Distributed Detection: Finite-Time Analysis and Impact of Network Topology

Shahin Shahrampour, Alexander Rakhlin, and Ali Jadbabaie, *Fellow, IEEE*

***Abstract*—This paper addresses the problem of distributed detection in multi-agent networks. Agents receive private signals about an *unknown* state of the world. The underlying state is globally identifiable, yet informative signals may be dispersed throughout the network. Using an optimization-based framework, we develop an iterative local strategy for updating individual beliefs. In contrast to the existing literature which focuses on asymptotic learning, we provide a *finite-time* analysis. Furthermore, we introduce a Kullback-Leibler cost to compare the efficiency of the algorithm to its centralized counterpart. Our bounds on the cost are expressed in terms of network size, spectral gap, centrality of each agent and relative entropy of agents' signal structures. A key observation is that distributing more informative signals to central agents results in a faster learning rate. Furthermore, optimizing the weights, we can speed up learning by improving the spectral gap. We also quantify the effect of link failures on learning speed in symmetric networks. We finally provide numerical simulations for our method which verify our theoretical results.**

***Index Terms*—Kullback-Leibler cost.**

I. Introduction

Recent years have witnessed an intense interest on *distributed* detection, estimation, prediction and optimization [1]–[7]. Decentralizing the computation burden among agents has been widely regarded in networks ranging from sensor and robot to social and economic networks [8]–[11]. In this broad class of problems, agents in a network need to perform a global task for which they only have *partial* information. Therefore, they recursively exchange information with their neighbors, and the global dispersion of information in the network provides them with adequate data to accomplish the task. In the big picture, many of these schemes can also be embedded in the context of *consensus* protocols which have gained a growing popularity over the past three decades [12]–[14].

Earlier works on decentralized detection have considered scenarios where each agent sends its observations to a fusion center that decides over the true value of a parameter [1], [2], [8]. In these situations, the fusion center faces a classical hypothesis testing (centralized detection) problem after collecting the data from agents. Distributed detection has been widely regarded in various works providing the asymptotic analysis. Cattivelli *et al.* [15] propose a fully distributed algorithm where no fusion center is necessary. The methodology builds on the connection of Neyman-Pearson detection and minimum-variance estimation to solve the problem. Jakovetić *et al.* [16] develop a consensus+innovations algorithm for detection under *Gaussian* observations. The method achieves an *asymptotic* exponential error rate even when communications of agents are noisy. In [17], the authors extend the consensus+innovations method to *generic* (non-Gaussian) observations over random networks. More recently, another model of learning and detection has been proposed by Jadbabaie *et al.* [18]. In this framework, the world is governed by a fixed true *state* or *hypothesis* that is aimed to be recovered by a network of agents. The state belongs to a *finite* set, and might represent a decision, an opinion, the price of a product or any quantity of interest. Each agent observes a stream of *private* signals generated by a marginal of the global likelihood *conditioned* on the true state. However, the signals might *not* be informative enough for the agent to distinguish the underlying state of the world. Therefore, agents use *local* diffusion to compensate for their imperfect knowledge about the environment. In the literature, a host of schemes build on this model to describe distributed learning [18]–[22]. Despite the wealth of results on the asymptotic behavior of these methods, the *finite-time* analysis remains elusive. Though appealing in certain cases, asymptotic analysis might not unveil all important factors for learning. Realistically, one always has finite time to make a decision; hence, studying non-asymptotic aspects of learning is an interesting complementary direction. For instance, let us think of a social network where individuals need to choose a product which best suits the network. Individuals might value the product differently, and they need to reach consensus in a few rounds of opinion exchange. Agents do not have an infinite horizon to make a decision; therefore, one needs to view this scenario as a finite-time problem.

We now elaborate on several works inspired by the model considered in [18]. The authors in [18] propose a non-Bayesian update rule in the context of social networks. Each individual averages her Bayesian posterior belief with the opinion of her neighbors. It is then shown that, under mild technical assumptions, agents' beliefs converge to the true state almost surely. Following up on the work of Duchi *et al.* [23] on distributed dual averaging, an optimization-based algorithm is

Manuscript received March 18, 2015; revised August 11, 2015 and November 27, 2015; accepted December 3, 2015. Date of publication December 8, 2015; date of current version October 25, 2016. This work was supported by ONR BRC under Grant N00014-12-1-0997. Recommended by Associate Editor M. Verhaegen.

S. Shahrampour is with the Department of Electrical Engineering, Harvard University, Cambridge, MA 02138 USA (e-mail: shahin@seas.harvard.edu).

A. Rakhlin is with the Department of Statistics, University of Pennsylvania, Philadelphia, PA 19104 USA (e-mail: rakhlin@wharton.upenn.edu).

A. Jadbabaie is with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104 USA (e-mail: jadbabai@seas.upenn.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TAC.2015.2506903

developed in [19]. The authors propose an update rule which is the solution of a distributed stochastic optimization. They demonstrate that the belief sequence is weakly consistent when agents use *gossip* communication protocol. A communication-efficient variant of the problem is studied in [20] where agents switch between Bayesian and non-Bayesian regimes to asymptotically learn the true state. Lalitha *et al.* [21] introduce another strategy where agents perform a local Bayesian update, and geometrically average the posteriors in their neighborhood. The authors then provide the convergence and rate analysis of their method. On the other hand, Rahnama Rad *et al.* [22] present a distributed estimation algorithm for *continuous* state space. They prove the convergence of the algorithm, and characterize the asymptotic efficiency of the method in comparison to any centralized estimator. In [18]–[21], the convergence occurs exponentially fast, and the *asymptotic* rate is characterized in terms of the *relative entropy* of individuals' signal structures and their *eigenvector centralities* (see [24] for the rate analysis of [18]). As an important consequence, the rate in [19] only recovers the *empirical* average of relative entropies since the method is restricted to undirected networks.

The asymptotic analysis presented in the above-discussed papers only describes the dominant factors that influence learning in the long run. In real world applications, however, the *decision* on the true state has to be made in a *finite* time. Therefore, it is crucial to study the finite-time variant of these schemes to gain insight into the interplay of *network parameters* which affect learning. To this end, we extend the work of Shahrampour *et al.* [19] to directed networks where agents are *not* equally central. Moreover, we introduce the notion of *Kullback-Leibler (KL)* cost to measure the learning rate of an individual agent versus an *expert* who has all available information for learning. The KL decentralization cost simply compares the performance of distributed algorithm to its centralized counterpart. We derive an upper bound on the cost which proves the *spectral gap* of the network is substantial beside agents' centralities. It turns out that the upper bound scales inversely in the spectral gap, and logarithmically with the network size, number of states and time horizon. The rate also scales with the inverse of the relative entropy of the conditional marginals. More specifically, the KL cost grows when signals do not provide enough evidence in favor of the true state versus some other state of the world.

Assuming that the network is realized with a *default* communication structure, each agent is endowed with a centrality. We establish that allocating more informative signals to more central agents can expedite learning. More interestingly, the importance of spectral gap opens new venues for *optimal* network design to facilitate agents' interactions. Each agent assigns different weights to its neighbors' information while communicating with them. We demonstrate how agents can modify these weights to achieve a faster learning rate. The key idea is to find the Markov chain with the best mixing behavior that is consistent with the network structure and agents' centralities. On the other hand, as a natural conjecture, we expect a more rapid learning rate in well-connected networks. We study the ramification of *link failures* in the network, and prove that in symmetric networks, less connectivity amounts to a sluggish learning process. We further apply our results on star, cycle and two-dimensional grid network, and observe that in each case the effect of spectral gap can be translated to the network *diameter*. Intuitively, a larger diameter makes the information propagation difficult around the network. Finally, we present numerical experiments which perfectly match our theoretical findings.

