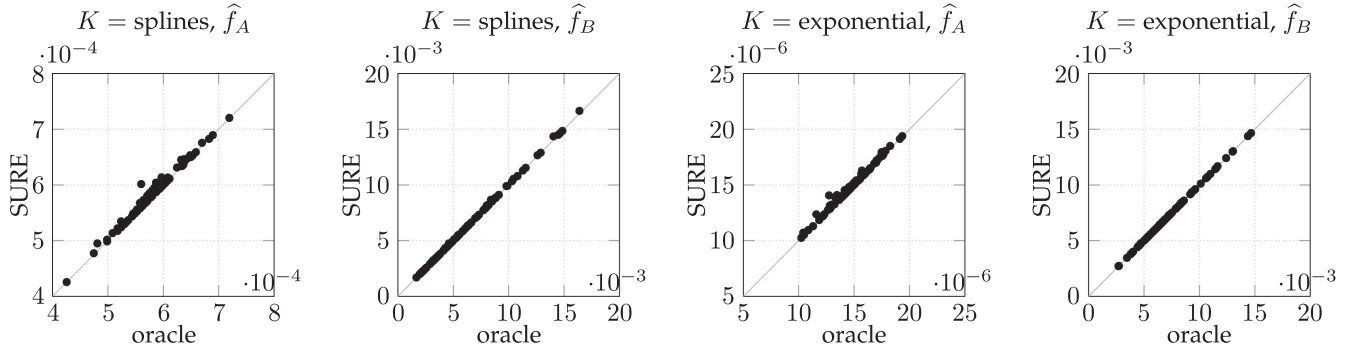


Fig. 2. Comparison of the MSE indexes obtained by the SURE- and oracle-based strategies. Each circle corresponds to the result of a certain Monte Carlo run (the x -axis being associated to oracle-based estimators, and the y -axis to SURE-based ones). The fact that the circles groups are close to the bisector of the first quadrant indicates that the performance of SURE is almost equivalent to that of the oracle.

solutions, such as [5], [6], for example, thus not suitable in our distributed framework. Measurements y_m are a series of monthly average precipitations measured at 367 stations within the rectangular longitude/latitude region $[-109.5, -101] \times [36.5, 41.5]$ remapped for convenience into the unitary square so that $x_m \in [0, 1]^2$ for every m .

We test the SURE strategies (28) and (29). When using \hat{f}_A we set $E = 20$ and Γ to the grid containing 50 values logarithmically spaced between 10^{-5} and 10^5 . When adopting \hat{f}_B we use $\Gamma = 0$ and $\Omega = \{2, 4, \dots, 20\}$, i.e., consider only E' as a regularization parameter. In both cases, we consider the Gaussian kernel

$$K(x, x') = \exp(-10\|x - x'\|^2).$$

The eigenfunctions are computed by assuming that the input locations are not known a-priori and extracted uniformly from the monitored region, i.e., $\mu(x)$ is a uniform distribution. We design a Monte Carlo study of 1000 runs where, at any run, we select randomly two months within the 1995-1997 dataset obtaining three different sets. The first one is a training set $\mathcal{D}_{\text{train}}$ of average precipitations obtained by selecting randomly and uniformly 2/3 of the measurements from the first selected month. The second is a test set $\mathcal{D}_{\text{test}}$ corresponding to the remaining 1/3 measurements from the first selected month. The last one is $\mathcal{D}_{\sigma_v^2}$ and contains measurements in the second selected month which are used to estimate the noise variance via least squares based on E eigenfunctions. This corresponds to using \hat{f}_A with $\gamma = 0$ obtaining as estimate of the noise variance

$$\hat{\sigma}_v^2 = \frac{1}{\dim(\mathcal{D}_{\sigma_v^2}) - E} \sum_{m=1}^{\dim(\mathcal{D}_{\sigma_v^2})} (\hat{f}_A(x_m; 0) - y_m)^2,$$

where $\dim(\mathcal{D}_{\sigma_v^2})$ is the cardinality of $\mathcal{D}_{\sigma_v^2}$. Overall, this represents a situation where noise levels are determined by a centralized approach before running the estimators \hat{f}_A and \hat{f}_B .

