

Distributed Bayesian Estimation of Continuous Variables over Time-varying Directed Networks

Parth Paritosh, Nikolay Atanasov, *Member, IEEE*, and Sonia Martinez, *Fellow, IEEE*

Abstract—This work presents a distributed Bayesian estimation algorithm for time-varying directed sensor networks. We consider a network of sensing agents aiming to estimate continuous variables of interest using direct observations as well as communication across the network. We aim to obtain a probability density function for the unknown variables that best explains the collectively gathered data. To account for point-to-point and broadcast communication, our formulation considers uniformly and strongly connected digraphs. Each agent pools neighbor densities via a weighted geometric average to achieve consensus. We deal with continuous variables via a novel application of large deviation analysis to the estimated probability ratios. Our analysis captures a large class of probability density functions, including Gaussian mixtures, and guarantees that the mode of the estimated density converges to the true parameter value at an exponential rate. The consistency and convergence rate of our algorithm are demonstrated in cooperative localization and distributed target tracking simulations.

Index Terms—Sensor networks, Distributed control, Agents-based systems

I. INTRODUCTION

INTERCONNECTED sensing devices are the bedrock of the information infrastructure in the Internet of Things and autonomous robots. In networked cyber-physical systems, multi-agent interactions enable estimating any quantities of interest. Scalability and robustness considerations motivate distributed algorithms relying on inter-agent communication to achieve similar accuracy and convergence speed as a centralized estimator.

Consider a sensor network aiming to estimate a variable of interest x_* , which dictates the distribution of the sensing agent's observations. The agents face a local identifiability problem in which any single agent's observations may not be sufficient to estimate a unique x_* . To resolve this, the agents thus need to exchange information. This setting has motivated consensus [5], social learning [8], and distributed hypothesis testing [9] techniques. Solutions vary as per the quantity of interest and the communication network.

The simplest communication network is a connected static graph typically represented by a doubly stochastic matrix [6].

We gratefully acknowledge support from ARL DCIST CRA W911NF-17-2-0181, NSF FRR CAREER 2045945, FA9550-18-1-0158, and ONR N00014-19-1-2471.

The authors are with the Contextual Robotics Institute, University of California San Diego, 9500 Gilman Dr, La Jolla, CA 92093 pparitos, natanasov, soniamd@ucsd.edu,

Less restrictive row and column stochastic representations requiring the knowledge of in and out-degrees respectively are practical for time-varying networks [12]. Uniformly strongly connected (C -connected) graphs contain a path between any two nodes in the union of network edges over a given time period. A further relaxed connectivity constraint imposes the same requirement on averaged graph adjacency matrices [16].

Distributed estimation can be posed as an optimization problem over discrete, continuous, or probability spaces and its variations depend on the underlying communication network. The seminal non-Bayesian inference algorithm in [8] estimates a probability mass function (pmf) by arithmetically averaging one-hop neighbor pmfs in a connected network. The algorithm assumptions related to belief transfer and independent observations are discussed in [11]. The distributed hypothesis testing algorithm developed in [18] uses geometric averaging of one-hop neighbor densities in static networks. The corresponding maximum likelihood estimates are shown to converge to the true parameter, if it is globally identifiable. Assuming that log-likelihood functions are concave, [18] shows convergence to the true parameter in probability, whereas [9] shows almost sure convergence. For linear observation models, these distributed Bayesian estimation algorithms specialize to the Kalman filter [15]. Relaxing the connectivity assumptions to C -connected networks, [13] uses geometric averages to show exponential probability decay at sub-optimal hypotheses.

We develop a distributed estimation algorithm for continuous variables in uniformly connected digraphs. Our work extends the results in Nedić et al. [13] on finite space estimation by relaxing the bounds on the agents' log-likelihood ratios, thus enabling estimation for continuous probability density functions (pdf), such as Gaussian mixtures and Gamma distributions. This was recognized and relaxed by Lalitha et al. [9] to achieve distributed hypothesis testing in static networks.

Statement of contributions: We develop a Bayesian distributed estimation algorithm and analyze convergence for continuous variables with unbounded log-likelihood ratios in C -connected networks. For continuous likelihoods, we prove that any large deviation of the probability ratio at any arbitrary to optimal hypothesis x_* decays exponentially. The corresponding rate of convergence depends on the sum of the KL-divergence between the agent observation models evaluated at the hypothesis and x_* . Using the Borel-Cantelli lemma, we show that a mode of the estimated pdf observationally equivalent to x_* exists almost surely. In discrete space, our result implies that the estimated probability mass vanishes

exponentially almost surely over the non-optimal domain.

II. PROBLEM FORMULATION

Consider a time-varying directed graph $\mathcal{G}_t = (\mathcal{V}, \mathcal{E}_t, A_t)$ with node set $\mathcal{V} = \{1, \dots, n\}$, edge set $\mathcal{E}_t \subseteq \mathcal{V} \times \mathcal{V}$, and adjacency matrix $A_t \in \mathbb{R}_{\geq 0}^{n \times n}$. An element $A_{t,ij}$ of the adjacency matrix is positive only if $i = j$ or when there is a directed edge from node i to node j , indicating that node i can message node j . A row stochastic adjacency matrix A_t can be used to model any communication graph [6], including fully distributed one-hop broadcast networks [7]. The graph \mathcal{G}_t is *strongly connected* at time t if there exists a path connecting any two nodes. In practice, the communication network may not be strongly connected at each time t . Instead, we consider a C -connected network [13] such that the C -step union of graphs is strongly connected.

Assumption 1. At any time t , the graph \mathcal{G}_t satisfies:

- 1.1 (Row stochastic weights) The adjacency matrix A_t is row-stochastic, i.e., $A_{t,ii} > 0$ for all $i \in \mathcal{V}$ and $A_t \mathbf{1} = \mathbf{1}$, where $\mathbf{1}$ is a vector of ones.
- 1.2 (C -connectivity) The C -step union $(\mathcal{V}, \cup_{k=0}^{C-1} \mathcal{E}_{t+k})$ of the graphs $\mathcal{G}_t, \dots, \mathcal{G}_{t+C-1}$ is strongly connected.

