

Learning distributed channel access policies for networked estimation: data-driven optimization in the mean-field regime

Marcos M. Vasconcelos*

MARCOV@VT.EDU

Virginia Polytechnic Institute and State University

Editors: R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, M. Kochenderfer

Abstract

The problem of communicating sensor measurements over shared networks is prevalent in many modern large-scale distributed systems such as cyber-physical systems, wireless sensor networks and the internet of things. Due to bandwidth constraints, the system designer must jointly optimize decentralized medium access transmission and estimation policies that accommodate a very large number of devices in extremely contested environments such that the collection of all observations is reproduced at the destination with the best possible fidelity. We formulate a remote estimation problem in the mean-field regime where a very large number of sensors communicate their observations to an access point, or base-station, under a strict constraint on the maximum fraction of transmitting devices. We show that in the mean-field regime, this problem exhibits a structure which enables tractable optimization algorithms. More importantly, we obtain a data-driven learning scheme and a characterization of its convergence rate.

Keywords: Networks, medium access control, stochastic optimization, convex-concave procedure, statistical learning

1. Introduction

State estimation is a fundamental building block of control systems subject to uncertain disturbances. Most modern applications rely on smart sensors that collect and transmit data to a remote location. In particular, for cyber-physical systems, where a large number of sensors collect and transmit data over a bandwidth limited network, there is a need to communicate efficiently and maintain the quality of the estimates as close as possible to optimal. Moreover, these conflicting goals must be achieved under the lack of knowledge of the underlying stochastic model of the environment. This task requires a combination of techniques in machine learning, decision-theory, and optimization.

Consider a distributed remote estimation system, where a very large number of sensors collect data and communicate with a destination over a constrained network. To obtain a tractable problem, we assume that the data is independent and identically distributed across sensors. When the number of sensors in the system tends to infinity, this assumption leads to an optimization problem for which the optimal communication policy is determined by designing a quantizer under a rate constraint, a problem known to be non-convex. We show how to obtain a locally optimal solution using an algorithm known as the Convex-Concave procedure, which is guaranteed to converge regardless of the underlying probability density function of the data.

* The work of M. M. Vasconcelos was supported by funding from the Commonwealth Cyber Initiative (CCI).

During the last decade, there has been significant research activity in systems with an asymptotically large number of coupled decision-making agents (Mahajan et al., 2013; Gagrani et al., 2020; Sanjari and Yüksel, 2021). Such problem formulations are very relevant in modern applications such as industrial internet of things, where many tiny, low-power devices sense the environment and communicate with a remote gateway, base-station or access point (Gatsis, 2021). Another application where the mean-field regime is appropriate for analysis is in robotic swarms (Zheng et al., 2021). State estimation for control of power systems in the mean-field regime has been considered in the literature (Chen et al., 2017, and references therein). The main idea behind studying teams with an infinite number of identical/symmetric agents is to allow the system designer to perform the optimization over a single policy instead of accounting for the overwhelming combinatorial-like complexity that arise from the interactions among finitely many discrete agents.

The problem of remote estimation over a collision channel was introduced in (Vasconcelos and Martins, 2017) for a finite number of sensors communicating over a channel with unit capacity. Recently, Zhang et al. (2021) extended the setup for a channel with arbitrary capacity, and introduced the first instance of such remote estimation problems when the distribution is unknown. Finally, Vasconcelos and Mitra (2021) considered using a nonparametric statistics technique known as *Kernel Density Estimation* (KDE) (Wasserman, 2006), to obtain a scheme that is more stable and sample-efficient than the one obtained via empirical-risk minimization, also known as sample average approximation in the stochastic programming literature (Shapiro et al., 2021). We use the fact that our algorithm is independent of the data’s pdf to come up with a data-driven scheme where the system designer relies only on a finite collection of data points using KDE. KDE enables a data-based algorithm and with the choice of a Gaussian kernel, which leads to an algorithm is free of integrals, implying in great numeric stability (Asi and Duchi, 2019). More importantly, we obtain the convergence rate that the probability that system will operate within the desired collision free regime, as a function of the state-dimension, the number of samples and a capacity back-off parameter. Our convergence rate result reveals an information-theoretic interpretation: The number of samples required to guarantee operation in the target regime is inversely proportional to how far we are willing to operate from the system capacity.

2. Problem setup

Consider a remote estimation system where a collection of n agents observe independent and identically distributed (iid) random vectors $\{X_i\}_{i=1}^n$, where $X_i \in \mathbb{R}^d$ and $X_i \sim f$, $i \in \{1, \dots, n\}$. Assume throughout this paper that the data’s probability density function (pdf) f is smooth, and supported over the entire \mathbb{R}^d .¹ The agents communicate their observations to a remote destination over a shared communication medium: only a portion $\kappa(n) < n$ of the agents is allowed to communicate noiselessly at any time. The agents behave strategically as to which observations to transmit. After observing X_i , the i -th agent makes a binary decision $U_i \in \{0, 1\}$ to transmit its observation ($U_i = 1$) or to remain silent ($U_i = 0$). The decision variable U_i is computed according to a policy $\gamma_i : \mathbb{R}^d \rightarrow \{0, 1\}$, such that $U_i = \gamma_i(X_i)$.