The rest of the paper is organized as follows: we describe the formal statement of the problem, and flesh out the distributed detection scheme in Section II. Section III is devoted to the finite-time analysis of the algorithm, whereas Section IV elaborates on the impact of network characteristics on the convergence rate. We discuss briefly about applications of the model, and provide our numerical experiments in Section V. Section VI concludes.

Notation: We adhere to the following notation in the exposition of our results:

$[n]$	The set $\{1, 2, ..., n\}$ for any integer n
$x^\top$	Transpose of the vector x
$x(k)$	The k-th element of vector x
$x_{[k]}$	The k-th largest element of vector x
I_m	Identity matrix of size m
Δ_m	The m-dimensional probability simplex
$\mathbf{e}_k$	Delta distribution on k-th component
$\langle \cdot, \cdot \rangle$	Standard inner product operator
$\|\cdot\|_p$	p-norm operator
$\mathbb{1}$	Vector of all ones
$\|\mu - \pi\|_{\text{TV}}$	Total variation distance between $\mu, \pi \in \Delta_m$
$D_{KL}(\mu\|\pi)$	KL-divergence of $\pi \in \Delta_m$ from $\mu \in \Delta_m$
$\lambda_i(W)$	The i-th largest eigenvalue of matrix W

For any $f \in \mathbb{R}^m$ and $\mu \in \Delta_m$, we let $\mathbb{E}_\mu[\cdot]$ represent the expectation of f under the measure μ, i.e., we have $\mathbb{E}_\mu[f] = \sum_{j=1}^m \mu(j) f(j)$. Throughout, all the vectors are assumed to be column vectors.

II. The Problem Description and Algorithm

In this section, we describe the observation and network model, and outline the centralized setting for the problem. Then, we provide a formal statement of the distributed setting, and characterize the decentralization cost.

A. Observation Model

The signal and observation model of this work closely follows the framework proposed in [18]. We consider an environment in which $\Theta = \{\theta_1, \theta_2, \ldots, \theta_m\}$ denotes a finite set of *states* of the world. We have a network of n agents that seek the *unique*, true state of the world $\theta_1 \in \Theta$. At each time $t \in [T]$, the belief of agent i is denoted by $\mu_{i,t} \in \Delta_m$, where Δ_m is a probability distribution over the set Θ. In particular, $\mu_{i,0} \in \Delta_m$ denotes the prior belief of agent $i \in [n]$ about the states of the world, and it is assumed to be uniform.[1]

[1]The assumption of uniform prior only avoids notational clutter. The analysis in the paper holds for any prior with full support.

The learning model is given by a conditional likelihood function $\ell(\cdot|\theta_k)$ which is governed by a state of the world $\theta_k \in \Theta$. For each $i \in [n]$, let $\ell_i(\cdot|\theta_k)$ denote the i-th marginal of $\ell(\cdot|\theta_k)$, and we use the vector representation $\ell_i(\cdot|\theta) = [\ell_i(\cdot|\theta_1), \dots, \ell_i(\cdot|\theta_m)]^\top$ to stack all states. At each time $t \in [T]$, the signal $s_t = (s_{1,t}, s_{2,t}, \dots, s_{n,t}) \in \mathcal{S}_1 \times \cdots \times \mathcal{S}_n$ is generated based on the true state θ_1. Therefore, for each $i \in [n]$, the signal $s_{i,t} \in \mathcal{S}_i$ is a sample drawn according to the likelihood $\ell(\cdot|\theta_1)$ where $\mathcal{S}_i$ is the sample space.

The signals are i.i.d. over time, and also the marginals are independent, i.e., $\ell(\cdot|\theta_k) = \Pi_{i=1}^n \ell_i(\cdot|\theta_k)$ for any $k \in [m]$. For the sake of convenience, we define $\psi_{i,t} := \log \ell_i(s_{i,t}|\theta)$ which is a sample corresponding to $\Psi_i := \log \ell_i(\cdot|\theta)$ for any $i \in [n]$.

A1. We assume that all log-marginals are uniformly bounded such that $\|\psi_{i,t}\|_\infty \le B$ for any $s_{i,t} \in \mathcal{S}_i$, i.e., we have $|\log \ell_i(\cdot|\theta_k)| \le B$ for any $i \in [n]$ and $k \in [m]$.

Based on assumption **A1**, every private signal has bounded information content. The assumption can also be interpreted as Radon-Nikodym derivative of every private signal (likelihood ratio) being bounded [25]. This bound can be found, for instance, when the signal space is discrete and provides a full support for distribution. Let us define $\bar{\Theta}_i$ as the set of states that are observationally equivalent to θ_1 for agent $i \in [n]$; in other words, $\bar{\Theta}_i = \{\theta_k \in \Theta : \ell_i(s_i|\theta_k) = \ell_i(s_i|\theta_1)\ \forall\, s_i \in \mathcal{S}_i\}$ almost surely with respect to the signal space.[2] As evident from the definition, any state $\theta_k \neq \theta_1$ in the set $\bar{\Theta}_i$ is not distinguishable from the true state by observation of samples from the i-th marginal. Let $\bar{\Theta} = \cap_{i=1}^n \bar{\Theta}_i$ be the set of states that are observationally equivalent to θ_1 from all agents perspective.

A2. We assume that no state in the world is observationally equivalent to the true state from the standpoint of the network, i.e., the true state is globally identifiable, and we have $\bar{\Theta} = \{\theta_1\}$.

Assumption **A2** guarantees that the global likelihood provides sufficient information to make the true state uniquely identifiable. In other words, for any false state $\theta_k \neq \theta_1$, there must exist an agent who is able to distinguish θ_1 from θ_k.

Let $\mathcal{F}_t$ be the smallest σ-field containing the information about all agents up to time t. Then, when the learning process continues for T rounds, the probability triple $(\Omega, \mathcal{F}, \mathbb{P})$ is defined as follows: the sample space $\Omega = \otimes_{t=1}^T (\otimes_{i=1}^n \mathcal{S}_i)$, the σ-field $\mathcal{F} = \cup_{t=1}^T \mathcal{F}_t$, and the true probability measure $\mathbb{P} = \otimes_{t=1}^T \ell(\cdot|\theta_1)$. Finally, the operator $\mathbb{E}$ denotes the expectation with respect to $\mathbb{P}$.

B. Network Model

The interaction between agents is captured by a directed graph $G = ([n], E)$, where $[n]$ is the set of nodes corresponding to agents, and E is the set of edges.[3] Agent i receives information from j only if the pair $(i,j) \in E$. We let $\mathcal{N}_i = \{j \in [n] : (i,j) \in E\}$ be the set of neighbors of agent i. Throughout the learning process agents truthfully report their information to their neighbors. We represent by $[W]_{ii} > 0$ the *self-reliance* of agent i, and by $[W]_{ij} > 0$ the weight that agent i assigns to information received from agent j in its neighborhood. Then, the matrix W is constructed such that $[W]_{ij}$ denotes the entry in its i-th row and j-th column. Therefore, W has nonnegative entries, and $[W]_{ij} > 0$ only if $(i,j) \in E$. For normalization purposes, we further assume that W is stochastic; hence

$$\sum_{j=1}^n [W]_{ij} = \sum_{j \in \mathcal{N}_i} [W]_{ij} = 1.$$

A3. We assume that the network is *strongly connected*, i.e., there exists a directed path from any agent $i \in [n]$ to any agent $j \in [n]$. We further assume that W is diagonalizable with real eigenvalues.[4]

The strong connectivity constraint in assumption **A3** guarantees the information flow in the network. The assumption implies that $\lambda_1(W) = 1$ is unique, and the other eigenvalues of W are strictly less than one in magnitude [26]. Given the matrix of social interactions W, the eigenvector centrality is a nonnegative vector π such that for all $i \in [n]$

$$\pi(i) = \sum_{j=1}^n [W]_{ji} \pi(j) \tag{1}$$

for $\|\pi\|_1 = 1$. Then, $\pi(i)$ denoting the i-th element of π is the eigenvector centrality of agent i. In the matrix form, the preceding relation takes the form $\pi^\top W = \pi^\top$, which means π is the stationary distribution of W. Assumption **A3** entails that the Markov chain W is irreducible and aperiodic, and the unique stationary distribution π has strictly positive components [26].