TABLE 1
SURE's Performance Index S_p Summarizing Four Monte Carlo Studies as a Function of the Number M of Available Measurements

Data set size	M = 100	M = 1000	M = 10000
S_p	0.93	0.987	0.99

A value of S_p close to 1 indicates that SURE's performance is close to that of the oracle. For $M = 10000$, S_p represents a resume of the entire Fig. 2.

The following tuning strategies are used:

- (1) " $\hat{f}_A + \text{oracle}$ " and " $\hat{f}_B + \text{oracle}$ ", where "oracle" now indicates that these approaches can select those hyperparameters minimizing the following prediction errors on the test set

$$\text{RSS}_A(\gamma) := \frac{1}{\dim(\mathcal{D}_{\text{test}})} \sum_{m=1}^{\dim(\mathcal{D}_{\text{test}})} (\hat{f}_A(x_m; \gamma) - y_m)^2 \quad (31)$$

$$\text{RSS}_B(\gamma, E') := \frac{1}{\dim(\mathcal{D}_{\text{test}})} \sum_{m=1}^{\dim(\mathcal{D}_{\text{test}})} (\hat{f}_B(x_m; \gamma, E') - y_m)^2, \quad (32)$$

where (x_m, y_m) are all elements of $\mathcal{D}_{\text{test}}$. Note that RSS_A and RSS_B can be seen as approximations of the MSEs (20) and (21) and that the oracle provides a lower bound on their values;

- (2) " $\hat{f}_A + \text{SURE}$ " and " $\hat{f}_B + \text{SURE}$ ", where the hyperparameters tuning step is performed minimizing the estimated risks $J_A(\gamma)$ and $J_B(\gamma, E')$ defined in (28) and (29).

Fig. 3 compares with a scatter-plot the prediction errors (31) and (32) achieved by the estimators after the 1000 Monte Carlo runs. The situation is not dissimilar from the case of synthetic data: as for the estimators " $\hat{f}_A + \text{SURE}$ " and " $\hat{f}_B + \text{SURE}$ ", the performance of the SURE strategies is close to that of the oracles.




Fig. 3. Comparison of the Residual sum of squares (RSS) prediction error indexes obtained by the oracle- and SURE-based strategies. Each opaque circle corresponds to the result of one of the 1000 Monte Carlo runs. The closer the circles are to the bisector of the first quadrant performance means that the closer the performance of that SURE-based or Nyström-SURE estimator is to the ones of the oracle-based estimator.




Fig. 4. Comparison of the predictive performance of the estimators \hat{f}_A (left panel) and \hat{f}_B (right panel) over the test set in Fig. 5 for $\gamma \in \Gamma$ and $E' \in \Omega$ against the SURE scores $J_A(\gamma)$ and $J_B(\gamma)$ in the first Monte Carlo run.

Specifically considering the SURE-based strategies, Fig. 4 also compares the estimated risks $J_A(\gamma)$ and $J_B(\gamma, E')$ against the indexes $RSS_A(\gamma)$ and $RSS_B(\gamma, E')$ in (31) and (32) in the first Monte Carlo run. The curves show that hyperparameters values have a major effect on the estimation performance and that our SURE approach leads to a good regularization tuning. The related function estimates are visible in Fig. 5.

5.6 Numerical Study on Field Data – UCI Datasets

The second study is performed on two datasets from the public UCI repository, and is executed using the Nyström-based strategy described in Section 5.3.3 to compute the basis functions for the estimators using all the input locations that define the training set. Our purpose is here to compare the proposed SURE-based strategy for tuning the regularization parameters against an oracle that selects as the best regularization parameters that ones that give the best fit performance in the test set. As for the kernel, we consider a Gaussian kernel with an unitary variance (not accurately tuned, since the purpose of this section is more checking that the proposed SURE strategy chooses the regularization parameters accurately rather than actually maximizing the generalization capabilities of the regression

TABLE 2
Summary of the Performance of the Proposed Parameters Calibration Strategies Against Oracles for Different Publicly Available Datasets

	CCPP		CPU	
	\hat{f}_A	\hat{f}_B	\hat{f}_A	\hat{f}_B
fit oracle	99.3	73.3	75.7	76.1
fit SURE	99.3	73.3	71.2	66.5
γ oracle	10^{-3}	-	10^{-4}	-
γ SURE	10^{-3}	-	10^{-5}	-
E' oracle	-	2.00	-	10.00
E' SURE	-	2.00	-	27.00

“CCPP” indicates the UCI Combined Cycle Power Plant regression dataset, with 9568 instances and a 4-dimensional input domain \mathcal{X} . “CPU” indicates the UCI Computer hardware regression dataset, with 209 instances and a 6-dimensional input domain \mathcal{X} (corresponding to the quantitative features available in the dataset). For each dataset 1/3 of the data has been used for test purposes.