The graph \mathcal{G}_t is used to model the communication among n agents, each associated with a node in \mathcal{G}_t . The agents aim to cooperatively estimate a parameter of interest, $\mathbf{x}_* \in \mathcal{X} \subseteq \mathbb{R}^m$. Each agent i is equipped with a sensor that provides observations $\mathbf{z}_{i,t} \in \mathbb{R}^\ell$ at each t sampled from an observation model specified by pdf $q_i(\mathbf{z}_{i,t} | \mathbf{x} = \mathbf{x}_*) \in \mathcal{F}_\ell$ conditioned on the true parameter value \mathbf{x}_* . The space \mathcal{F}_ℓ of pdfs is given as:

$$\mathcal{F}_\ell = \left\{ \mathbf{g} \in L^1(\mathbb{R}^\ell) \mid \int g(\mathbf{y}) d\mathbf{y} = 1, g(\mathbf{y}) \geq 0, \forall \mathbf{y} \in \mathbb{R}^\ell \right\}.$$

For any agent i , the known conditional density $q_i(\mathbf{z}_{i,t} | \mathbf{x})$ serves as the likelihood model. The parameter values in Agent i 's optimal set \mathcal{X}_i^* minimize the divergence between the true $q_i(\cdot | \mathbf{x}_*) \in \mathcal{F}_\ell$ and the evaluated $q_i(\cdot | \mathbf{x}) \in \mathcal{F}_\ell$ observation models. The optimal parameters common across all agents form the set \mathcal{X}_* given as:

$$\mathcal{X}_* \equiv \cap_{i=1}^n \mathcal{X}_i^*, \quad \mathcal{X}_i^* = \arg \min_{\mathbf{x} \in \mathcal{X}} H_i(\mathbf{x}, \mathbf{x}_*). \quad (1)$$

Here, the *KL-divergence* term for agent i is $H_i(\mathbf{x}, \mathbf{x}_*) = \text{KL}(q_i(\cdot | \mathbf{x}_*) \| q_i(\cdot | \mathbf{x})) = \int q_i(\mathbf{z} | \mathbf{x}_*) \log \frac{q_i(\mathbf{z} | \mathbf{x}_*)}{q_i(\mathbf{z} | \mathbf{x})} d\mathbf{z}$, which quantifies the difference between conditional densities induced by true and arbitrary values of \mathbf{x} .

Assuming conditional independence of observations, the joint likelihood model of the sensor network is $q(\mathbf{z}_t | \mathbf{x}) \triangleq \prod_{i \in \mathcal{V}} q_i(\mathbf{z}_{i,t} | \mathbf{x}) \in \mathcal{F}_{n\ell}$. Here, the variable \mathbf{z}_t represents the collection of observations $\mathbf{z}_{i,t}$ over all n agents at time t .

Assumption 2. (Independent observations) The measurements $\mathbf{z}_{i,t} \sim q_i(\cdot | \mathbf{x}_*)$ collected by agent i at time t are independent across time and agents.

If agent i 's likelihood model assigns equal probabilities at distinct values $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{X}$ for any observation, then the values are *observationally equivalent*, i.e. $q_i(\cdot | \mathbf{x}_1) = q_i(\cdot | \mathbf{x}_2)$. In

the case of data generated from densities parametrized by a common \mathbf{x}_* at each agent, we formulate the problem of finding a value observationally equivalent to \mathbf{x}_* .

Problem 1. How can each agent i in a C -connected network estimate a common parameter \mathbf{x}_* using local observations from $q_i(\cdot | \mathbf{x}_*)$ and estimates communicated by its neighbors?

III. DISTRIBUTED ESTIMATION FOR CONTINUOUS VARIABLES

In this section, we find the parameters in the set $\mathcal{X}^* \subset \mathbb{R}^m$ by estimating a pdf $p^* \in \mathcal{F}_m$ over the parameter space \mathcal{X} . The product of the known likelihood model with this pdf p^* best approximates the joint observation model $q(\cdot | \mathbf{x}_*) = \prod_{i \in \mathcal{V}} q_i(\cdot | \mathbf{x}_*)$ in the following sense,

$$p^* \in \arg \min_{p \in \mathcal{F}_m} \int \text{KL}(q(\cdot | \mathbf{x}_*) \| q(\cdot | \mathbf{x})) p(\mathbf{x}) d\mathbf{x}. \quad (2)$$

The objective is a linear, convex function in the minimizing argument p . To understand the objective, we note that, $q(\mathbf{z} | \mathbf{x}_*) \neq q(\mathbf{z} | \mathbf{x})$ a.e. in \mathbf{z} for all $\mathbf{x} \notin \mathcal{X}^*$, implies that $\text{KL}(q(\cdot | \mathbf{x}_*) \| q(\cdot | \mathbf{x})) > 0$ for all $\mathbf{x} \notin \mathcal{X}^*$. Since the KL-divergence is a continuous functional, the objective is positive for any function $p(\mathbf{x})$ whose support contains a positive measure subset of $\mathcal{X} \setminus \mathcal{X}^*$. In other words, only those pdfs $p(\mathbf{x})$ that place their mass entirely over \mathcal{X}^* will result into zero expected divergence. Thus, we can reduce the problem of finding the values \mathcal{X}^* to finding an optimal pdf p . The entropy term $q(\cdot | \mathbf{x}_*) \log(q(\cdot | \mathbf{x}_*))$ is constant w.r.t. the minimizing argument $p \in \mathcal{F}_m$ and $\int p = 1$, implying that we can drop this term from the objective:

$$p^* \in \arg \min_{p \in \mathcal{F}_m} - \int \left[\int q(\mathbf{z} | \mathbf{x}_*) \log(q(\mathbf{z} | \mathbf{x})) d\mathbf{z} \right] p(\mathbf{x}) d\mathbf{x}.$$

We can switch the order of integration in the objective using Fubini's theorem on the finite cross entropy integral. Denoting $F(p, \mathbf{z}) = \mathbb{E}_{\mathbf{x} \sim p} [-\log(q(\mathbf{z} | \mathbf{x}))]$, the objective becomes:

$$p^* \in \arg \min_{p \in \mathcal{F}_m} \left\{ \int F(p, \mathbf{z}) q(\mathbf{z} | \mathbf{x}_*) d\mathbf{z} \right\}. \quad (3)$$

Since we learn about the data generating density $q(\cdot | \mathbf{x}_*)$ by sequentially sampling data $\{\mathbf{z}_t\}_{t=1}^T$ from it, we can approximate the objective function by its sample average:

$$\min_{p \in \mathcal{F}_m} \left\{ \frac{1}{T} \sum_{t=1}^T F(p, \mathbf{z}_t) \right\}.$$

To accommodate sequential observations and achieve online inference, we consider the Stochastic Mirror Descent (SMD) algorithm. The SMD algorithm is a generalization of stochastic gradient descent, using noisy gradient estimates and a decaying weight sequence $\{\alpha_t\}$ to solve stochastic optimization problems, such as (3); see [14] for details:

$$p_{t+1} \in \arg \min_{p \in \mathcal{F}_m} \{ \alpha_t \langle p, -\log q(\mathbf{z}_t | \mathbf{x}) \rangle + \text{KL}(p \| p_t) \}. \quad (4)$$

As shown in [17], (4) has a closed-form solution which resembles a Bayesian update:

$$p_{t+1}(\mathbf{x}) \propto q(\mathbf{z}_t | \mathbf{x})^{\alpha_t} p_t(\mathbf{x}), \quad (5)$$

where \propto indicates proportionality and the weight α_t balances the contributions of the likelihood and the prior terms.