If the number of transmissions is less than or equal to the capacity $\kappa(n)$, all the transmitted observations are perfectly recovered at the receiver, and only the non-received observations need to be estimated. If the number of transmissions is larger than the channel capacity $\kappa(n)$, nothing can

1. For instance, this assumption is satisfied if a Gaussian mixture describes the data well and encompasses a broad class of distributions.

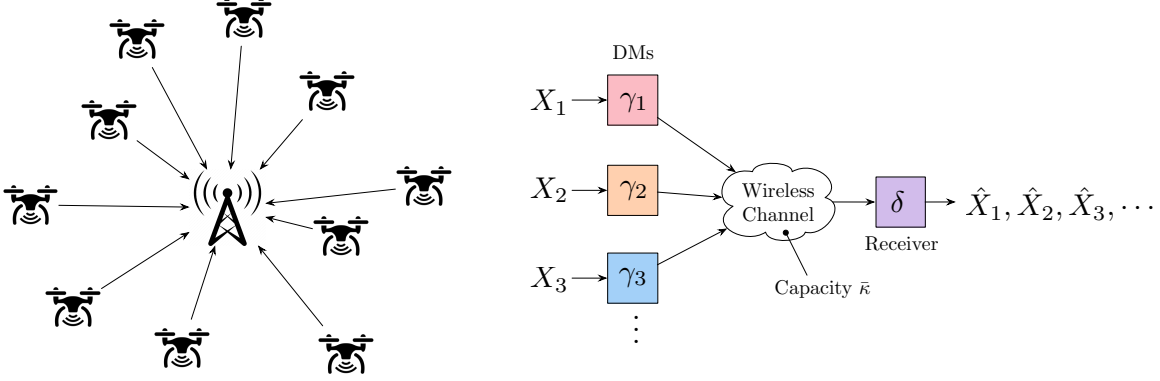


Figure 1: System diagram: A large number of decision makers in a remote estimation system communicating their observations to a base station or access point without coordination, and its block diagram abstraction considered for analytical problem formulation herein.

be decoded at the receiver, and *all* of the observations need to be estimated. This is mathematically described by the following estimation rule:

$$\hat{X}_i = \begin{cases} X_i, & \text{if } \sum_{j=1}^n U_j \leq \kappa(n), U_i = 1 \\ \theta_i, & \text{if } \sum_{j=1}^n U_j \leq \kappa(n), U_i = 0 \\ \alpha_i, & \text{if } \sum_{j=1}^n U_j > \kappa(n) \end{cases} \quad (1)$$

where $\theta_i, \alpha_i \in \mathbb{R}^d$, $i \in \{1, \dots, n\}$. Our goal is to jointly design the communication policies $\gamma \stackrel{\text{def}}{=} \{\gamma_i\}_{i=1}^n$ and the estimation parameters $\theta \stackrel{\text{def}}{=} \{\theta_i\}_{i=1}^n$ and $\alpha \stackrel{\text{def}}{=} \{\alpha_i\}_{i=1}^n$ that minimize the normalized mean-squared error (NMSE) between the observations and their estimates:

$$\mathcal{J}_n(\gamma, \theta, \alpha) \stackrel{\text{def}}{=} \frac{1}{n} \mathbf{E} \left[\sum_{i=1}^n \|X_i - \hat{X}_i\|^2 \right]. \quad (2)$$

Theorem 1 *If $\lim_{n \rightarrow \infty} \kappa(n)/n = \bar{\kappa} \in (0, 1)$, the problem of jointly optimizing the NMSE with respect to the transmission and estimation policies for a system where $n \rightarrow \infty$ is equivalent² to solving:*

$$\underset{\gamma \in \Gamma, \theta \in \mathbb{R}^d}{\text{minimize}} \quad \mathbf{E}[\|X - \theta\|^2 | U = 0] \mathbf{P}(U = 0) \quad (3)$$

$$\text{subject to} \quad \mathbf{P}(U = 1) \leq \bar{\kappa}, \quad (4)$$

where $X \sim f$ and $\Gamma = \{\gamma \mid \gamma : \mathbb{R}^d \rightarrow \{0, 1\}\}$.