C. Centralized Detection

To motivate the development of distributed scheme, we commence by introducing centralized detection.[5] In this case, the scenario could be described as a two player repeated game between Nature and a *centralized* agent (expert) that has *global* information to learn the true state. More specifically, the expert observes the sequence of signals $\{s_t\}_{t=1}^T$ that are in turn revealed by Nature, and knows the entire network characteristics. At any round $t \in [T]$, the expert accumulates a *weighted average* of log-marginals, and forms the *belief* $\mu_t \in \Delta_m$ about the states, where $\Delta_m = \{\mu \in \mathbb{R}^m | \mu \succeq 0,\ \sum_{k=1}^m \mu(k) = 1\}$ denotes the m-dimensional probability simplex. Letting

$$\psi_t := \sum_{i=1}^n \pi(i) \psi_{i,t} = \sum_{i=1}^n \pi(i) \log \ell_i(s_{i,t}|\theta) \tag{2}$$

[2] The likelihoods can differ on a set of measure zero.

[3] Without loss of generality, we assume that $n > 1$ to have a well-defined network.

[4] Note that diagonalizability is not necessary for convergence analysis, and it only simplifies the results by avoiding Jordan blocks. In the absence of this assumption, our theoretical results will depend on the size of the largest Jordan block of W, which only complicates the message of the paper.

[5] The method can be cast as a special case of *Follow the Regularized Leader* [27] and *Mirror Descent* [28] algorithm.

the sequence of interactions could be depicted in the form of the following algorithm:

Centralized Detection

Input: A uniform prior belief μ_0, a learning rate $\eta > 0$.
Initialize: Let $\phi_0(k) = 0$ for all $k \in [m]$.
At time $t = 1, \ldots, T$:
Observe the signal $s_t = (s_{1,t}, s_{2,t}, \ldots, s_{n,t})$, update the vector function ϕ_t, and form the belief μ_t as follows:

$$\phi_t = \phi_{t-1} + \psi_t$$
$$\mu_t = \arg\min_{\mu \in \Delta_m} \left\{ -\mu^\top \phi_t + \frac{1}{\eta} D_{KL}(\mu \| \mu_0) \right\}. \tag{3}$$

Weighting the marginals based on the eigenvector centrality (2), the centralized detector aggregates a geometric average of marginals in ϕ_t. At each time $t \in [T]$, the goal is to maximize the expected sum while sticking to the default belief μ_0, i.e., minimizing the divergence. The trade-off between the two behavior is tuned with the *learning rate* η.

Let us note that according to Jensen's inequality for the concave function $\log(\cdot)$, we have for every $i \in [n]$ and $k \in [m]$ that

$$-D_{KL}\left(\ell_i(\cdot|\theta_1) \| \ell_i(\cdot|\theta_k)\right) = \mathbb{E}\left[\log \frac{\ell_i(\cdot|\theta_k)}{\ell_i(\cdot|\theta_1)}\right] \leq \log \mathbb{E}\left[\frac{\ell_i(\cdot|\theta_k)}{\ell_i(\cdot|\theta_1)}\right] = 0$$

where the inequality turns to equality if and only if $\ell_i(\cdot|\theta_1) = \ell_i(\cdot|\theta_k)$, i.e., iff $\theta_k \in \bar{\Theta}_i$. Therefore, it holds that $\mathbb{E}[\log \ell_i(\cdot|\theta_k)] \leq \mathbb{E}[\log \ell_i(\cdot|\theta_1)]$, and recalling that the stationary distribution π consists of positive elements, we have for any $k \neq 1$ that

$$\mathbb{E}\left[\sum_{i=1}^{n} \pi(i)\Psi_i(k)\right] = \mathbb{E}\left[\sum_{i=1}^{n} \pi(i) \log \ell_i(\cdot|\theta_k)\right] < \mathbb{E}\left[\sum_{i=1}^{n} \pi(i) \log \ell_i(\cdot|\theta_1)\right] = \mathbb{E}\left[\sum_{i=1}^{n} \pi(i)\Psi_i(1)\right]$$

where the strict inequality is due to uniqueness of the true state θ_1, and the fact that $\bar{\Theta} = \cap_{i=1}^{n} \bar{\Theta}_i = \{\theta_1\}$ based on assumption **A2**. In the sequel, without loss of generality, we assume the following descending order, i.e.,

$$\mathbb{E}\left[\sum_{i=1}^{n} \pi(i)\Psi_i(1)\right] > \mathbb{E}\left[\sum_{i=1}^{n} \pi(i)\Psi_i(2)\right] \geq \cdots \geq \mathbb{E}\left[\sum_{i=1}^{n} \pi(i)\Psi_i(m)\right]. \tag{4}$$

We shall see that the ordering will only simplify the derivation of technical results throughout the paper.

D. Distributed Detection

We now extend the previous section to distributed setting modeled based on a network of agents. In the distributed scheme, each agent $i \in [n]$ only observes the stream of private signals $\{s_{i,t}\}_{t=1}^{T}$ generated based on the parametrized likelihood $\ell_i(\cdot|\theta_1)$. That is, agent $i \in [n]$ does not directly observe $s_{j,t}$ for any $j \neq i$. As a result, it gathers the local information by averaging the log-likelihoods in its neighborhood, and forms the belief $\mu_{i,t} \in \Delta_m$ at round $t \in [T]$ as follows:

Distributed Detection

Input: A uniform prior belief $\mu_{i,0}$, a learning rate $\eta > 0$.
Initialize: Let $\phi_{i,0}(k) = 0$ for all $k \in [m]$ and $i \in [n]$.
At time $t \in [T]$:
Observe the signal $s_{i,t}$, update the function $\phi_{i,t}$, and form the belief $\mu_{i,t}$ as follows:

$$\phi_{i,t} = \sum_{j \in \mathcal{N}_i} [W]_{ij} \phi_{j,t-1} + \psi_{i,t}$$
$$\mu_{i,t} = \arg\min_{\mu \in \Delta_m} \left\{ -\mu^\top \phi_{i,t} + \frac{1}{\eta} D_{KL}(\mu \| \mu_{i,0}) \right\}. \tag{5}$$

As outlined above, each agent updates its belief using purely local diffusion. We are interested in measuring the efficiency of the distributed algorithm via a metric comparing that to its centralized counterpart. The centralized detector (expert) collects *all* log-marginals and weights them according to centralities. A distributed detector (an agent) collects *local* log-marginals, and does not have access to centralities. At any round $t \in [T]$, let us postulate that the cost which agent $i \in [n]$ needs to pay to have the same opinion as the expert is $D_{KL}(\mu_{i,t} \| \mu_t)$; then, the total *decentralization cost* that the agent incurs after T rounds is as follows:

$$\mathbf{Cost}_{i,T} := \sum_{t=1}^{T} D_{KL}(\mu_{i,t} \| \mu_t) = \sum_{t=1}^{T} \mathbb{E}_{\mu_{i,t}}\left[\log \frac{\mu_{i,t}}{\mu_t}\right]. \tag{6}$$

At each round, the outputs of the centralized and decentralized algorithm are probability distributions over state space. The KL-divergence captures the dissimilarity of two probability distributions; hence, it could be a reasonable metric to measure the difference between two algorithms. The function quantifies the difference between the agent that observes private signals $\{s_{i,t}\}_{t=1}^{T}$ and an expert that has $\{s_t\}_{t=1}^{T}$ and π available. In other words, it shows how well the decentralized algorithm copes with the partial information. Note importantly that $\mathbf{Cost}_{i,T}$ is a random quantity since the expectation is *not* taken with respect to randomness of signals.