Now, we consider a distributed-problem counterpart, where each agent minimizes a portion of the centralized objective. The independence of observation models in Assumption 2 with $q(z_t|\mathbf{x}) = \prod_{i \in \mathcal{V}} q_i(z_{i,t}|\mathbf{x})$ decomposes the centralized objective into a separable problem based on agents' models,

$$F(p, \mathbf{z}_t) = \sum_{i \in \mathcal{V}} F_i(p, \mathbf{z}_{i,t}), \quad (6)$$

where, $F_i(p, \mathbf{z}_{i,t}) \equiv \mathbb{E}_{\mathbf{x} \sim p}[-\log(q_i(z_{i,t}|\mathbf{x}))]$. Similarly to the centralized objective, the gradient of agent i 's local linear objective $\frac{\delta F_i}{\delta p}(p, \mathbf{z}_{i,t})$ is the log likelihood sample $-\log q_i(z_{i,t}|\mathbf{x})$. Apart from minimizing their objective, the agents need to maintain consensus across the network by achieving equality among their estimates, i.e., $p_i = p_j$, for $i, j \in \mathcal{V}$. To enable consensus, we introduce KL-divergence terms in the SMD algorithm in (4), regularizing the deviation of agent i 's pdf p from the neighbors' prior densities:

$$p_{i,t+1} \in \arg \min_{p \in \mathcal{F}_m} \{ \alpha \langle p, -\log q_i(z_{i,t}|\mathbf{x}) \rangle + \sum_{j \in \mathcal{V}} A_{t,ij} \text{KL}(p||p_{j,t}) \},$$

where $A_{t,ij}$ are the network-dependent coefficients from Assumption 1 and the step size $\alpha_t = \alpha > 0$ is constant. Our analysis in Sec. IV guarantees the convergence of this formulation despite the constant step size and, hence, has faster convergence rate than updates with decaying weights. A minimizer follows from equating the gradient of the right-hand side to zero, which leads to $1 + \log p_{i,t+1} = \alpha \log q_i(z_{i,t}|\mathbf{x}) + \sum_{j \in \mathcal{V}} A_{t,ij} \log(p_{j,t})$. Thus, the update rule for agent i is,

$$p_{i,t+1}(\mathbf{x}) \propto q_i(z_{i,t}|\mathbf{x})^\alpha \prod_{j=1}^n p_{j,t}^{A_{t,ij}}(\mathbf{x}), \quad (7)$$

where α captures the relative importance between the agent's private observation and a geometric average of the neighbors' priors. For pdfs, geometric averaging is advantageous to algebraic averaging (e.g., in [16]) because the average is less dispersed and captures the component pdf modes [1]. A similar algorithm in [13] uses geometric averaging for distributed estimation over a finite discrete space. The following section analyzes the proposed algorithm in a continuous domain.

IV. LARGE DEVIATION ANALYSIS

This section analyzes the pointwise convergence and the mode of the estimated pdf. We extend the large deviation analysis in [9] to time-varying networks for estimating continuous space pdfs with possibly infinite support. Typical proof techniques compare the estimated probability at an arbitrary value to the optimal \mathbf{x}_* , leading to log-probability and log-likelihood ratios,

$$r_{i,t}(\mathbf{x}) = \log \left[\frac{p_{i,t}(\mathbf{x})}{p_{i,t}(\mathbf{x}_*)} \right], \quad g_{i,t}(\mathbf{x}) = \log \left[\frac{q_i(z_{i,t}|\mathbf{x})}{q_i(z_{i,t}|\mathbf{x}_*)} \right].$$

This characterization has two benefits: first, the normalization factor of (7) is simplified. Second, the update rule (7) becomes

linear in terms of the log likelihoods $g_{i,t}$,

$$\mathbf{r}_{t+1}(\mathbf{x}) = A_{t:0} \mathbf{r}_0(\mathbf{x}) + \alpha \sum_{k=1}^t A_{t:k} \mathbf{g}_k(\mathbf{x}). \quad (8)$$

Here, we use the shorthand matrix-product notation $A_{t:k} = A_t \dots A_k$, the stacked vector of log probability ratios $\mathbf{r}_t = [\dots r_{i,t}(\mathbf{x}) \dots]^\top$, and the log likelihood ratios $\mathbf{g}_t = [\dots g_{i,t}(\mathbf{x}) \dots]^\top$. Since each communication matrix A_t is row stochastic, there exists a vector sequence as follows:

Definition 1 (Absolute probability sequence, [19, Def. 1]). *For any sequence of row-stochastic matrices $\{A_t\}$, an absolute probability sequence is a sequence of stochastic vectors $\{\phi(t)\}$ satisfying $\phi(t)^\top = \phi(t+1)^\top A_t$, for all t .*

The vector $\phi(k)$ is related to the point of convergence for the matrix product $A_{t:k}$ with known rate of convergence as:

Lemma 1 (C -step contraction, [12, Lemmas 2, 4]). *Assume that \mathcal{G}_t is C -connected. Then for each time $\bar{t} \geq 0$, there exists a stochastic vector $\phi(\bar{t})$ such that for all $i, j \in \mathcal{V}$ and $t \geq \bar{t}$,*

- 1) $|[A_t \dots A_{\bar{t}}]_{ij} - \phi_j(\bar{t})| \leq 2\lambda^{t-\bar{t}}$, and,
- 2) $\phi_j(\bar{t}) \geq \delta > 0$,

with $\lambda = (1 - \frac{1}{n^{nC}})^{1/C} \in (0, 1)$ and $\delta = 1/n^{nC+1}$.

The initial log-probability ratio exists only if the agents' prior pdf is positive on the optimal parameter space \mathcal{X}_* .