2. The optimal policies and the minimum of the objective function are equal for both problems.

Proof Using the law of total expectation, we expand the NMSE in Eq. (2) as:

$$\begin{aligned} \mathcal{J}_n(\boldsymbol{\gamma}, \boldsymbol{\theta}, \boldsymbol{\alpha}) &= \frac{1}{n} \sum_{i=1}^n \mathbf{E} \left[\left\| X_i - \hat{X}_i \right\|^2 \middle| U_i = 0, \sum_{j \neq i}^n U_j \leq \kappa(n) \right] \mathbf{P}(U_i = 0) \mathbf{P} \left(\sum_{j=1}^n U_j \leq \kappa(n) \right) \\ &+ \frac{1}{n} \sum_{i=1}^n \mathbf{E} \left[\left\| X_i - \hat{X}_i \right\|^2 \middle| U_i = 1, \sum_{j \neq i}^n U_j \leq \kappa(n) - 1 \right] \mathbf{P}(U_i = 1) \mathbf{P} \left(\sum_{j \neq i}^n U_j \leq \kappa(n) - 1 \right) \\ &+ \frac{1}{n} \mathbf{E} \left[\sum_{i=1}^n \left\| X_i - \hat{X}_i \right\|^2 \middle| \sum_{j=1}^n U_j > \kappa(n) \right] \mathbf{P} \left(\sum_{j=1}^n U_j > \kappa(n) \right). \quad (5) \end{aligned}$$

The iid assumption implies that $\sum_{j \neq i} U_j \perp\!\!\!\perp X_i$. Then, using the estimator's definition in Eq. (1), and dropping the conditioning on the event $\{\sum_{j \neq i} U_j \leq \kappa(n)\}$ from the first conditional expectation in Eq. (5), we get

$$\begin{aligned} \mathcal{J}_n(\boldsymbol{\gamma}, \boldsymbol{\theta}, \boldsymbol{\alpha}) &= \frac{1}{n} \sum_{i=1}^n \mathbf{E} \left[\left\| X_i - \theta_i \right\|^2 \middle| U_i = 0 \right] \mathbf{P}(U_i = 0) \mathbf{P} \left(\sum_{j=1}^n U_j \leq \kappa(n) \right) \\ &+ \frac{1}{n} \mathbf{E} \left[\sum_{i=1}^n \left\| X_i - \alpha_i \right\|^2 \middle| \sum_{j=1}^n U_j > \kappa(n) \right] \mathbf{P} \left(\sum_{j=1}^n U_j > \kappa(n) \right). \quad (6) \end{aligned}$$

The iid assumption on the observation sequence $\{X_i\}_{i=1}^n$ leads to symmetry in the problem such that each agent is indistinguishable from the others. Therefore, we may assume without loss of optimality that every agent uses the same policy $\gamma_i = \gamma$, $i \in \{1, \dots, n\}$. Thus, the number of agents that decide to communicate their observations to the receiver is a Binomial random variable. Let $p = \mathbf{P}(U_i = 1)$. For $\bar{\kappa} \in (0, 1)$, the following holds:

$$\lim_{n \rightarrow \infty} \frac{\kappa(n)}{n} = \bar{\kappa} \implies \mathbf{P} \left(\sum_{j=1}^n U_j > \kappa(n) \right) \longrightarrow \mathbf{1}(p > \bar{\kappa}), \quad \text{a.e.} \quad (7)$$

The last step is proved using the following sequence of identities:

$$\begin{aligned} \lim_{n \rightarrow \infty} \mathbf{P} \left(\sum_{j=1}^n U_j > \kappa(n) \right) &= \lim_{n \rightarrow \infty} \mathbf{E} \left[\mathbf{1} \left(\frac{1}{n} \sum_{j=1}^n U_j > \frac{\kappa(n)}{n} \right) \right] \\ &\stackrel{(a)}{=} \mathbf{E} \left[\lim_{n \rightarrow \infty} \mathbf{1} \left(\frac{1}{n} \sum_{j=1}^n U_j > \frac{\kappa(n)}{n} \right) \right] \stackrel{(b)}{=} \mathbf{E} \left[\mathbf{1}(p > \bar{\kappa}) \right] \stackrel{(c)}{=} \mathbf{1}(p > \bar{\kappa}), \quad (8) \end{aligned}$$

where (a) follows from the Lebesgue dominated convergence theorem, (b) follows from the strong law of large numbers, and (c) follows from the fact that the argument of the expectation operator is a deterministic function. Therefore, using the fact that $\{X_i\}$ is iid, we argue that $\theta_1 = \theta_2 = \dots$. Moreover, it is possible to show that the optimal value of α_i^* satisfies:

$$\lim_{n \rightarrow \infty} \alpha_i^* = \mathbf{E} \left[X_i \middle| \sum_{j=1}^n U_j > \kappa(n) \right] = \mathbf{E}[X], \quad i \in \{1, 2, \dots\}. \quad (9)$$

Observe from Eq. (7) that by an appropriate choice of p , we may operate in one of two regimes: the *collision-free* regime ($p \leq \bar{\kappa}$) or the *constant-collision* regime ($p > \bar{\kappa}$). And consequently, the NMSE in the regime when $n \rightarrow \infty$ becomes:

$$\mathcal{J}_\infty(\gamma, \theta) = \begin{cases} \mathbf{E}[\|X - \theta\|^2 | U = 0] \mathbf{P}(U = 0) & \text{if } \mathbf{P}(U = 1) \leq \bar{\kappa} \\ \mathbf{E}[\|X - \mathbf{E}[X]\|^2] & \text{otherwise.} \end{cases} \quad (10)$$

Clearly, the constant-collision regime offers worse performance than the collision-free regime. ■

Remark 2 Notice that in the asymptotic regime, the optimization variables α are not present, meaning that the detrimental effect of collisions can be mitigated by adjusting the probability of transmission to operate below the asymptotic channel capacity parameter $\bar{\kappa}$.