We conclude this section with the following lemma which reiterates that both algorithms are reminiscent of the well-known *Exponential Weights* algorithm.

Lemma 1: The update rules (3) and (5) have the explicit form solutions

$$\mu_t(k) = \frac{\exp\{\eta\phi_t(k)\}}{\langle \mathbb{1}, \exp\{\eta\phi_t\}\rangle} \text{ and } \mu_{i,t}(k) = \frac{\exp\{\eta\phi_{i,t}(k)\}}{\langle \mathbb{1}, \exp\{\eta\phi_{i,t}\}\rangle}$$

respectively, for any $i \in [n]$ and $k \in [m]$. Moreover

$$\phi_{i,t} = \sum_{\tau=1}^{t}\sum_{j=1}^{n}[W^{t-\tau}]_{ij}\psi_{j,\tau}.$$

One can observe from above that

$$\begin{aligned}\sum_{i=1}^{n}\pi(i)\phi_{i,t} &= \sum_{\tau=1}^{t}\sum_{j=1}^{n}\sum_{i=1}^{n}\pi(i)[W^{t-\tau}]_{ij}\psi_{j,\tau}\\ &= \sum_{\tau=1}^{t}\sum_{j=1}^{n}\pi(j)\psi_{j,\tau} = \phi_t\end{aligned}$$

which connects the centralized and decentralized update via eigenvector centrality (1). As explored in [21] and [24], we shall see that centrality plays an important role in the convergence rate.

III. Finite-time Analysis of Beliefs and Cost Functions

In this section, we investigate the convergence of agents' beliefs to the true state in the network. Agents exchange information over time, and reach consensus about the true state. The connectivity of the network plays an important role in the learning as $W^t \to \mathbb{1}\pi^\top$ as $t \to \infty$. To examine the learning rate, we need to have knowledge about the mixture behavior of Markov chain W. The following lemma sheds light on the mixture rate, and we invoke it later for technical analysis.

Lemma 2: Let the strong connectivity of network (Assumption **A3**) hold, and define $\lambda_{\max}(W) := \max\{|\lambda_n(W)|, |\lambda_2(W)|\}$. Then, for any $t \in [T]$, the stochastic matrix W satisfies

$$\sum_{\tau=1}^{t}\sum_{j=1}^{n}\left|[W^{t-\tau}]_{ij} - \pi(j)\right| \leq \frac{4\log n}{1-\lambda_{\max}(W)}$$

for any $i \in [n]$ where $0 \leq \lambda_{\max}(W) < 1$.

We now establish that agents have arbitrarily close opinions in a strongly connected network. Furthermore, the convergence rate is governed by cardinality of state space and network characteristics.

Lemma 3: Let the sequence of beliefs $\{\mu_{i,t}\}_{t=1}^{T}$ for each agent $i \in [n]$ be generated by the distributed detection algorithm with the learning rate η. Given bounded log-marginals (Assumption **A1**), global identifiability of the true state (Assumption **A2**), and strong connectivity of the network (Assumption **A3**), for each individual agent $i \in [n]$ it holds that

$$\frac{1}{\eta}\log\|\mu_{i,t} - \mathbf{e}_1\|_{TV} \leq -\mathcal{I}(\theta_1,\theta_2)t + \sqrt{2B^2 t\log\frac{m}{\delta}} + \frac{8B\log n}{1-\lambda_{\max}(W)} + \frac{\log m}{\eta}$$

with probability at least $1-\delta$, where for $k \geq 2$

$$\mathcal{I}(\theta_1,\theta_k) := \sum_{i=1}^{n}\pi(i)D_{KL}\left(\ell_i(\cdot|\theta_1)\|\ell_i(\cdot|\theta_k)\right).$$

In particular, we have $\|\mu_{i,t} - \mathbf{e}_1\|_{TV} \longrightarrow 0$ almost surely.

Beside providing an any-time bound in the high probability sense, the lemma verifies that the belief $\mu_{i,t}$ of each agent $i \in [n]$ is *strongly* consistent, i.e., it converges almost surely to a delta distribution on the true state. We also remark that the asymptotic rate of $\mathcal{I}(\theta_1,\theta_2)$ was also discovered in [19], [21], and [24] for the updates under study. However, Lemma 3 provides a non-asymptotic version of the convergence rate. Let us proceed to the next lemma to derive a total variation bound on the decentralization cost (6).

Lemma 4: The instantaneous KL cost associated to the distributed detection algorithm with the learning rate η satisfies for any $t \in [T]$

$$D_{KL}(\mu_{i,t}\|\mu_t) \leq 2\|\mathbf{e}_1 - \mu_t\|_{TV}$$

as long as $\eta\|q_{i,t}\|_\infty \leq 1/4$ at each round, where $q_{i,t} := \phi_{i,t} - \phi_t$.

The bound in Lemma 4 is evocative of a reverse Pinsker's inequality. It provides a total variation bound on the cost function which is of the KL-divergence form. Let us remark that an appropriate choice of learning rate η warrants the condition $\eta\|q_{i,t}\|_\infty \leq 1/4$. We now present the main result of the paper in the following theorem.

Theorem 5: Let the sequence of beliefs $\{\mu_{i,t}\}_{t=1}^{T}$ for each agent $i \in [n]$ be generated by the distributed detection algorithm with the choice of learning rate $\eta = 1 - \lambda_{\max}(W)/16B\log n$. Given bounded log-marginals (Assumption **A1**), global identifiability of the true state (Assumption **A2**), and strong connectivity of the network (Assumption **A3**), we have

$$\mathbf{Cost}_{i,T} \leq \frac{18B^2}{\mathcal{I}^2(\theta_1,\theta_2)}\max\left\{\log\frac{6m}{\delta}, \frac{3B\sqrt{2}}{\mathcal{I}(\theta_1,\theta_2)}\right\} + \frac{48B\log n}{\mathcal{I}(\theta_1,\theta_2)}\frac{\log m + 2}{1-\lambda_{\max}(W)}$$

with probability at least $1-\delta$.

Regarding Theorem 5 the following comments are in order: the rate is related to the inverse of $\mathcal{I}(\theta_1,\theta_2)$ which is a weighted average of KL-divergence of observations under θ_2 (the second best alternative) from observations under θ_1 (the true state). Also, from the definition of $\mathcal{I}(\theta_1,\theta_2)$ in Lemma 3, the weights turn out to be agents' centralities. Intuitively, when signals hardly reveal the difference between the best two candidates for the true state, agents must make more effort to distinguish the two. In turn, this results in suffering a larger cost caused by slower learning. The decentralization cost always scales logarithmically with the number of states m. Now define

$$\gamma(W) := 1 - \lambda_{\max}(W) \tag{7}$$

as the *spectral gap* of the network. Then, Theorem 5 suggests that for large networks, the cost scales inversely in the spectral gap, and logarithmically with the network size n. Finally, the detection cost is time-independent (with high probability), proving the best possible bound with respect to time. Therefore, the average expected cost (per iteration cost) asymptotically tends to zero.

IV. The Impact of Network Topology

The results of previous section verify that network characteristics govern the learning process. We now discuss the role of agents' centralities and the network spectral gap.