Assumption 3 (Positive priors). *The agents' prior pdf is positive $p_{i,0}(\mathbf{x}_*) > 0$ at the optimal values $\mathbf{x}_* \in \mathcal{X}_*$.*

We now show that the prevalent assumption $\|g_{i,t}\| < L < \infty, \forall i \in \mathcal{V}$ does not hold for continuous space densities. To this end, let us assume that agent-observation models are given as $\pi_i(z|\mu_i, 1)$, modeling the sampled observation z from a Gaussian with mean μ_i and unit variance. As the log-likelihood ratio is linearly dependent on z , it is unbounded,

$$\log(\pi_1(z)/\pi_2(z)) = 2z(\mu_1 - \mu_2) + (\mu_2^2 - \mu_1^2).$$

We instead rely on *moment generating functions* (mgf) to bypass the boundedness assumption. The mgf for a random variable X w.r.t. its pdf p_X , is the function $\psi(b) = \mathbb{E}[\exp(bX)]$, for $b \in \mathbb{R}$. If the observations are sampled from $\pi_*(z|\mu_*, 1)$, the mgf of the random variable defined by the log-ratio $g_{12}(z) = \log(\frac{\pi_1(z)}{\pi_2(z)})$ w.r.t. pdf $\pi_*(z)$ is bounded,

$$\begin{aligned} \mathbb{E}[\exp(bg_{12}(z))] &= \int_z \frac{e^{-(z-\mu_*)^2}}{\sqrt{2\pi}} e^{2bz(\mu_1-\mu_2)+b(\mu_2^2-\mu_1^2)} dz, \\ &= c \int \exp(-z^2 + 2z(\mu_* + b(\mu_1 - \mu_2))) dz < \infty. \end{aligned}$$

Assumption 4 (Finite mgf). *The mgf of the log-likelihood ratio $g_{i,t}(\mathbf{x})$ is finite for any $\mathbf{x} \in \mathcal{X}$ and agents $i \in \mathcal{V}$.*

Large-deviation analysis has been used in conjunction with mgfs to characterize exponentially decreasing bounds on rare events. A general application of mgfs is the following Cramer's theorem [4]. This result upper bounds the deviation of the sum of i.i.d. variables $S_t = X_1 + \dots + X_t$ from their mean with a probability converging to 1 at an exponential rate.

Lemma 2 (Cramer's theorem [4]). *Assume that the mgf $\psi(b)$ of a random variable X_t is finite for some $b > 0$ and let $\mu = \mathbb{E}[X_t]$. Then, for any $a > \mu$, $\mathbb{P}(S_t > at) \leq \exp(-tI(a))$, where $I(a) = \sup_{b>0}\{ab - \log(\psi(b))\} > 0$.*

Cramer's theorem was employed in [9] to prove convergence of distributed estimation in static networks to account for observation models with unbounded support. The theorem cannot be directly applied to time-varying networks as the weighted sum of observations are not i.i.d. Therefore, we develop a theorem with similar guarantees on a sequence of independent random variables. Fix any agent $i \in \mathcal{V}$, then from the log-linear ratio update in (8), we define,

$$e_0 = [A_{t:0}\mathbf{r}_0]_i, e_k = \alpha[A_{t:k}\mathbf{g}_k]_i, \psi_k(b) = \mathbb{E}[\exp(be_k)]. \quad (9)$$

The terms e_k are independent but not i.i.d. Thus, we define a function similar to $I(a)$ in Lemma 2 as follows,

$$J_t(a) = \sup_{b>0} \left(D_t(a, b) \equiv ab - \frac{1}{t} \sum_{k=0}^t \log(\psi_k(b)) \right) \quad (10)$$

In the following theorem, large deviations are used to show that the ratio of estimated probability at $\mathbf{x} \notin \mathcal{X}^*$ and \mathbf{x}_* converges to zero at an exponential rate with probability 1.

Theorem 1. *Let Assumptions 1-4 hold. For each $\mathbf{x} \notin \mathcal{X}_*$, $\mathbf{x}_* \in \mathcal{X}_*$, there is a $t_0 \in \mathbb{N}$ s.t. $\forall t \geq t_0$, $p_{i,t}$ in (7) satisfies,*

$$\mathbb{P}\left(\frac{p_{i,t}(\mathbf{x})}{p_{i,t}(\mathbf{x}_*)} > \exp(\bar{a}(\mathbf{x}, \mathbf{x}_*)t)\right) \leq \exp(-tJ_{t_0}(\bar{a}(\mathbf{x}, \mathbf{x}_*))).$$

The exponential rate of convergence $\bar{a}(\mathbf{x}, \mathbf{x}_*) = -c\delta\|H(\mathbf{x}, \mathbf{x}_*)\|_1 < 0$ is defined via the bound $\delta \in (0, 1)$ from Lemma 1 and sum of KL-divergence terms $\|H(\mathbf{x}, \mathbf{x}_*)\|_1 = \sum_{j \in \mathcal{V}} \text{KL}(\mathbf{q}_i(\cdot|\mathbf{x}_*)\|\mathbf{q}_i(\cdot|\mathbf{x}))$. Any choice of $c \in (0, 1)$ ensures $J_{t_0}(\bar{a}(\mathbf{x}, \mathbf{x}_*))$ is positive.

Proof. Fix $\mathbf{x} \notin \mathcal{X}_*$, and use shorthand notation $\mathbf{r} \equiv \mathbf{r}(\mathbf{x})$, $\mathbf{g} \equiv \mathbf{g}(\mathbf{x})$. Since the terms e_k defined for an arbitrary agent i in (9) are not identically distributed, we work with the running sum $S_t = e_0 + \sum_{k=1}^t e_k$. For any $a \in \mathbb{R}$ and $b > 0$,

$$\begin{aligned} \mathbb{P}(S_t > at) &= \mathbb{P}(\exp(bS_t - bat) > 1), \\ &\leq \mathbb{E}[\exp(bS_t - bat)], \quad (\text{Markov's inequality}) \\ &= \exp(-bat)\mathbb{E}\left[\prod_{k=0}^t \exp(be_k)\right], \\ &= \exp(-bat)\prod_{k=0}^t \mathbb{E}[\exp(be_k)]. \quad (\text{Independence}) \end{aligned}$$

From Assumptions 1 and 4, we know that the mgf $\psi_k(b)$ exists and satisfies $\psi_k(0) = 1$. Using $D_t(a, b)$ from (10), the preceding inequality is equivalent to $\mathbb{P}(S_t > at) \leq \exp(-tD_t(a, b))$. Since this holds for any a, b , we have,

$$\mathbb{P}(S_t > at) \leq \exp(-tJ_t(a)). \quad (11)$$

Now, we prove the existence of $J_t(a) > 0$ for some choice of $a \equiv \bar{a}(\mathbf{x}, \mathbf{x}_*) < 0$ and all $t \geq t_0 \equiv t_0(\mathbf{x}, \mathbf{x}_*) > 0$. If $D_t(a, 0) = 0$ and the term $\frac{dD_t}{db}(a, b)|_{b=0}$ is positive, then there exists $b > 0$ for which $D_t(a, b) > 0$ for all a and $t > t_0$.