3. Solving the optimization problem

Once the constrained optimization problem for the remote estimation system in the asymptotic regime is established, we may obtain the structure of its optimal solutions. Notice that there are two variables in the optimization problem: the function γ , and a vector θ . Therefore, this is an infinite dimensional problem. Moreover, the objective in the optimization problem is non-convex. Writing the Lagrangian associated with Eq. (3), we characterize its saddle points. Let $\lambda \geq 0$ and define:

$$\mathcal{L}(\gamma, \theta, \lambda) \stackrel{\text{def}}{=} \int_{\mathbb{R}^d} \|x - \theta\|^2 \mathbf{1}(\gamma(x) = 0) f(x) dx + \lambda \left(\int_{\mathbb{R}^d} \mathbf{1}(\gamma(x) = 1) f(x) dx - \bar{\kappa} \right) \quad (11)$$

First, we have the following structural result, which establishes that the communication policy in every saddle point is of the threshold type (Vasconcelos and Martins, 2017, and references therein).

Proposition 3 If $(\theta^*, \gamma^*, \lambda^*)$ is a saddle-point of $\mathcal{L}(\theta, \gamma, \lambda)$, then

$$\gamma^*(x) = \begin{cases} 0, & \text{if } \|x - \theta^*\|^2 \leq \lambda^* \\ 1, & \text{otherwise.} \end{cases} \quad (12)$$

Furthermore,

$$\mathbf{P}(\|X - \theta^*\|^2 > \lambda^*) = \bar{\kappa} \quad \text{and} \quad \theta^* = \frac{1}{1 - \bar{\kappa}} \int_{\mathbb{R}^d} x \mathbf{1}(\|x - \theta^*\|^2 \leq \lambda^*) f(x) dx. \quad (13)$$

Proof From the definition of the Lagrangian in Eq. (11), for any $\theta \in \mathbb{R}^d$ and $\lambda \geq 0$, the communication policy that minimizes it may be constructed by assigning to a set $\mathcal{A}_0 \subset \mathbb{R}^d$, all the points $x \in \mathbb{R}^d$ such that $\|x - \theta\|^2 \leq \lambda$. Otherwise, the point is assigned to $\mathcal{A}_1 = \mathbb{R}^d \setminus \mathcal{A}_0$. If there are any sets with non-zero measure for which $\|x - \theta\|^2 \leq \lambda$ assigned to \mathcal{A}_1 , we can reduce the Lagrangian even further and $(\theta^*, \gamma^*, \lambda^*)$ would not constitute a saddle-point.

To establish the second result, we take the derivative of Eq. (11) with respect to λ and use Eq. (12), which leads to:

$$\frac{\partial}{\partial \lambda} \mathcal{L}(\gamma^*, \theta^*, \lambda^*) = 0 \implies \int_{\mathbb{R}^d} \mathbf{1}(\gamma^*(x) = 1) f(x) dx = \mathbf{P}(\|X - \theta^*\|^2 > \lambda^*) = \bar{\kappa}. \quad (14)$$

Lastly, taking the gradient of Eq. (11) with respect to θ , we have:

$$\nabla_{\theta} \mathcal{L}(\gamma^*, \theta^*, \lambda^*) = 0 \implies \int_{\mathbb{R}^d} -2(x - \theta^*) \mathbf{1}(\gamma^*(x) = 0) f(x) dx = 0. \quad (15)$$

Then, using the characterization in Eqs. (12) and (13), we have:

$$\int_{\mathbb{R}^d} x \mathbf{1}(\|x - \theta^*\|^2 \leq \lambda^*) f(x) dx = \theta^* \underbrace{\int_{\mathbb{R}^d} \mathbf{1}(\|x - \theta^*\|^2 \leq \lambda^*) f(x) dx}_{=1-\bar{\kappa}}. \quad (16)$$

■

The structural results stated in Proposition 3 help us define an algorithm that converges to a saddle-point of the Lagrangian and therefore defines a locally optimal policy for the remote estimation system in the asymptotic regime. First, consider the Laplacian when λ and θ , are arbitrarily fixed. Then, using Eq. (12), we obtain the function $\tilde{\mathcal{L}} : \mathbb{R}^{d+1} \rightarrow \mathbb{R}$ defined as:

$$\tilde{\mathcal{L}}(\theta, \lambda) \stackrel{\text{def}}{=} \mathcal{L}(\gamma^*, \theta, \lambda) = \mathbf{E} \left[\min \{ \|X - \theta\|^2, \lambda \} \right] - \lambda \bar{\kappa}. \quad (17)$$