A. Effect of Agent Centrality

To examine centrality, let us return to the definition of $\mathcal{I}(\theta_1, \theta_2)$ in Lemma 3, and imagine that the network is *collaborative* in the sense that the network designer wants to expedite learning. Then, to have the best information dispersion, the marginal which collects the most evidence in favor of θ_1 against θ_2 should be allocated to the most central agent. By the same token, in an *adversarial* network where Nature aims to delay the learning process, such marginal should be assigned to the least central agent. To sum up, let us put forth the concept of network *regularity* as defined in [24] in the context of social learning. Recalling the definition of eigenvector centrality (1), we say a network G is more regular than G' if π' majorizes π, i.e., if for all $j \in [n]$

$$\sum_{i=1}^{j} \pi_{[i]} \le \sum_{i=1}^{j} \pi'_{[i]} \tag{8}$$

where $\pi_{[i]}$ denotes the i-th largest element of π. Letting

$$u := [D_{KL}(\ell_1(\cdot|\theta_1)\|\ell_1(\cdot|\theta_2)), \ldots, D_{KL}(\ell_n(\cdot|\theta_1)\|\ell_n(\cdot|\theta_2))]^\top$$

it is a straightforward consequence of Lemma 1 proved in [24] that

$$\sum_{i=1}^{n} \pi_{[i]} u_{[i]} \le \sum_{i=1}^{n} \pi'_{[i]} u_{[i]}$$

when π' majorizes π. Therefore, spreading more informative signals among central agents speeds up the learning procedure.

B. Optimizing the Spectral Gap

We now turn our attention to the spectral gap of network (7). Suppose that agents are given a default communication matrix W which determines their neighborhood and centrality. The problem is to find the optimal spectral gap assuming that the neighborhood and centrality of each agent are fixed. The key idea is to change the mixing behavior of the Markov chain W. It is well-known, for instance, that we could do so using lazy ran dom walks [29] which replaces W with $(1/2)(W+I_n)$. To generalize the idea, let us define a modified communication matrix

$$W' := \alpha W + (1-\alpha) I_n \quad \alpha \in [0,1] \tag{9}$$

which has the same eigenstructure as W. Then, the eigenvalues of W' are weighted averages of those of W with one. From standpoint of network design, one can exploit the freedom in choosing α to optimize the spectral gap.

Proposition 6: The optimal spectral gap of the modified communication matrix W' (9) is as follows:

$$\gamma^* = \frac{2 - 2\lambda_2(W)}{2 - \lambda_n(W) - \lambda_2(W)} \text{ for } \alpha^* = \frac{2}{2 - \lambda_n(W) - \lambda_2(W)}$$

when $\lambda_n(W) + \lambda_2(W) < 0$

Proof: To optimize the spectral gap, we need to minimize the second largest eigenvalue of W' in magnitude, that is, to solve the min-max problem

$$\min_{\alpha \in [0,1]} \max \{ |\alpha\lambda_2(W) + 1 - \alpha|, |\alpha\lambda_n(W) + 1 - \alpha| \}. \tag{10}$$

The functions $|\alpha\lambda_2(W) + 1 - \alpha|$ and $|\alpha\lambda_n(W) + 1 - \alpha|$ are both convex with respect to α. Therefore, the point-wise maximum of the two is also convex, and achieves its minimum on a compact set. Since $\lambda_n(W) < -\lambda_2(W)$ by hypothesis, the minimum occurs at the intersection of the following lines:

$$\alpha\lambda_2(W) + 1 - \alpha = -\alpha\lambda_n(W) + \alpha - 1$$

yielding $\alpha^* = 2/(2 - \lambda_n(W) - \lambda_2(W))$. Plugging α^* into the min-max problem (10), we calculate the optimal value $\lambda^*_{\max}$ as

$$\lambda^*_{\max} = \frac{\lambda_2(W) - \lambda_n(W)}{2 - \lambda_n(W) - \lambda_2(W)}$$

and since $\gamma^* = 1 - \lambda^*_{\max}$ the proof follows immediately. ■

We remark that when the Markov chain is symmetric, the problem can be formulated as a convex optimization [30]. Moreover, for gossip protocols where the expected communication matrix is symmetric, the problem can be posed as a semidefinite program [31]. However, in our setting the chain is not necessarily symmetric and these results are not applicable.

C. Sensitivity to Link Failure

It is intuitive that in a network with more links, agents are offered more opportunities for communication. Adding links provides more avenues for spreading information, and improves the learning quality. We study this phenomenon for symmetric networks where a pair of agents assign similar weights to each other, i.e., $W^\top = W$. In particular, we explore the connection of spectral gap with the link failure. In this regard, let us introduce the following positive semi-definite matrix:

$$\Delta W(i,j) := (\mathbb{e}_i - \mathbb{e}_j)(\mathbb{e}_i - \mathbb{e}_j)^\top \tag{11}$$

where $\mathbb{e}_i$ is the i-th unit vector in the standard basis of $\mathbb{R}^n$. Then, for $i, j \in [n]$ the matrix

$$\bar{W}(i,j) := W + [W]_{ij} \Delta W(i,j) \tag{12}$$

corresponds to a new communication matrix that removes edges (i,j) and (j,i) from the network, and adds $[W]_{ij} = [W]_{ji}$ to the self-reliance of agent i and agent j.

Proposition 7: Consider the communication matrix $\bar{W}(i,j)$ in (12). Then, for any $i, j \in [n]$ the following inequality holds:

$$\lambda_{\max}(W) \le \lambda_{\max}\left(\bar{W}(i,j)\right)$$

as long as W is positive semi-definite.

Proof: We recall that $\Delta W(i,j)$ in (11) is positive semi-definite with $\lambda_n(\Delta W(i,j)) = 0$. Applying Weyl's eigenvalue inequality on (12), we obtain for any $k \in [n]$

$$\lambda_k(W) \le \lambda_k\left(\bar{W}(i,j)\right)$$

which holds in particular for $k = 2$. On the other hand, the matrix W is positive semi-definite, so we have that $\lambda_{\max}(W) = \lambda_2(W)$. Combining with the fact that $\bar{W}(i,j)$ is symmetric and positive semi-definite, the proof is completed. ■

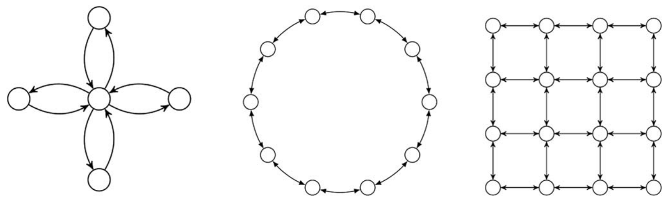

Fig. 1. Illustration of networks: star, cycle and grid networks with n agents. For each network, each individual agent possesses a self-reliance of $\omega \in (0, 1)$.

The proposition immediately implies that removing a link reduces the spectral gap. In this case, in view of the bound in Theorem 5, the decentralization cost has more latitude to vary. Therefore, to keep the costs small, agents tend to maintain their connections. Let us take note of the delicate point that monotone increase in the upper bound does *not* necessarily imply a monotone increase in the cost; however, one can *roughly* expect such behavior. We elaborate on this issue in the numerical experiments. Notice that the positive semi-definiteness constraint on W is not strong, since it can be easily satisfied by replacing a lazy random walk $(1/2)(W + I_n)$ with W. Finally, we remark that link failures in distributed optimization [32] and consensus protocols [33] has been previously studied in the literature. We refer the interested reader to these references where the impact of *random* link failure is considered.

D. Star, Cycle and Grid Networks

We now examine the spectral gap impact for some interesting networks (Fig. 1), and derive explicit bounds for decentralization cost. In the star network (regardless of the network size), existence of a central agent always preserves the network diameter, and therefore, we expect a benign scaling with network size. On the other side of the spectrum lies the cycle network where the diameter grows linearly with the network size. We should, hence, observe how the poor communication in cycle network affects the learning rate. Finally, as a possible model for sensor networks, we study the grid network where the network size scales quadratically with the diameter.