We notice that $D_t(a, 0) = 0$ and $D_t(a, b)$ is finite for any $b > 0$, its derivative at $b = 0$ is,

$$\frac{dD_t}{db}\bigg|_{b=0} = a - \frac{1}{t} \sum_{k=0}^t \frac{\psi'_k(b)}{\psi_k(b)}\bigg|_{b=0} = a - \frac{1}{t} \sum_{k=0}^t \mathbb{E}[e_k].$$

We will show the existence of a time t_0 such that the running average $\frac{1}{t} \sum_{k=0}^t \mathbb{E}[e_k]$ is bounded above by some $a < 0$ for all times $t > t_0$. Adding and subtracting the expected weighted likelihoods with weights $\phi(k)$ from Lemma 1,

$$e_k = \alpha[(A_{t:k} - \mathbf{1}\phi(k)^\top)\mathbf{g}_k]_i + \alpha\phi(k)^\top\mathbf{g}_k. \quad (12)$$

The expected value of the stochastic average of the log likelihood samples $\phi(k)^\top\mathbf{g}_k$ is evaluated by computing the expectation w.r.t. the true observation model $\mathbf{q}(\mathbf{z}_t|\mathbf{x}_*)$,

$$\begin{aligned} \mathbb{E}[\phi(k)^\top\mathbf{g}_k] &= \int \mathbf{q}(\mathbf{z}_t|\mathbf{x}_*) \sum_{i \in \mathcal{V}} \phi_i(k) \log \left[\frac{\mathbf{q}_i(\mathbf{z}_{i,t}|\mathbf{x})}{\mathbf{q}_i(\mathbf{z}_{i,t}|\mathbf{x}_*)} \right] d\mathbf{z}_t, \\ &= - \sum_{i \in \mathcal{V}} \phi_i(k) \text{KL}(\mathbf{q}_i(\cdot|\mathbf{x}_*)\|\mathbf{q}_i(\cdot|\mathbf{x})) = -\phi(k)^\top H(\mathbf{x}, \mathbf{x}_*). \end{aligned}$$

Since a product of stochastic matrices remains stochastic, we can upper bound $\mathbb{E}[A_{t:0}\mathbf{r}_0]_i \leq |\mathbf{r}_0(\mathbf{x})|_1$. The matrix product is also independent of the observations, so we can use Lemma 1 to bound the first term of e_k in (12),

$$\begin{aligned} \mathbb{E}[(A_{t:k} - \mathbf{1}\phi(k)^\top)\mathbf{g}_k]_i &= [(A_{t:k} - \mathbf{1}\phi(k)^\top)H(\mathbf{x}, \mathbf{x}_*)]_i \\ &\in (-\lambda^{t-k}\|H(\mathbf{x}, \mathbf{x}_*)\|_1, \lambda^{t-k}\|H(\mathbf{x}, \mathbf{x}_*)\|_1). \end{aligned} \quad (13)$$

From Lemma 1, the stochastic vector terms satisfy $\phi_i(k) > \delta > 0$ for all agents $i \in \mathcal{V}$ and time $k \geq 1$, implying,

$$\frac{1}{t} \sum_{k=0}^t \mathbb{E}[e_k] \leq \frac{1}{t} |\mathbf{r}_0(\mathbf{x})|_1 + \frac{\alpha}{t} \|H(\mathbf{x}, \mathbf{x}_*)\|_1 \sum_{k=1}^t (\lambda^{t-k} - \delta).$$

The upper bound is strictly negative for all $t > t_0 = \frac{\lambda}{\delta(1-\lambda)} + \frac{|\mathbf{r}_0(\mathbf{x})|_1}{\alpha\delta\|H(\mathbf{x}, \mathbf{x}_*)\|_1}$. The initial time t_0 is same for all agents in \mathcal{V} , and depends on the network characteristics (δ, λ) , the initial probability ratio $|\mathbf{r}_0(\mathbf{x})|_1$ and the divergence sum at $(\mathbf{x}, \mathbf{x}_*)$. We can choose $\bar{a}(\mathbf{x}, \mathbf{x}_*)$ using any $c \in (0, 1)$ as,

$$\bar{a}(\mathbf{x}, \mathbf{x}_*) = -c\delta\|H(\mathbf{x}, \mathbf{x}_*)\|_1 < 0. \quad (14)$$

Given the decreasing upper bound on the term $\frac{1}{t} \sum_{k=1}^t \mathbb{E}[e_k]$, this choice for $a = \bar{a}(\mathbf{x}, \mathbf{x}_*)$ implies that $\frac{dJ_t}{db}|_{b=0} > 0$ for all $t > t_0$. Upon choosing $\bar{a}(\mathbf{x}, \mathbf{x}_*)$ and J_{t_0} in (11), we show that the probability of log probability ratio exceeding a linearly decreasing value diminishes exponentially at \mathbf{x} ,

$$\mathbb{P}(r_{i,t+1}(\mathbf{x}) > \bar{a}(\mathbf{x}, \mathbf{x}_*)t) \leq \exp(-tJ_{t_0}(\bar{a}(\mathbf{x}, \mathbf{x}_*))). \quad \blacksquare$$

Remark 1. *If the algorithm weighs the likelihood terms by square-summable α_k in (7), then our analysis does not guarantee convergence of the probability ratio in (8) to zero.*

Remark 2. *As per our analysis, higher sum of divergence $\|H(\mathbf{x}, \mathbf{x}_*)\|_1$ implies higher rate of convergence $\bar{a}(\mathbf{x}, \mathbf{x}_*)$ and lower starting time t_0 , meaning that the estimate starts converging sooner and at a faster rate.*

This probabilistic result on convergence of log-ratio probability in Theorem 1 holds over several distinct sequences of

estimates $p_{i,t}$. Therefore, we use Borel-Cantelli Lemma to gain insight into the convergence of an individual sequence.