We are now interested in minimizing the function $\tilde{\mathcal{L}}$ over θ for any $\lambda \geq 0$. The function in Eq. (17) is non-convex in θ . In fact, this new objective function is the expectation of a clipped convex function, which is known to lead to NP hard problems for discrete random variables (Barratt et al., 2020). Therefore, there is little hope that we will be able to obtain globally optimal solutions without making additional assumptions on the problem's structure. However, the function possess a very convenient structure, namely, that the term in the expectation can be decomposed as a difference of convex functions, also known as a *DC decomposition*. The DC decomposition enables the use of an algorithm known as the *Convex-Concave Procedure* (CCP) (Lipp and Boyd, 2016). For the purpose of optimization over θ , we may temporarily ignore the term $-\lambda \bar{\kappa}$.

3.1. Convex-Concave Procedure

To obtain a DC decomposition of the objective function, simply notice that the expectation of the minimum can be expressed as follows:

$$\mathbf{E} \left[\min \{ \|X - \theta\|^2, \lambda \} \right] = \mathbf{E} \left[\|X - \theta\|^2 + \lambda \right] - \mathbf{E} \left[\max \{ \|X - \theta\|^2, \lambda \} \right]. \quad (18)$$

Let $\mathcal{G}_{\lambda}(\theta) \stackrel{\text{def}}{=} \mathbf{E} \left[\max \{ \|X - \theta\|^2, \lambda \} \right]$. The CCP consists of approximating the objective function by replacing $\mathcal{G}_{\lambda}(\theta)$ with its affine approximation at a point $\theta_k \in \mathbb{R}^d$, then minimizing the resulting approximate convex function. Conveniently, the decomposition above leads to:

$$\theta_{k+1} = \arg \min_{\theta \in \mathbb{R}^d} \left\{ \mathbf{E} [\|X - \theta\|^2] + \lambda - \left(\mathcal{G}_{\lambda}(\theta_k) + g_{\lambda}(\theta_k)^T (\theta - \theta_k) \right) \right\}, \quad (19)$$

where g_{λ} is any subgradient of \mathcal{G}_{λ} . The first benefit of using the DC decomposition above where the first function is quadratic, is that solving the optimization problem above is equivalent to finding the solution of the first-order optimality condition

$$\nabla_{\theta} \tilde{\mathcal{L}}(\theta_{k+1}, \lambda) = -\mathbf{E} [2(X - \theta_{k+1})] - g_{\lambda}(\theta_k) = 0 \implies \theta_{k+1} = \frac{1}{2} g_{\lambda}(\theta_k) + \mathbf{E}[X]. \quad (20)$$

The second benefit of using the DC decomposition stems from the fact that computing a subgradient g_λ of the expectation of pointwise maximum of functions is extremely simple. For instance, using elementary subgradient calculus, the following is a suitable choice:

$$g_\lambda(\theta) = -2\mathbf{E}\left[(X - \theta)\mathbf{1}(\|X - \theta\|^2 > \lambda)\right], \quad (21)$$

which leads to:

$$\theta_{k+1} = \mathbf{E}\left[X\mathbf{1}(\|X - \theta_k\|^2 \leq \lambda)\right] + \theta_k\mathbf{P}(\|X - \theta_k\|^2 > \lambda). \quad (22)$$

The recursion in Eq. (22) is guaranteed to converge (Lanckriet and Sriperumbudur, 2009; Lipp and Boyd, 2016). Moreover, this convergence is irrespective of the pdf f of the random vector X . Then, by considering the limit of the sequence $\{\theta_k\} \rightarrow \theta^*$, we have:

$$\theta^* = \frac{\mathbf{E}\left[X\mathbf{1}(\|X - \theta^*\|^2 \leq \lambda)\right]}{1 - \mathbf{P}(\|X - \theta^*\|^2 > \lambda)}. \quad (23)$$

If λ is chosen such that $\mathbf{P}(\|X - \theta^*\|^2 > \lambda) = \bar{\kappa}$, the CCP in Eq. (22) converges to a saddle point of the Lagrangian and defines a locally optimal solution of the optimization problem in the asymptotic regime. Algorithm 1 summarizes the discussion above.

4. Data-driven algorithm

One remarkable feature of this problem is the number of structural and convergence results we can obtain without the explicit knowledge of the pdf f . However, our algorithm relies on the availability of a smooth pdf to compute the sequence of thresholds λ_k . For that reason, a data-driven version of our algorithm cannot rely on non-smooth approximations, e.g. using the empirical distribution of the data. Our problem requires a smooth data-driven approximation. One alternative is to use Kernel Density Estimation (KDE), where the pdf f is approximated by:

$$\hat{f}_M(x) = \frac{1}{M} \sum_{m=1}^M \prod_{j=1}^d \frac{1}{\sqrt{2\pi}h_M} \exp\left(-\frac{(x_j - x_{mj})^2}{2h_M}\right) \quad \text{with } h_M = \frac{1.06}{M^{1/5}} \min\left\{s, \frac{Q}{1.34}\right\}, \quad (24)$$