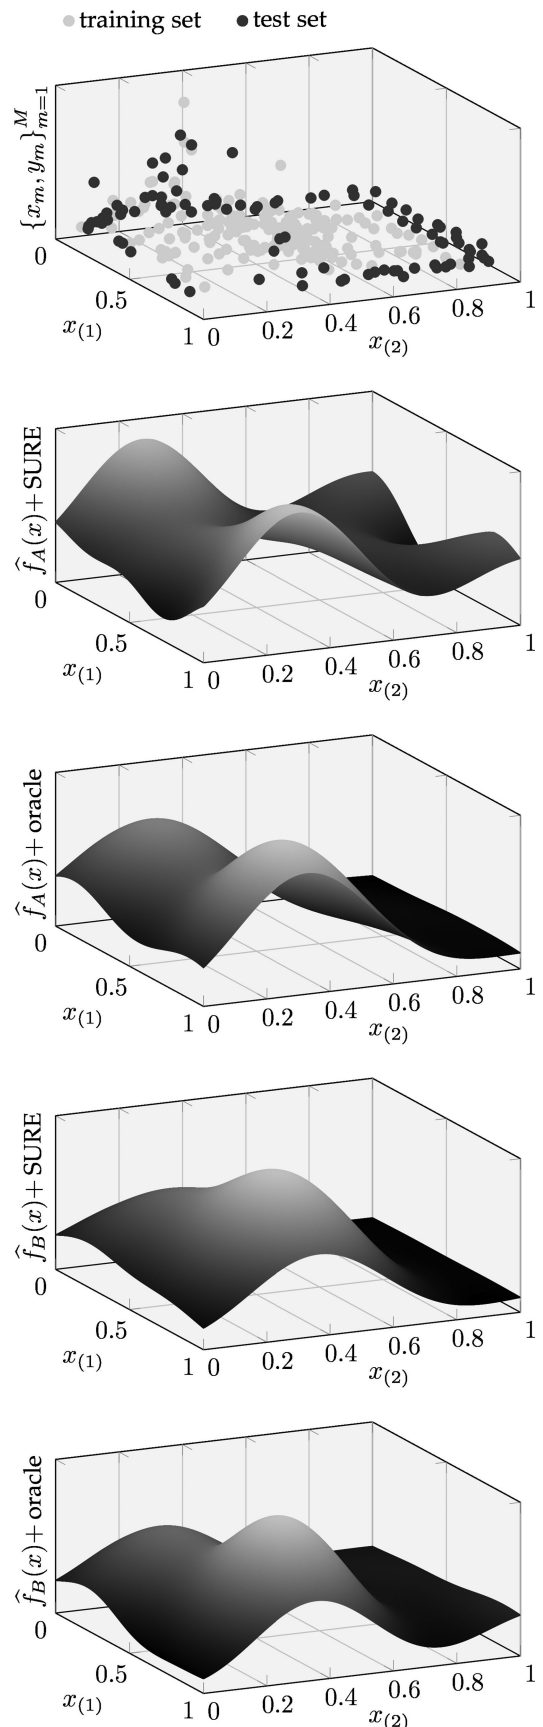


Fig. 5. Visualization of the training and test sets (top panel, respectively 173 and 87 samples), and of the estimates returned by “ $\hat{f}_A + \text{SURE}$ ”, “ $\hat{f}_A + \text{oracle}$ ”, “ $\hat{f}_B + \text{SURE}$ ”, and “ $\hat{f}_B + \text{oracle}$ ” in the first Monte Carlo run.

algorithms). As for the grid for tuning γ , we then consider the set $\Gamma = \{0, 10^{-5}, 10^{-4}, \dots, 10^2\}$; as for E' , we consider $\Omega = \{1, 2, \dots, 30\}$. Table 2 summarizes then the obtained results, and numerically confirms the efficacy of the proposed parameters tuning strategies.