Definition 2. For a sequence of events $\{E_t\}$, we define the events (a) E_t occurs infinitely often, (i.o.) $\equiv \limsup_{t \rightarrow \infty} E_t \equiv$

$\bigcap_{t=0}^{\infty} \bigcup_{k=t}^{\infty} E_k$, and (b) E_t occurs eventually, (e.v.) $\equiv \liminf_{t \rightarrow \infty} E_t \equiv$
 $\bigcup_{t=0}^{\infty} \bigcap_{k=t}^{\infty} E_k$. Also, we have E_t i.o. $= (E_t^c \text{ e.v.})^c$.

Lemma 3 (Borel-Cantelli Lemma [3]). For any event sequence $\{E_t\}_{t=1}^{\infty}$,

- 1) if $\sum_{t=1}^{\infty} \mathbb{P}(E_t) < \infty$, then $\mathbb{P}(E_t \text{ i.o.}) = 0$, and
- 2) if $\sum_{t=1}^{\infty} \mathbb{P}(E_t) = \infty$ and the event sequence $\{E_t\}$ is independent, then $\mathbb{P}(E_t \text{ e.v.}) = 1$.

Proposition 1. As $t \rightarrow \infty$, a mode of the pdf $p_{i,t}(\mathbf{x})$ estimated by agent i almost surely lies in the set of optimal parameters \mathcal{X}_* as defined in (1).

Proof. We proceed by contradiction. For an arbitrary agent i , suppose that all modes of $\lim_{t \rightarrow \infty} p_{i,t}(\mathbf{x})$ almost surely lie outside of \mathcal{X}_* . Hence, for any $\delta_0 > 0$, there exists t_1 such that any mode \mathbf{x}_1 almost surely satisfies $p_{i,t}(\mathbf{x}_1) > p_{i,t}(\mathbf{x}_*) + \delta_0$ for all $\mathbf{x}_* \in \mathcal{X}_*$ and all $t \geq t_1$. We show that this assumption is contradicted by the fact that $|p_{i,t}(\mathbf{x}_1) - p_{i,t}(\mathbf{x}_*) \exp(\bar{a}(\mathbf{x}_1, \mathbf{x}_*)t)| \xrightarrow{\text{a.s.}} 0$ as established in Theorem 2.

Two random sequences Y_t and Z_t satisfy $|Y_t - Z_t| \xrightarrow{\text{a.s.}} 0$ if and only if $\forall \epsilon \geq 0$, $\mathbb{P}[|Y_t - Z_t| \leq \epsilon \text{ e.v.}] = 1$, which holds if and only if $\forall \epsilon \geq 0$, $\mathbb{P}[|Y_t - Z_t| > \epsilon \text{ i.o.}] = 0$.

Let $E_t(\epsilon)$ denote the event that $|p_{i,t}(\mathbf{x}_1) - p_{i,t}(\mathbf{x}_*) \exp(\bar{a}(\mathbf{x}_1, \mathbf{x}_*)t)| > \epsilon$. From Theorem 1, for $t > t_0$, $\mathbb{P}[E_t(\epsilon)] \leq \mathbb{P}[E_t(0)] \leq \exp(-tJ_t(a))$ with $J_t(a) > 0$. Since $\sum_{t=1}^{\infty} \exp(-tJ_t(a)) < \infty$, we have that $\sum_{t=1}^{\infty} \mathbb{P}[E_t(\epsilon)] < \infty$. By the Borel-Cantelli Lemma, this implies that $\mathbb{P}[E_t(\epsilon) \text{ i.o.}] = 0$ for all $\epsilon \geq 0$, and, hence, $|Y_t - Z_t| \rightarrow 0$ a.s. In other words,

$$\mathbb{P}[\lim_{t \rightarrow \infty} |p_{i,t}(\mathbf{x}_1) - p_{i,t}(\mathbf{x}_*) \exp(\bar{a}(\mathbf{x}_1, \mathbf{x}_*)t)| = 0] = 1. \quad (15)$$

The above result implies that there exists t_2 such that almost surely $|p_{i,t}(\mathbf{x}_1) - p_{i,t}(\mathbf{x}_*) \exp(\bar{a}(\mathbf{x}_1, \mathbf{x}_*)t)| \leq \delta_0$ for all $t \geq t_2$. Since $\bar{a}(\mathbf{x}_1, \mathbf{x}_*) < 0$, $\exp(\bar{a}(\mathbf{x}_1, \mathbf{x}_*)t) \leq 1$ and we have:

$$p_{i,t}(\mathbf{x}_1) - p_{i,t}(\mathbf{x}_*) \leq p_{i,t}(\mathbf{x}_1) - p_{i,t}(\mathbf{x}_*) \exp(\bar{a}(\mathbf{x}_1, \mathbf{x}_*)t) \leq \delta_0.$$

However, for any $t \geq \max\{t_1, t_2\}$ the above result contradicts with the assumption that $p_{i,t}(\mathbf{x}_1) > p_{i,t}(\mathbf{x}_*) + \delta_0$. \blacksquare

Corollary 1 (Uniqueness). If the optimal hypothesis set \mathcal{X}_* is globally identifiable i.e. $\mathcal{X}_* = \{\mathbf{x}_*\}$, then the unique mode of the estimated pdf almost surely lies at \mathbf{x}_* .

Proof. The claim follows from eventually almost sure existence of mode in pdf estimates in Proposition 1. \blacksquare

Corollary 2 (Discrete probabilities). If the estimated probability density $p_{i,t}$ is bounded above by some $\gamma > 0$ as is the case for probability mass functions, then the probability estimated at any $\mathbf{x}_1 \in \mathcal{X} \setminus \mathcal{X}_*$ satisfy, $p_{i,t}(\mathbf{x}_1) \rightarrow 0$ a.s.

Proof. From Theorem 2, we know that $|p_{i,t}(\mathbf{x}_1) - p_{i,t}(\mathbf{x}_*) \exp(\bar{a}(\mathbf{x}_1, \mathbf{x}_*)t)| \xrightarrow{\text{a.s.}} 0$. With the property $p_{i,t}(\mathbf{x}) < \gamma$ for all $\mathbf{x} \in \mathcal{X}$ and our choice $\bar{a}(\mathbf{x}_1, \mathbf{x}_*) < 0$, there exists some

$t_2 > 0$ such that we have for any arbitrary $\delta_0 > 0$ and $t > t_2$, $p_{i,t}(\mathbf{x}_1) \leq \gamma \exp(\bar{a}(\mathbf{x}_1, \mathbf{x}_*)t) + \delta_0$ almost surely. With $t \rightarrow \infty$ and arbitrary $\delta_0 > 0$, $p_{i,t}(\mathbf{x}_1) \rightarrow 0$ a.s. \blacksquare

V. EVALUATION

This section demonstrates the proposed algorithm in two examples: cooperative localization and target tracking. In the first example, the sensing agents measure relative position to their one-hop neighbors to infer sensor locations. For a linear observation model w.r.t. the agent positions, a Gaussian algorithm is derived. In the second example, a sensors in a C -connected network apply particle version of our algorithm on non-linear range measurements to track a moving target.