which corresponds to KDE using a product Gaussian kernel (Wasserman, 2006). The quantity M is the number of data samples in a batch, s is the sample standard deviation and Q is the interquartile range³. In addition to the smoothness of the estimate \hat{f}_M , one special feature of using KDE to approximate the pdf is that the estimation quality can be quantified via the mean-squared integration error (MISE):

$$\text{MISE} = \mathbf{E}\left[\int_{\mathbb{R}^d} |\hat{f}_M(x) - f(x)|^2 dx\right] = \mathcal{O}\left(M^{-4/(d+4)}\right), \quad (25)$$

where the expectation is taken with respect to the true distribution f from which the data samples $\{x_m\}_{m=1}^M$ are drawn. It can be shown that no other non-parametric density estimation scheme can produce estimates with a better convergence rate if the estimation quality is measured using the MISE criterion (Van der Vaart, 2000).

3. The interquartile range is the difference between the 75th and 25th percentiles of the data.

Algorithm 1 Data-driven policy optimization for remote estimation in the mean-field regime

1. Given a batch of data $\mathcal{D} = \{x_m\}_{m=1}^M$, estimate \hat{f}_M . If f is known, $\hat{f}_M \leftarrow f$

2. **Initialization:** Compute:

$$\theta_0 \leftarrow \int_{\mathbb{R}^d} x \hat{f}_M(x) dx$$

and solve for λ^* the following equation:

$$\int_{\mathbb{R}^d} \mathbf{1}(\|x - \theta_0\|^2 \leq \lambda^*) \hat{f}_M(x) dx = 1 - \bar{\kappa}$$

Set $\lambda_0 \leftarrow \lambda^*$.

3. **Iteration:** For $k = 0, 1, 2, \dots$

a. Repeat the following recursion until convergence to θ^*

$$\theta_{\ell+1} \leftarrow \int_{\mathbb{R}^d} x \mathbf{1}(\|x - \theta_\ell\|^2 \leq \lambda_k) \hat{f}_M(x) dx - \bar{\kappa} \theta_\ell$$

Set $\theta_{k+1} \leftarrow \theta^*$

b. Solve for λ^* the following equation

$$\int_{\mathbb{R}^d} \mathbf{1}(\|x - \theta_k\|^2 \leq \lambda^*) \hat{f}_M(x) dx = 1 - \bar{\kappa}$$

Set $\lambda_{k+1} \leftarrow \lambda^*$

c. Repeat until convergence

Since the algorithm described in the previous section is agnostic to the pdf f , a data-driven version of it simply uses a KDE \hat{f}_M instead of f . Following [Wasserman \(2006\)](#), we have $\hat{f}_M \xrightarrow{P} f$, $M \rightarrow \infty$. Therefore, Algorithm 1 produces good estimates of locally optimal solutions to Eq. (3) as the size of the data-set increases. In the next section, we provide a convergence rate analysis for obtaining good data-driven locally optimal solutions to Eq. (3) in the absence of knowledge of the true probabilistic model f .

5. Convergence rate analysis

Consider a random batch of data samples $\{X_m\}_{m=1}^M$. Assuming that $(\hat{\theta}_M, \hat{\lambda}_M)$ are computed using the procedure in Algorithm 1, they are considered in this section as random variables. Thus, using the approximate pdf to obtain a channel access policy parametrized by $(\hat{\theta}_M, \hat{\lambda}_M)$ may yield a true probability of channel access that violates the constraint in Eq. (3). Such violation may be very detrimental. Here, we provide a convergence rate analysis that the actual probability of channel access will not violate the channel capacity constraint as a function of the number of data samples.

Let $\mathbf{P}(U = 1) = 1 - \int_{\hat{\theta}_M - \sqrt{\hat{\lambda}_M}}^{\hat{\theta}_M + \sqrt{\hat{\lambda}_M}} f(x) dx$, and notice that this quantity is itself a random variable.

Theorem 4 (Convergence rate) *Consider a remote estimation system designed according to Algorithm 1. Assume that the pdf of the observations $f(x)$ is smooth with full support on \mathbb{R}^d . For a wireless network with asymptotic capacity $\bar{\kappa}$ and a channel access policy designed to operate with probability of transmission $\bar{\kappa} - \delta$, for $\delta \in (0, \bar{\kappa})$. Then, the true probability that the the capacity constraint will be violated decays to zero with a rate that satisfies:*

$$\mathbf{P}\left(\mathbf{P}(U = 1) > \bar{\kappa}\right) \leq \mathcal{O}\left(\frac{1}{\delta M^{2/(d+4)}}\right). \quad (26)$$