6 CONCLUSIONS

Distributed function estimation is an important problem where agents with limited computational, data storage and communication capabilities collect noisy measurements and have to reconstruct an unknown map in a collaborative way. In this context, we have studied Gaussian regression providing rigorous statistical bounds on the performance of two distributed estimators, also characterizing their asymptotic behavior. On the practical side, our study indicates how the dimension E of the adopted estimator has to depend on the number of measurements M collected by the agents to guarantee the desired statistical performance. The analysis clarifies merits and limitations of the two approaches also in function of the different amount of information exchange required to the network (linear or quadratic in E). We have also introduced novel distributed strategies which learn from data possibly unknown hyperparameters entering the estimators, and that do not necessarily require solving potentially numerically intensive eigenfunctions-eigenvalues decompositions of kernel functions. For the first time, to our knowledge, this paper has shown how it is possible to estimate the regularization parameter and the unknown function via a single average consensus operation.

Overall, the theoretical achievements and the numerical strategies here described provide sound tools to reconstruct static functions in distributed networks. An important future direction is to extend all the analysis to an even more challenging situation where the unknown map may change in time and has to be tracked in an on-line manner.

ACKNOWLEDGMENTS

The research leading to these results has received funding from the Swedish research council Norrbottens Forskningsråd.

REFERENCES

- [1] T. Poggio and F. Girosi, "Networks for approximation and learning," *Proc. IEEE*, vol. 78, no. 9, pp. 1481–1497, Sep. 1990. [Online]. Available: <http://ieeexplore.ieee.org/document/58326/>
- [2] F. Cucker and S. Smale, "On the mathematical foundations of learning," *Bulletin Amer. Math. Soc.*, vol. 39, no. 101, pp. 1–49, 2001.
- [3] M. Lei, L. Shiyan, J. Chuanwen, L. Hongling, and Z. Yan, "A review on the forecasting of wind speed and generated power," *Renewable Sustain. Energy Rev.*, vol. 13, no. 4, pp. 915–920, 2009.
- [4] L. Parolini, B. Sinopoli, B. H. Krogh, and Z. Wang, "A cyber-physical systems approach to data center modeling and control for energy efficiency," *Proc. IEEE*, vol. 100, no. 1, pp. 254–268, Jan. 2012.
- [5] A. Gelfand and S. A. G. S. Banerjee, "Spatial process modelling for univariate and multivariate dynamic spatial data," *Environmetrics*, vol. 16, pp. 465–479, 2005.
- [6] A. Datta, S. Banerjee, A. O. Finley, and A. E. Gelfand, "Hierarchical nearest-neighbor gaussian process models for large geostatistical datasets," *J. Amer. Statist. Assoc.*, vol. 111, no. 514, pp. 800–812, 2016.
- [7] Yunfei Xu, Jongeun Choi, S. Dass, and T. Maiti, "Sequential Bayesian prediction and adaptive sampling algorithms for mobile sensor networks," *IEEE Trans. Autom. Control*, vol. 57, no. 8, pp. 2078–2084, Aug. 2012. [Online]. Available: <http://ieeexplore.ieee.org/document/6099560/>
- [8] B. Schölkopf and A. J. Smola, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*, Cambridge, MA, USA: MIT Press, 2001.
- [9] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA, USA: MIT Press, Apr. 2006, vol. 14, no. 2. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/15112367>
- [10] R. Tibshirani, "Regression shrinkage and selection via the lasso," *J. Royal Statist. Soc., Series B*, vol. 58, no. 1, pp. 267–288, 1996.
- [11] N. Meinshausen and B. Yu, "Lasso-type recovery of sparse representations for high-dimensional data," *Annals Statist.*, vol. 37, no. 1, pp. 246–270, 2009.