Cooperative localization: Consider $n = 10$ sensors positioned at $\{\mathbf{x}_i \in \mathbb{R}^2, i \in \mathcal{V}\}$. Sensor 1 is an anchor with known position $\mathbf{x}_1 = [0, 0]^\top$, while the positions \mathbf{x}_i of the remaining sensors are unknown and need to be estimated. Each sensor i measures the relative position of its neighbors in a static connected measurement graph $(\mathcal{V}, \mathcal{E}^m)$. The relative position measurement $\mathbf{z}_{i,j,t}$ made by sensor i of sensor j at time t follows a Gaussian distribution:

$$\mathbf{z}_{i,j,t} \sim \mathcal{N}(\mathbf{x}_j - \mathbf{x}_i, \mathbb{I}_2), \quad \forall (i, j) \in \mathcal{E}^m. \quad (16)$$

All measurements received by sensor i at time t are:

$$\mathbf{z}_{i,t} \sim \mathcal{N}(H_i \mathbf{x}, (\Omega_i^z)^{-1}), \quad (17)$$

where H_i and $(\Omega_i^z)^{-1}$ are obtained by stacking the expressions in (16). While the sensor receive relative-position measurements $\mathbf{z}_{i,t}$ at every time t , the communication among them is unreliable. The communication network is described by a randomly-generated uniformly-connected graph $\mathcal{G}_t = (\mathcal{V}, \mathcal{E}_t, A_t)$. The updates to the information matrix $\Omega_{i,t}$ and the mean $\mu_{i,t}$ at agent i , derived from (7) are:

$$\begin{aligned} \Omega_{i,t+1} &= \sum_{j \in \mathcal{V}} A_{t,ij} \Omega_{j,t} + \alpha H_i^\top \Omega_i^z H_i, \\ \mu_{i,t+1} &= \Omega_{i,t+1}^{-1} \left(\sum_{j \in \mathcal{V}} A_{t,ij} \Omega_{j,t} \mu_{j,t} + \alpha H_i^\top \Omega_i^z \mathbf{z}_{i,t} \right). \end{aligned} \quad (18)$$

The update equations appear similar to those of a distributed Gaussian filter [15] but differ due to the term α , weighting the effect of the measurement, and due to the time-varying communication weights $A_{t,ij}$. Fig. 1 presents the neighbors observing relative positions, a communication network sample, and the estimated mean of agent 2's position by other agents for increasing values of the likelihood weight α .

Target tracking: Consider the problem of estimating the center $\mathbf{x}_* \in \mathbb{R}^2$ of a circular maneuver of a target. Other fixed parameters defining the target motion are the initial angle $\theta_0 = 0$, radius $r = 1$ and angular velocity $\omega = 0.2$. The target position \mathbf{y}_t^d at time t is,

$$\theta_k = \theta_{k-1} + \omega \Delta t, \quad \mathbf{y}_t^d = \mathbf{x}_* + r[\cos(\theta_t), \sin(\theta_t)]^\top \quad (19)$$

We aim to distribute estimation of variable \mathbf{x} over a network of n range sensors using Time-of-Arrival measurements. The noisy measurements for sensor i located at \mathbf{y}_i^s is,

$$z_{i,t}(\mathbf{y}_i^s, \mathbf{y}_t^d) = |\mathbf{y}_i^s - \mathbf{y}_t^d|_2 + \eta, \quad \eta \sim \mathcal{N}(0, 1). \quad (20)$$

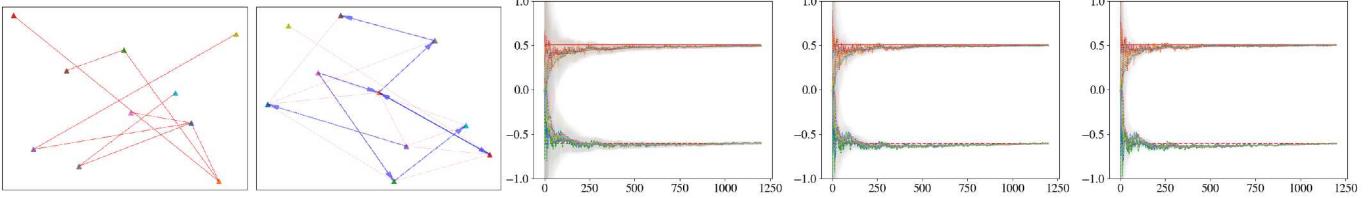


Fig. 1: Cooperative localization in a ten-agent time-varying network. Left to right: observation network, communication network at $t = 2$, estimates of agent 2's position for $\alpha \in \{0.5, 1, 1.5\}$ over 1250 iterations. The horizontal red lines indicate the true position of agent 2.

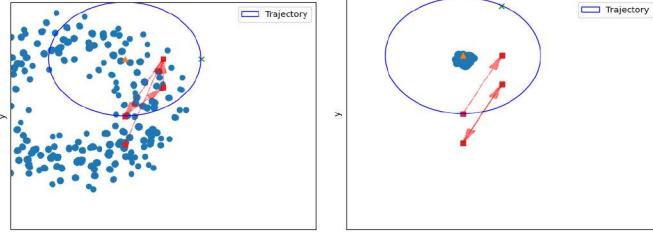


Fig. 2: Estimating the center of a circular trajectory (orange triangle at $[1, 4]$) using a time-varying uniformly connected network of four sensors (red squares at $[1, 1], [1, 2], [2, 4], [2, 3]$). The subfigures show the cooperatively estimated particle-filter distribution of the circle center after 1 (left) and 200 (right) iterations.