Proof To guarantee that the probability constraint will be met with high probability, we need to back off from the capacity by a pre-specified constant $\delta \in (0, \bar{\kappa})$. Let $\hat{\mathbf{P}}(U = 1) \stackrel{\text{def}}{=} 1 - \int_{\hat{\theta}_M - \sqrt{\hat{\lambda}_M}}^{\hat{\theta}_M + \sqrt{\hat{\lambda}_M}} \hat{f}_M(x) dx = \bar{\kappa}$ (by design). Then, notice that

$$\mathbf{P}\left(\hat{\mathbf{P}}(U = 1) - \mathbf{P}(U = 1) \geq \delta\right) \leq \mathbf{P}\left(|\hat{\mathbf{P}}(U = 1) - \mathbf{P}(U = 1)| \geq \delta\right). \quad (27)$$

The following inequalities hold:

$$\begin{aligned} \left|\mathbf{P}(U = 1) - \hat{\mathbf{P}}(U = 1)\right| &= \left|\int_{\hat{\theta}_M - \sqrt{\hat{\lambda}_M}}^{\hat{\theta}_M + \sqrt{\hat{\lambda}_M}} [\hat{f}_M(x) - f(x)] dx\right| \\ &\stackrel{(a)}{\leq} \int_{\mathbb{R}} \mathbf{1}\left(\hat{\theta}_M - \sqrt{\hat{\lambda}_M} \leq x \leq \hat{\theta}_M + \sqrt{\hat{\lambda}_M}\right) |\hat{f}_M(x) - f(x)| dx \\ &\stackrel{(b)}{\leq} \left(2\sqrt{\hat{\lambda}_M}\right)^{1/2} \left(\int_{\mathbb{R}} |\hat{f}_M(x) - f(x)|^2 dx\right)^{1/2} \end{aligned}$$

where (a) follows from the Triangle inequality and (b) follows from Hölder's inequality. Finally, we apply the expectation operator with respect to the data $\{X_m\}_{m=1}^M$, obtaining via the Cauchy-Schwarz inequality the following bound:

$$\mathbf{E}\left[\left|\mathbf{P}(U = 1) - \hat{\mathbf{P}}(U = 1)\right|\right] \leq \left(2\mathbf{E}\left[\sqrt{\hat{\lambda}_M}\right]\right)^{1/2} \left(\mathbf{E}\left[\int_{\mathbb{R}^d} |\hat{f}_M(x) - f(x)|^2 dx\right]\right)^{1/2}. \quad (28)$$

Now, since we obtain $\hat{\lambda}_M$ from an algorithmic procedure that depends on a batch of sampled data, we have very little control over it, but we can guarantee that this term is bounded for any number of samples M . Therefore, using Eq. (25), we obtain: $\mathbf{E}\left[\left|\mathbf{P}(U = 1) - \hat{\mathbf{P}}(U = 1)\right|\right] = \mathcal{O}(M^{-2/(d+4)})$. The following holds:

$$\mathbf{P}\left(\mathbf{P}(U = 1) > \bar{\kappa}\right) = \mathbf{P}\left(\mathbf{P}(U = 1) - \hat{\mathbf{P}}(U = 1) > \delta\right) \stackrel{(c)}{\leq} \frac{\mathbf{E}\left[\left|\mathbf{P}(U = 1) - \hat{\mathbf{P}}(U = 1)\right|\right]}{\delta}, \quad (29)$$

where (c) follows from Markov's inequality. ■

Remark 5 (Interpretation) *The result above mimics results related to the capacity in the Shannon sense. To guarantee a decay in the probability of violating the channel capacity constraint we back off from $\bar{\kappa}$ by δ . Notice that as $\delta \rightarrow 0$, our scheme based on KDE requires an increasingly large number of samples.*

6. Numerical results

Consider data samples generated according to the following Gaussian mixture:

$$X \sim 0.2\mathcal{N}(-2, 0.2) + 0.2\mathcal{N}(-1, 0.075) + 0.1\mathcal{N}(0, 0.1) + 0.3\mathcal{N}(1, 0.1) + 0.2\mathcal{N}(2, 0.1). \quad (30)$$

Assuming that this distribution is known, and setting $\bar{\kappa} = 0.5$, Algorithm 1 results in $\theta^* = 0.0592$ and $\lambda^* = 1.5063$. The value of the optimal NMSE in the asymptotic regime is $\mathcal{J}_\infty^* = 0.3411$. Using the data-driven approach based on KDE for several different sample sizes ($M = 10^3, 10^4, 10^5$), we obtain the curves displayed in Fig. 2, where we vary the capacity back-off parameter from $\delta = 10^{-5}$ to 10^{-1} . The estimation performance of the system degrades with δ , but the probability that the capacity constraint will be violated decreases sharply to zero even for moderate values of the sample size Fig. 2 (right, top). However, the KDE is very sensitive to M , which can be seen from Fig. 2 (left), yielding in a large variance in performance for smaller values of M , e.g. $M = 10^3$