Corollary 8: Under conditions of Theorem 5 and the choice of learning rate $\eta = \gamma(\cdot)/16B \log n$, for n large enough we have the following bounds on the decentralization cost:

(**a**) For the star network in Fig. 1

$$\mathbf{Cost}_{i,T} \leq \mathcal{O}\left(\frac{\log[nm]}{\min\{1-\omega, 1-|2\omega-1|\}}\right).$$

(**b**) For the cycle network in Fig. 1

$$\mathbf{Cost}_{i,T} \leq \mathcal{O}\left(\frac{\log[nm]}{\min\left\{1-|2\omega-1|, 2(1-\omega)\sin^2\frac{\pi}{n}\right\}}\right).$$

(**c**) For the grid network in Fig. 1

$$\mathbf{Cost}_{i,T} \leq \mathcal{O}\left(\frac{\log[nm]}{\min\left\{1-|2\omega-1|, 2(1-\omega)\sin^2\frac{\pi}{\sqrt{n}}\right\}}\right).$$

Proof: The spectrum of the Laplacian of star and cycle graphs are well-known [34]. We have the eigenvalue set corresponding to communication matrix of star and cycle graphs as

$$\{1, \omega, \ldots, \omega, 2\omega-1\} \text{ and } \left\{\omega + (1-\omega)\cos\frac{2\pi i}{n}\right\}_{i=0}^{n-1}$$

respectively. Therefore, the proof of (**a**) and (**b**) follows immediately. The grid graph is the Cartesian product of two rings of size $\sqrt{n}$ (due to wraparounds at the edges), and hence, its eigenvalues are derived by summing the eigenvalues of two $\sqrt{n}$-rings [34]. Therefore, the eigenvalue set takes the form

$$\left\{\omega + (1-\omega)\cos\frac{\pi(i+j)}{\sqrt{n}}\cos\frac{\pi(i-j)}{\sqrt{n}}\right\}_{i,j=0}^{\sqrt{n}-1}$$

and the proof of (**c**) is completed. ■

Let us use the notation $\tilde{\mathcal{O}}(\cdot)$ to hide the poly log factors. Then, the bounds derived in Corollary 8 indicate that the algorithm requires $\tilde{\mathcal{O}}(1)$ iterations to achieve a near optimal log-distance from the true state in the star network. However, the rate deteriorates to $\tilde{\mathcal{O}}(n^2)$ (respectively, $\tilde{\mathcal{O}}(n)$) in the cycle (respectively, grid) network. In all cases, the rate is proportional to the diameter squared, and diameter is a natural indicator of information dissemination quality.

V. Numerical Experiment: Binary Signal Detection

In this section, we discuss our numerical experiments. Note that, as mentioned in the footnote of assumption **A3**, in our convergence results the communication matrix need not be diagonalizable, and the assumption is only for convenience. In what follows, we disregard diagonalizability (in the construction of network) for the first section. Therefore, we verify the generality of convergence for arbitrary strongly connected networks.

A. Convergence of Beliefs

We generate a random network of $n = 50$ agents based on the Erdös-Rényi model. In our example, each link exists with probability 0.3 independent of other links. We verify the strong connectivity of the network before running the experiment. Though generated randomly, the network is fixed throughout the process. Assume that there exist $m = 51$ states in the world and agents are to discover the true state θ_1. At time $t \in [T]$, a signal $s_{i,t} \in \{0, 1\}$ is generated based on the true state such that $\ell_i(\cdot|\theta_1) = \ell_i(\cdot|\theta_{i+1})$. In other words, for agent $i \in [n]$, we have $\bar{\Theta}_i = \{\theta_1, \theta_{i+1}\}$ and θ_{i+1} is observationally equivalent to the true state. Therefore, each agent $i \in [n]$ fails to distinguish θ_1 from θ_{i+1} once relying on the private signals. However, since we have $\bar{\Theta} = \cap_{i=1}^{n}\bar{\Theta}_i = \{\theta_1\}$, the true state is globally identifiable. Consequently, in view of Lemma 3, all agents reach a consensus on the true state (Fig. 2), and learn the truth exponentially fast.

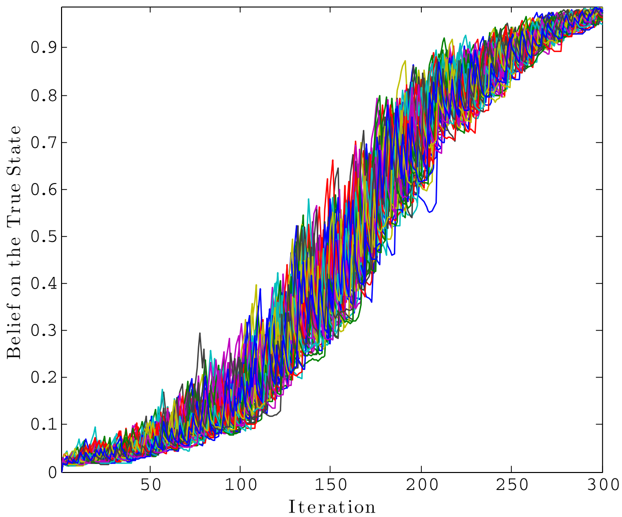

Fig. 2. The belief evolution for all 50 agents in the network. The global identifiability of the true state and strong connectivity of the network result in learning.

B. Optimizing the Spectral Gap

To verify the result of Proposition 6, we must construct a communication matrix that is diagonalizable, yet not symmetric. We let

$$W_1 = \begin{bmatrix} 0 & 0.95 & 0.05 \\ 0.95 & 0 & 0.05 \\ 0.05 & 0.95 & 0 \end{bmatrix} \text{ and } W_2 = \begin{bmatrix} 0.5 & 0.5 \\ 0.3 & 0.7 \end{bmatrix}$$

and set $W = W_1 \otimes (W_2 \otimes W_2)$. One can verify that W is row stochastic, diagonalizable and asymmetric. Also, $W^t \to \mathbb{1}\pi^\top$ as $t \to \infty$, where π consists of positive elements. The resulting network has a specific structure, but it suits our purposes since it satisfies all the conditions without being symmetric. The signal generating process is precisely the same as the previous section. We now turn to optimizing the spectral gap to speed up learning. We proved in Proposition 6 that every default communication matrix can be adjusted to a matrix W' which has the optimal spectral gap when centralities are fixed. Setting the parameter α in (9) equal to α^* derived in Proposition 6, we obtain the optimal network. In this example we have $\gamma(W) = 0.05$, $\alpha^* = 0.7273$ and $\gamma^* = 0.5818$. The dependence of decentralization cost to the spectral gap was theoretically proved in Theorem 5. Applying the results of Proposition 6 verifies that in the optimal network, agents suffer a lower decentralization cost comparing to the default network (Fig. 3). On the other hand, we proved theoretically in Theorem 5 that the cost bound is time-independent with high probability. Interestingly, the plot verifies the high probability upper bound on the cost for both cases. In particular, we observe that, for each agent, the asymptotic value of the optimal network is smaller than the default network, while both curves remain constant after almost 300 iterations. The difference between the asymptotic values, in fact, reflects the quality of finite-time performance. The reason is that for larger iterations (say, $t > 300$) both centralized and decentralized algorithms have almost converged to the true state, and there is no cost accumulation, whereas in smaller iterations (say, $t < 300$) the network structure is still impactful in collecting the cost.

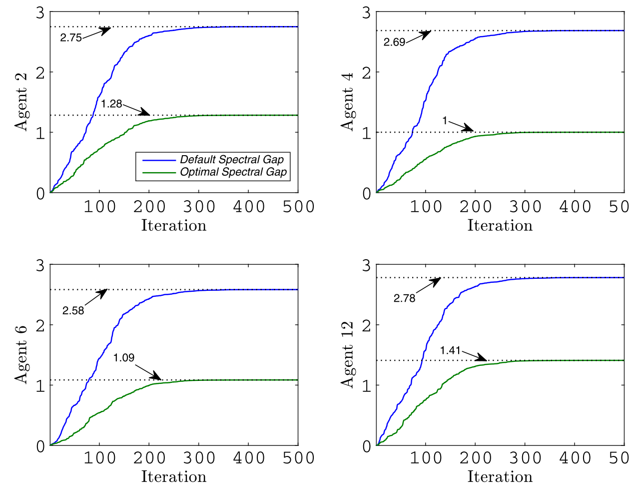

Fig. 3. The plot of decentralization cost versus time horizon for agents 2, 4, 6, and 12 in the network. The cost in the network with the optimal spectral gap (green) is always less than the network with default weights (blue).