- [12] H. Zou and T. Hastie, "Regularization and variable selection via the elastic net," *J. Roy. Statist. Soc.: Series B*, vol. 67, no. 2, pp. 301–320, 2005.
- [13] D. Wipf and S. Nagarajan, "A new view of automatic relevance determination," in *Proc. NIPS*, 2007, pp. 1625–1632.
- [14] A. Aravkin, J. Burke, A. Chiuso, and G. Pillonetto, "Convex vs. nonconvex approaches for sparse estimation: GLASSO, multiple kernel learning, and HGLASSO," *J. Mach. Learn. Res.*, vol. 15, pp. 217–252, 2014.
- [15] S. Boyd, P. Neal, E. Chu, B. Peleato, and E. Jonathan, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, 2011.
- [16] J. Quiñero-Candela and C. E. Rasmussen, "A unifying view of sparse approximate Gaussian process regression," *J. Mach. Learn. Res.*, vol. 6, no. Dec, pp. 1939–1959, 2005.
- [17] E. Snelson and Z. Ghahramani, "Sparse Gaussian processes using pseudo-inputs," in *Proc. 18th Int. Conf. Neural Inf. Process. Syst.*, 2006, pp. 1257–1264. [Online]. Available: <http://citeseer.ist.psu.edu/viewdoc/summary?doi=10.1.1.104.5333>
- [18] M. Lázaro-Gredilla, J. Quiñero-Candela, C. E. Rasmussen, and A. R. Figueiras-Vidal, "Sparse spectrum Gaussian process regression," *J. Mach. Learn. Res.*, vol. 11, pp. 1865–1881, 2010.
- [19] S. Ambikasaran, D. Foreman-Mackey, L. Greengard, D. W. Hogg, and M. O'Neil, "Fast direct methods for Gaussian processes," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 2, pp. 252–265, Feb. 2016. [Online]. Available: <http://ieeexplore.ieee.org/document/7130620/>
- [20] F. R. Bach and M. I. Jordan, "Predictive low-rank decomposition for kernel methods," in *Proc. 22nd Int. Conf. Mach. Learn.*, 2005, pp. 33–40. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=1102351.1102356>
- [21] B. Kulis, M. Sustik, and I. Dhillon, "Learning low-rank kernel matrices," in *Proc. 23rd Int. Conf. Mach. Learn.*, 2006, pp. 505–512. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=1143844.1143908>
- [22] C. Williams and M. Seeger, "Using the Nyström method to speed up kernel machines," *Adv. Neural Inf. Process. Syst.*, vol. 13, pp. 682–688, 2001. [Online]. Available: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.18.7519>
- [23] K. Zhang and J. T. Kwok, "Clustered Nyström method for large scale manifold learning and dimension reduction," *IEEE Trans. Neural Netw.*, vol. 21, no. 10, pp. 1576–1587, Oct. 2010. [Online]. Available: <http://ieeexplore.ieee.org/document/5559473/>
- [24] A. J. Smola, A. J. Smola, and B. Schölkopf, "Sparse greedy matrix approximation for machine learning," in *Proc. 17th Int. Conf. Mach. Learn.*, 2000, pp. 911–918. [Online]. Available: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.43.3153>
- [25] B. C. Levy, "Karhunen-loève expansion of gaussian processes," in *Principles of Signal Detection and Parameter Estimation*. Berlin, Germany: Springer, 2008, pp. 1–47.
- [26] G. Ferrari-Trecate, C. K. I. Williams, and M. Opper, "Finite-dimensional approximation of Gaussian processes," in *Proc. Conf. Adv. Neural Inf. Process. Syst.*, 1999, pp. 218–224.
- [27] H. Zhu, C. K. I. Williams, R. J. Rohwer, and M. Morciniec, "Gaussian regression and optimal finite dimensional linear models," in *Neural Networks and Machine Learning*, C. M. Bishop, Ed. Berlin, Germany: Springer, 1998.
- [28] L. Xiao and S. Boyd, "Fast linear iterations for distributed averaging," *Syst. Control Lett.*, no. 53, pp. 65–78, 1998.
- [29] F. Garin and L. Schenato, "A survey on distributed estimation and control applications using linear consensus algorithms," in *Networked Control Systems*. Berlin, Germany: Springer, 2010, pp. 75–107.
- [30] N. Bof, R. Carli, and L. Schenato, "Average consensus with asynchronous updates and unreliable communication," in *Proc. IFAC Word Congress*, vol. 1, 2017, pp. 601–606.