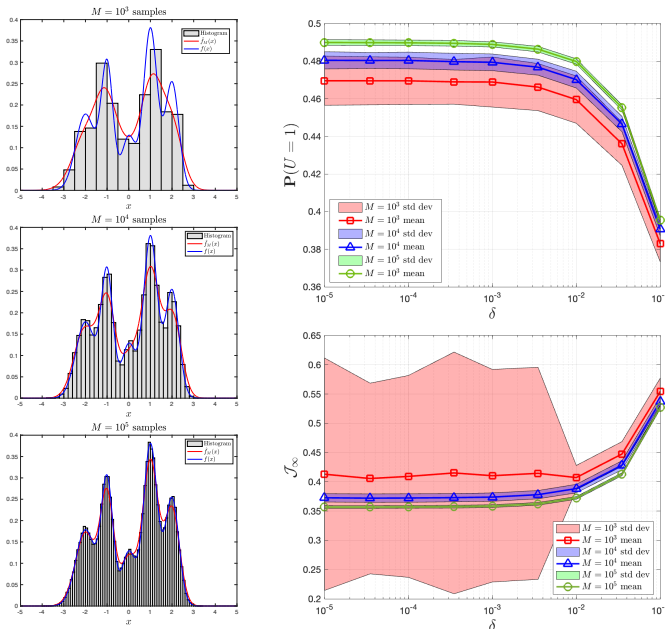


Figure 2: Numerical results.

Fig. 2 (right, bottom). An important observation is that the performance for systems designed using a random batch of data, quickly concentrate around their means, e.g. $M > 10^4$.⁴

7. Conclusions and future work

Scalability issues often challenge modern communication networks and their medium access control mechanisms. We argue that when it comes to applications such as the Internet of Things or in robotic swarms, it is beneficial from the analytical perspective to study the design of policies in the mean-field regime, i.e., systems with an infinite number of devices. One problem of interest consists of estimating an iid process where each random variable is accessed by a different device and transmitted over a constrained network to a remote location. Here, we showed that the interference between simultaneously transmitting agents can be effectively mitigated with appropriate design. More importantly, if we back off from the channel capacity by a constant amount, we can guarantee that the data-driven design based on Kernel Density Estimation will perform satisfactorily, meaning, will not violate a probability constraint with high probability.

4. All the code used to obtain the numerical results contained herein can be found on <https://github.com/mullervasconcelos/L4DC22.git>

References

- Hilal Asi and John C. Duchi. The importance of better models in stochastic optimization. *Proceedings of the National Academy of Sciences*, 116(46):22924–22930, 2019.
- Shane Barratt, Guillermo Angeris, and Stephen Boyd. Minimizing a sum of clipped convex functions. *Optimization Letters*, 14(8):2443–2459, 2020.
- Yue Chen, Ana Bušić, and Sean P. Meyn. State estimation for the individual and the population in mean field control with application to demand dispatch. *IEEE Transactions on Automatic Control*, 62(3):1138–1149, 2017.
- Mukul Gagrani, Sagar Sudhakara, Aditya Mahajan, Ashutosh Nayyar, and Yi Ouyang. Thompson sampling for linear quadratic mean-field teams. *arXiv:2011.04686*, 2020.
- Konstantinos Gatsis. Linear regression over networks with communication guarantees. In *Proceedings of the 3rd Conference on Learning for Dynamics and Control*, number 144, pages 767–778. PMLR, 2021.
- Gert Lanckriet and Bharath K Sriperumbudur. On the convergence of the concave-convex procedure. *Advances in neural information processing systems*, 22:1759–1767, 2009.
- Thomas Lipp and Stephen Boyd. Variations and extension of the convex–concave procedure. *Optimization and Engineering*, 17(2):263–287, 2016.
- Aditya Mahajan, Nuno C. Martins, and Serdar Yüksel. Static LQG teams with countably infinite players. In *52nd IEEE Conference on Decision and Control*, pages 6765–6770, 2013.
- Sina Sanjari and Serdar Yüksel. Optimal solutions to infinite-player stochastic teams and mean-field teams. *IEEE Transactions on Automatic Control*, 66(3):1071–1086, 2021.
- Alexander Shapiro, Darinka Dentcheva, and Andrzej Ruszczyński. *Lectures on stochastic programming: modeling and theory*. SIAM, 2021.
- Aad W. Van der Vaart. *Asymptotic statistics*. Cambridge university press, 2000.
- Marcos M. Vasconcelos and Nuno C. Martins. Optimal estimation over the collision channel. *IEEE Transactions on Automatic Control*, 62(1):321–336, January 2017.
- Marcos M. Vasconcelos and Urbashi Mitra. A sample-efficient scheme for channel resource allocation in networked estimation. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2021.
- Larry Wasserman. *All of nonparametric statistics*. Springer Science & Business Media, 2006.
- Xu Zhang, Marcos M. Vasconcelos, Wei Cui, and Urbashi Mitra. Distributed remote estimation over the collision channel with and without local communication. *IEEE Transactions on Control of Network Systems (early access)*, 2021.
- Tongjia Zheng, Qing Han, and Hai Lin. Transporting robotic swarms via mean-field feedback control. *IEEE Transactions on Automatic Control (early access)*, 2021.