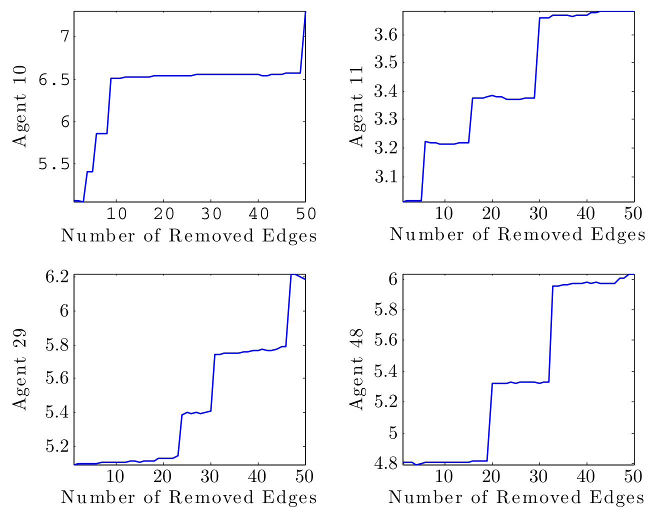

Fig. 4. The decentralization cost at round $T = 300$ for agents 10, 11, 29, and 48 in the network. Removing the links causes poor communication among agents and increase the decentralization cost.

C. Sensitivity to Link Failure

To evaluate the result of Proposition7, we need a symmetric network. The upper triangle of W is generated using Erdös-Rényi model (similar to the first section), and the matrix is then symmetrized. In this case every agent is equally central, and we have $\pi = \mathbb{1}/n$. To study the impact of link failure, we sequentially select random pairs of agents in the network, and remove their connection. Each time that a link is discarded, we compute the decentralization cost in the new network at iteration $T = 300$, and continue the process until 50 bi-directional edges are eliminated from the network. In view of Proposition7, we expect a monotone decrease in the spectral gap which amounts to a larger decentralization cost. We plot the cost for four agents in the network, and observe that the behavior is almost (not quite) monotonic (Fig. 4). The monotone dependence of the upper bound to the spectral gap (Theorem 5) does not necessarily guarantee a monotone relationship between cost and the spectral gap. However, we can intuitively expect that removing edges makes the network less connected, causing the performance of distributed algorithm to deteriorate.

VI. Conclusion

We considered a distributed detection model where a network of agents aim to learn the underlying state of the world. The private signals do not provide enough information for agents about the true state. Hence, agents engage in a local communication to compensate for their imperfect knowledge. Each agent iteratively forms a belief about the state space using the collected data in its neighborhood. We analyzed the learning procedure for a *finite* time horizon. To study the efficiency of our algorithm versus its centralized counterpart, we brought forward the idea of KL cost. It turned out that network size, spectral gap, centrality of each agent and relative entropy of agents' signal structures are the key parameters that affect distributed detection. We established that allocating more informative signals to central agents as well as optimizing the spectral gap can speed up learning. We also proved that the learning rate deteriorates in the case of link failures, which can be seen as a side effect of poor communication. Finally, we would like to address a few issues in future works. In this paper, we discussed a communication model in which agents exchange information at every round. In some networks, all-time communication is potentially costly or unnecessary. It would be interesting to study the trade-off between communication and learning in finite time. As another direction, we can consider scenarios where the signal distributions are not stationary. This generalizes the model to dynamic parameters where we can investigate detection robustness in changing environments.

Appendix
Omitted Proofs

Proof of Lemma 1: The proof is elementary, and it is only given to keep the paper self-contained. We write the Lagrangian associated to the update (3) as

$$L(\mu,\lambda) = -\mu^\top\phi_t + \frac{1}{\eta}\left\langle \mu, \log\frac{\mu}{\mu_0}\right\rangle + \lambda\mu^\top\mathbb{1} - \lambda$$

where we left the positivity constraint implicit. Differentiating above with respect to μ and λ, and setting the derivatives equal to zero, we get

$$\mu_t(k) = \mu_0(k)\exp\left\{\eta\left(\phi_t(k)-\lambda\right)-1\right\} \text{ and } \mu_t^\top\mathbb{1} = 1$$

respectively, for any $k \in [m]$. Combining the equations above and noting that μ_0 is uniform, we have

$$\frac{1}{m}\exp\{-\eta\lambda - 1\}\sum_{k=1}^{m}\exp\left\{\eta\phi_t(k)\right\} = 1$$

which allows us to solve for λ and calculate the optimal solution μ_t as follows:

$$\mu_t(k) = \frac{\exp\left\{\eta\phi_t(k)\right\}}{\sum_{k=1}^{m}\exp\left\{\eta\phi_t(k)\right\}}.$$

The proof for $\mu_{i,t}$ follows precisely in the same fashion. To calculate $\phi_{i,t}$, notice that in view of the first update in (5) we have

$$\begin{bmatrix}\phi_{1,t}\\ \phi_{2,t}\\ \vdots\\ \phi_{n,t}\end{bmatrix} = (W\otimes I_m)\begin{bmatrix}\phi_{1,t-1}\\ \phi_{2,t-1}\\ \vdots\\ \phi_{n,t-1}\end{bmatrix} + \begin{bmatrix}\psi_{1,t}\\ \psi_{2,t}\\ \vdots\\ \psi_{n,t}\end{bmatrix}$$

where $\otimes$ denotes the Kronecker product. The equation above represents a discrete-time linear system. Given the fact that $\phi_{i,0}(k) = 0$ for all $k \in [m]$ and $i \in [n]$, the closed-form solution of the system takes the form

$$\begin{bmatrix}\phi_{1,t}\\ \phi_{2,t}\\ \vdots\\ \phi_{n,t}\end{bmatrix} = \sum_{\tau=1}^{t}(W\otimes I_n)^{t-\tau}\begin{bmatrix}\psi_{1,\tau}\\ \psi_{2,\tau}\\ \vdots\\ \psi_{n,\tau}\end{bmatrix}$$
$$= \sum_{\tau=1}^{t}(W^{t-\tau}\otimes I_n)\begin{bmatrix}\psi_{1,\tau}\\ \psi_{2,\tau}\\ \vdots\\ \psi_{n,\tau}\end{bmatrix}.$$

Therefore, extracting $\phi_{i,t}$ for each $i \in [n]$ from the preceding relation completes the proof. ■