- [31] J. B. Predd, S. R. Kulkarni, and H. V. Poor, "A collaborative training algorithm for distributed learning," *IEEE Trans. Inf. Theory*, vol. 55, no. 4, pp. 1856–1871, Apr. 2009. [Online]. Available: <http://ieeexplore.ieee.org/document/4802309/>
- [32] F. Perez-Cruz and S. R. Kulkarni, "Robust and low complexity distributed kernel least squares learning in sensor networks," *IEEE Signal Process. Lett.*, vol. 17, no. 4, pp. 355–358, Apr. 2010. [Online]. Available: <http://ieeexplore.ieee.org/document/5395679/>
- [33] P. Honeine, M. Essoloh, C. Richard, and H. Snoussi, "Distributed regression in sensor networks with a reduced-order Kernel model," in *Proc. IEEE Global Telecommun. Conf.*, 2008, pp. 1–5. [Online]. Available: <http://ieeexplore.ieee.org/document/4697804/>
- [34] S. Martinez, "Distributed interpolation schemes for field estimation by mobile sensor networks," *IEEE Trans. Control Syst. Technol.*, vol. 18, no. 2, pp. 491–500, Mar. 2010. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5229132>
- [35] Y. Xu, J. Choi, S. Dass, and T. Maiti, "Efficient Bayesian spatial prediction with mobile sensor networks using Gaussian Markov random fields," *Autom.*, vol. 49, no. 12, pp. 3520–3530, 2013.
- [36] J. Cortés, "Distributed Kriged Kalman filter for spatial estimation," *IEEE Trans. Autom. Control*, vol. 54, no. 12, pp. 2816–2827, Dec. 2009. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5325717
- [37] J. Choi, S. Oh, and R. Horowitz, "Distributed learning and cooperative control for multi-agent systems," *Autom.*, vol. 45, no. 12, pp. 2802–2814, Dec. 2009. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S000510980900435X>
- [38] D. Varagnolo, G. Pillonetto, and L. Schenato, "Distributed parametric and nonparametric regression with on-line performance bounds computation," *Autom.*, vol. 48, no. 10, pp. 2468–2481, 2012. [Online]. Available: http://automatica.dei.unipd.it/tl_files/utente/lucaschenato/Papers/Journal/AUTO11_Distributed_parametric_and_Nonparametric_regression.pdf
- [39] T. Zhang, "Learning bounds for kernel regression using effective data dimensionality," *Neural Comput.*, vol. 17, no. 9, pp. 2077–2098, Sep. 2005. [Online]. Available: <http://www.mitpressjournals.org/doi/10.1162/0899766054323008>
- [40] C. M. Stein, "Estimation of the mean of a multivariate normal distribution," *Annals Statist.*, vol. 9, no. 6, pp. 1135–1151, 1981. [Online]. Available: http://projecteuclid.org/download/pdf_1/euclid.aos/1176345632
- [41] J. Rice, "Choice of smoothing parameter in deconvolution problems," *Contemporary Math.*, vol. 59, pp. 137–151, 1986.
- [42] G. Golub, M. Heath, and G. Wahba, "Generalized cross-validation as a method for choosing a good ridge parameter," *Technometrics*, vol. 21, no. 2, pp. 215–223, 1979.
- [43] J. S. Maritz and T. Lwin, *Empirical Bayes Method*. London, United Kingdom: Chapman and Hall, 1989.
- [44] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed. New York, NY, USA: Springer, 2009.
- [45] Y. Xu and J. Choi, "Adaptive sampling for learning Gaussian processes using mobile sensor networks," *Sens.*, vol. 11, no. 3, pp. 3051–3066, 2011.
- [46] J. Snoek, H. Larochelle, and R. P. Adams, "Practical bayesian optimization of machine learning algorithms," in *Proc. 25th Int. Conf. Neural Inf. Process. Syst.*, 2012, pp. 2951–2959.
- [47] W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, *{Markov} Chain {Monte Carlo} in Practice*. London, United Kingdom: Chapman and Hall, 1996.
- [48] P. Magni, R. Bellazzi, and G. D. Nicolao, "Bayesian function learning using mcmc methods," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 20, no. 12, pp. 1319–1331, Dec. 1998.
- [49] B. M. Bell and G. Pillonetto, "Estimating parameters and stochastic functions of one variable using nonlinear measurement models," *Inverse Problems*, vol. 20, no. 3, pp. 627–646, Jun. 2004. [Online]. Available: <http://stacks.iop.org/0266-5611/20/i=3/a=001?key=crossref.867ad40d4bc8a85b48424adef0274c8>
- [50] G. De Nicolao and G. F. Trecate, "Consistent identification of narx models via regularization networks," *IEEE Trans. Autom. Control*, vol. 44, no. 11, pp. 2045–2049, Nov. 1999.
- [51] G. Wahba, *Spline Models for Observational Data*. Philadelphia, PA, USA: SIAM, 1990.
- [52] D. J. C. MacKay, "Bayesian interpolation," *Neural Comput.*, vol. 4, no. 3, pp. 415–447, 1992.
- [53] G. De Nicolao, G. Sparacino, and C. Cobelli, "Nonparametric input estimation in physiological systems: Problems, methods and case studies," *Autom.*, vol. 33, pp. 851–870, 1997.
- [54] G. Kimeldorf and G. Wahba, "A correspondence between bayesian estimation on stochastic processes and smoothing by splines," *Annals Math. Statist.*, vol. 41, no. 2, pp. 495–502, 1970.
- [55] B. Schölkopf, R. Herbrich, and A. J. Smola, "A generalized representer theorem," *Neural Netw. Comput. Learn. Theory*, vol. 81, pp. 416–426, 2001.
- [56] P. Drineas and M. Mahoney, "On the Nyström method for approximating a gram matrix for improved kernel-based learning," *J. Mach. Learn. Res.*, vol. 6, pp. 2153–2175, Dec. 2005.
- [57] T. Yang, Y. Li, M. Mahdavi, R. Jin, and Z. Zhou, "Nyström method vs random Fourier features: A theoretical and empirical comparison," in *Proc. 25th Int. Conf. Neural Inf. Process. Syst.*, 2012, pp. 476–484.
- [58] "Colorado monthly meteorology dataset 1995–1997," <https://www.image.ucar.edu/Data/US.monthly.met/CO.shtml>, Accessed on: Jun. 20, 2017.