Since the observation model is non-linear w.r.t. positions \mathbf{x}_* , we recall particle filter method to represent the estimated pdf as a particle distribution, transforming the intractable integration into a summation. Each sensor i employs M particles $\mathbf{x}_{i,0}^m, m \in \{1, \dots, M\}$ initially sampled using Latin hypercube sampling [10]. The particles are assigned equal weights $\rho_{i,0}^m = 1/M$ initially, resulting in a prior pdf $p_{i,0}(\mathbf{x}_*) = \sum_{m=1}^M \rho_{i,0}^m \delta(\mathbf{x}_* - \mathbf{x}_{i,0}^m)$ on the circle's center. Here, δ is a Dirac-delta function. The likelihood update for each particle based on the observation $\mathbf{z}_{i,t}$ with $\alpha = 1$ is,

$$p_{i,t+1|t}(\mathbf{x}_*) \propto q_i(\mathbf{z}_{i,t}|\mathbf{x}_*) \sum_{m=1}^M \rho_{i,t}^m \delta(\mathbf{x}_* - \mathbf{x}_{i,t}^m),$$

$$\rho_{i,t+1}^m = \left(q_i(\mathbf{z}_{i,t}|\mathbf{x}_{i,t}^m) \rho_{i,t}^m / \sum_{m=1}^M q_i(\mathbf{z}_{i,t}|\mathbf{x}_{i,t}^m) \rho_{i,t}^m \right).$$

In a standard particle filter, methods like stratified resampling [2] sample particles with a frequency corresponding to their weights. To keep the number of particles fixed in our distributed formulation, we modify the resampling method to sample M particles at agent i from the $M|\mathcal{V}_i|$ -particle density $\sum_{j \in \mathcal{V}_i} A_{ij} \sum_{m=1}^M \rho_{j,t}^m \delta(\mathbf{x}_* - \mathbf{x}_{j,t}^m)$. The target trajectory is a circle centered at $\mathbf{x}_0 = [1, 4]$. Fig. 2 presents four sensors sampling distance to the target every half second starting $t = 0$. The circular particle distribution at $t = 1$ is caused by the uni-dimensional range observations. By $t = 200$ steps, the particles coalesce to the actual center \mathbf{x}_* . Fig. 3 shows the particle mean and the log maximum covariance eigenvalue.

VI. CONCLUSION

This paper addresses a distributed estimation problem in which agents need to cooperate over a time-varying network to consistently estimate a quantity of interest. A novel application of large deviation analysis allows estimating parameters in continuous space with an exponential rate of convergence for probability ratios. In case of a globally identifiable problem, we prove that a mode of the estimated density function at each agent coincides with the true value.

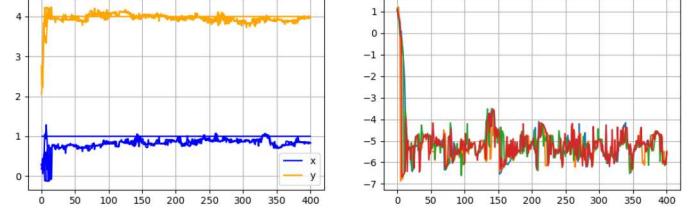


Fig. 3: Evolution of the mean and log-maximum eigenvalue of the covariance of the particle-filter estimates from Fig. 2.

REFERENCES

- [1] S. Bandyopadhyay and S.-J. Chung. Distributed estimation using bayesian consensus filtering. In *Proc. Am. Control Conf.*, pages 634–641, 2014.
- [2] J. Carpenter, P. Clifford, and P. Fearnhead. Improved particle filter for nonlinear problems. *IEE Proc.-Radar, Sonar and Navigation*, 146(1):2–7, 1999.
- [3] K. L. Chung and P. Erdos. On the application of the Borel-Cantelli lemma. *Trans. Am. Math. Soc.*, 72(1):179–186, 1952.
- [4] H. Cramér. Sur un nouveau théorème-limite de la théorie des probabilités. *Actual. Sci. Ind.*, 736:5–23, 1938.
- [5] M. H. DeGroot. Reaching a consensus. *J. Am. Stat. Assoc.*, 69(345):118–121, 1974.
- [6] B. Gharesifard and J. Cortés. When does a digraph admit a doubly stochastic adjacency matrix? In *Proc. Am. Control Conf.*, pages 2440–2445, 2010.
- [7] J. M. Hendrickx and J. N. Tsitsiklis. Fundamental limitations for anonymous distributed systems with broadcast communications. In *Annu. Allert. Conf. Commun. Control Comput.*, pages 9–16, 2015.
- [8] A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi. Non-Bayesian social learning. *Games Econ. Behav.*, 76(1):210–225, 2012.
- [9] A. Lalitha, T. Javidi, and A. D. Sarwate. Social learning and distributed hypothesis testing. *IEEE Trans. Inf. Theory*, 64(9):6161–6179, 2018.
- [10] W.-L. Loh. On latin hypercube sampling. *The annals of statistics*, 24(5):2058–2080, 1996.
- [11] P. Molavi, A. Tahbaz-Salehi, and A. Jadbabaie. Foundations of non-Bayesian social learning. *Columbia Business School Research Paper*, pages 15–95, 2017.
- [12] A. Nedić and A. Olshevsky. Distributed optimization over time-varying directed graphs. *IEEE Trans. Autom. Control*, 60(3):601–615, 2014.
- [13] A. Nedić, A. Olshevsky, and C. A. Uribe. Nonasymptotic convergence rates for cooperative learning over time-varying directed graphs. In *Proc. Am. Control Conf.*, pages 5884–5889, 2015.
- [14] A. Nemirovski. Tutorial: Mirror descent algorithms for large-scale deterministic and stochastic convex optimization. Citeseer, 2012.
- [15] R. Olfati-Saber. Distributed kalman filtering for sensor networks. In *IEEE Conference on Decision and Control*, pages 5492–5498, 2007.
- [16] R. Parasnis, M. Franceschetti, and B. Touri. Uniform strong connectivity is not necessary for non-bayesian social learning on time-varying directed graphs. In *Proc. IEEE Conf. Decis. Control*, pages 4933–4940, 2020.
- [17] P. Paritosh, N. Atanasov, and S. Martinez. Marginal density averaging for distributed node localization from local edge measurements. In *Proc. IEEE Conf. Decis. Control*, pages 2404–2410, 2020.
- [18] K. R. Rad and A. Tahbaz-Salehi. Distributed parameter estimation in networks. In *IEEE Conf. on Decision and Control*, 2010.
- [19] F. Saadatnaki, R. Xin, and U. A. Khan. Decentralized optimization over time-varying directed graphs with row and column-stochastic matrices. *IEEE Trans. Autom. Control*, 65(11):4769–4780, 2020.