Gianluigi Pillonetto received the doctoral degree in computer science engineering *summa cum laude* from the University of Padova, in 1998 and the PhD degree in bioengineering from the Polytechnic of Milan, in 2002. In 2000 and 2002 he was visiting scholar and visiting scientist, respectively, with the Applied Physics Laboratory, University of Washington, Seattle. In 2005, he became an assistant professor of control and dynamic systems with the Department of Information Engineering, University of Padova where he currently serves as an associate professor. His research interests include system identification, estimation, and machine learning. On these topics, he has published around 60 papers on peer reviewed international journals. He currently serves as associate editor for *Automatica* and the *IEEE Transactions on Automatic Control*. In 2003 he received the Paolo Durst award for the best Italian PhD thesis in Bioengineering, and he was the 2017 recipient of the Automatica Prize. He has been plenary speaker at the 18th IFAC Symposium on System Identification at Stockholm in 2018.



Luca Schenato received the Dr Eng degree in electrical engineering from the University of Padova, in 1999 and the PhD degree in electrical engineering and computer sciences from the U.C. Berkeley, in 2003. He held a post-doctoral position in 2004 and a visiting professor position in 2013-2014 with U.C. Berkeley. Currently he is an associate professor with the Information Engineering Department, University of Padova. His interests include networked control systems, multi-agent systems, wireless sensor networks, smart grids and cooperative robotics. He has been awarded the 2004 Researchers Mobility Fellowship by the Italian Ministry of Education, University and Research (MIUR), and the 2006 Eli Jury Award in U.C. Berkeley and the EUCA European Control Award in 2014, and IEEE fellow in 2017. He served as an associate editor for the *IEEE Transactions on Automatic Control* from 2010 to 2014 and he is currently an associate editor for the *IEEE Transactions on Control of Network Systems*.



Damiano Varagnolo received the MS degree in automation engineering and the PhD degree in information engineering from the University of Padova, Italy, respectively, in 2005 and 2011. He worked as a research engineer with Tecnogamma S.p.A., Treviso, Italy during 2006-2007 and visited UC Berkeley as a scholar researcher in 2010. From March 2012 to December 2013 he worked as a post-doctoral scholar with KTH, Royal Institute of Technology, Stockholm. Currently he is a senior lecturer with LTU, Luleå University of Technology. His interests include distributed optimization, distributed estimation, identification and control of heating and cooling systems.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.