Proof of Lemma 2: Since the network is strongly connected and the corresponding W is irreducible and aperiodic, by standard properties of stochastic matrices (see, e.g., [26]), the diagonalizable matrix W satisfies

$$\left\|\mathbf{e}_i^\top W^t - \pi^\top\right\|_1 \le n\lambda_{\max}(W)^t \tag{13}$$

for any $i \in [n]$, where π is the stationary distribution of a Markov chain with transition kernel W. Let us observe the following inequality:

$$n\lambda_{\max}(W)^{t-\tau} \le 2 \text{ for } t-\tau \ge \tilde{t} := \frac{\log\frac{n}{2}}{\log\lambda_{\max}(W)^{-1}}$$

and recall that the inequality $\|\mathbf{e}_i^\top W^{t-\tau} - \pi^\top\|_1 \le 2$ always holds since any power of W is stochastic. With that in mind, we use (13) to break the following sum into two parts to get:

$$\sum_{\tau=1}^{t}\left\|\mathbf{e}_i^\top W^{t-\tau} - \pi^\top\right\|_1 = \sum_{\tau=1}^{t-\tilde{t}}\left\|\mathbf{e}_i^\top W^{t-\tau} - \pi^\top\right\|_1 + \sum_{\tau=t-\tilde{t}+1}^{t}\left\|\mathbf{e}_i^\top W^{t-\tau} - \pi^\top\right\|_1$$
$$\le \sum_{\tau=1}^{t-\tilde{t}} n\lambda_{\max}(W)^{t-\tau} + 2\tilde{t}$$
$$\le \frac{n\lambda_{\max}(W)^{\tilde{t}}}{1-\lambda_{\max}(W)} + \frac{2\log\frac{n}{2}}{\log\lambda_{\max}(W)^{-1}}$$

for any $i \in [n]$. Note that $1 - \lambda_{\max}(W) \leq \log \lambda_{\max}(W)^{-1}$ and $2 + 2\log(n/2) \leq 4\log n$, since $n > 1$. It follows by plugging $\tilde{t}$ into above that:

$$\sum_{\tau=1}^{t}\sum_{j=1}^{n}\left|[W^{t-\tau}]_{ij} - \pi(j)\right| = \sum_{\tau=1}^{t}\left\|\mathbf{e}_i^\top W^{t-\tau} - \pi^\top\right\|_1 \leq \frac{4\log n}{1-\lambda_{\max}(W)}$$

which completes the proof. ∎

We use the following inequality in [35] in the proof of Lemma 3.

Lemma 9 (McDiarmid's Inequality): Let $X_1, \ldots, X_N \in \chi$ be independent random variables and consider the mapping $H : \chi^N \mapsto \mathbb{R}$. If for $i \in \{1, \ldots, N\}$, and every sample $x_1, \ldots, x_N, x_i' \in \chi$, the function H satisfies

$$|H(x_1,\ldots,x_i,\ldots,x_N) - H(x_1,\ldots,x_i',\ldots,x_N)| \leq c_i$$

then for all $\varepsilon > 0$

$$\mathbb{P}\left\{H(x_1,\ldots,x_N) - \mathbb{E}\left[H(X_1,\ldots,X_N)\right] \geq \varepsilon\right\} \leq \exp\left\{\frac{-2\varepsilon^2}{\sum_{i=1}^N c_i^2}\right\}.$$

Proof of Lemma 3: According to Lemma 1, we have

$$\begin{aligned}\mu_{i,t}(1) &= \frac{\exp\{\eta\phi_{i,t}(1)\}}{\sum_{k=1}^m \exp\{\eta\phi_{i,t}(k)\}} \\ &= \left(1 + \sum_{k=2}^m \exp\{\eta\phi_{i,t}(k) - \eta\phi_{i,t}(1)\}\right)^{-1} \\ &\geq 1 - \sum_{k=2}^m \exp\{\eta\phi_{i,t}(k) - \eta\phi_{i,t}(1)\} \qquad (14)\end{aligned}$$

where we used the fact that $(1+x)^{-1} \geq 1 - x$ for any $x \geq 0$. Since we know

$$\|\mu_{i,t} - \mathbf{e}_1\|_{TV} = \frac{1}{2}\left(1 - \mu_{i,t}(1) + \sum_{k=2}^m \mu_{i,t}(k)\right) = 1 - \mu_{i,t}(1)$$

we can combine above with (14) to obtain

$$\|\mu_{i,t} - \mathbf{e}_1\|_{TV} \leq \sum_{k=2}^m \exp\{\eta\phi_{i,t}(k) - \eta\phi_{i,t}(1)\}. \qquad (15)$$

For any $k \in [m]$, define

$$\Phi_{i,t}(k) := \sum_{\tau=1}^{t}\sum_{j=1}^{n}[W^{t-\tau}]_{ij}\log \ell_j(\cdot|\theta_k)$$

and note that $\Phi_{i,t}(k)$ is a function of nt random variables. As required in McDiarmid's inequality in Lemma 9, set $H = \Phi_{i,t}(k) - \Phi_{i,t}(1)$, fix the samples for $nt-1$ random variables, and draw two different samples $s_{j,\tau}$ and $s'_{j,\tau}$ for some $j \in [n]$ and some $\tau \in [t]$. The fixed samples are simply cancelled in the subtraction, and we have

$$\begin{aligned}&\left|H(\ldots,s_{j,\tau},\ldots) - H\left(\ldots,s'_{j,\tau},\ldots\right)\right| \\ &= \left|[W^{t-\tau}]_{ij}\left(\log\frac{\ell_j(s_{j,\tau}|\theta_k)}{\ell_j(s_{j,\tau}|\theta_1)} - \log\frac{\ell_j\left(s'_{j,\tau}|\theta_k\right)}{\ell_j\left(s'_{j,\tau}|\theta_1\right)}\right)\right| \\ &\leq [W^{t-\tau}]_{ij}2B\end{aligned}$$

where we used assumption **A1**. Since any power of W is stochastic, summing over $j \in [n]$ and $\tau \in [t]$, we get

$$\sum_{\tau=1}^{t}\sum_{j=1}^{n}\left([W^{t-\tau}]_{ij}2B\right)^2 \leq 4B^2t.$$

We now apply McDiarmid's inequality in Lemma 9 to obtain

$$\mathbb{P}\left(\phi_{i,t}(k) - \phi_{i,t}(1) > \mathbb{E}\left[\Phi_{i,t}(k) - \Phi_{i,t}(1)\right] + \varepsilon\right) \leq \exp\left\{\frac{-\varepsilon^2}{2B^2t}\right\}$$

for each fixed k. Setting the probability above to δ/m and taking a union bound over all states, the following event holds:

$$\phi_{i,t}(k) - \phi_{i,t}(1) \leq \mathbb{E}\left[\Phi_{i,t}(k) - \Phi_{i,t}(1)\right] + \sqrt{2B^2t\log\frac{m}{\delta}} \qquad (16)$$

simultaneously for all $k = 2, \ldots, m$, with probability at least $1 - \delta$. On the other hand, in view of assumption **A1**, we have

$$\begin{aligned}&\mathbb{E}\left[\Phi_{i,t}(k) - \Phi_{i,t}(1)\right] \\ &= \sum_{\tau=1}^{t}\sum_{j=1}^{n}[W^{t-\tau}]_{ij}\mathbb{E}\left[\log\ell_j(\cdot|\theta_k) - \log\ell_j(\cdot|\theta_1)\right] \\ &= \sum_{\tau=1}^{t}\sum_{j=1}^{n}\left([W^{t-\tau}]_{ij} - \pi(j)\right)\mathbb{E}\left[\log\ell_j(\cdot|\theta_k) - \log\ell_j(\cdot|\theta_1)\right] \\ &\quad + \sum_{\tau=1}^{t}\sum_{j=1}^{n}\pi(j)\mathbb{E}\left[\log\ell_j(\cdot|\theta_k) - \log\ell_j(\cdot|\theta_1)\right] \\ &\leq 2B\sum_{\tau=1}^{t}\sum_{j=1}^{n}\left|[W^{t-\tau}]_{ij} - \pi(j)\right| \\ &\quad - t\sum_{j=1}^{n}\pi(j)D_{KL}\left(\ell_j(\cdot|\theta_1)\|\ell_j(\cdot|\theta_k)\right) \\ &= 2B\sum_{\tau=1}^{t}\sum_{j=1}^{n}\left|[W^{t-\tau}]_{ij} - \pi(j)\right| - \mathcal{I}(\theta_1,\theta_k)t \\ &\leq \frac{8B\log n}{1-\lambda_{\max}(W)} - \mathcal{I}(\theta_1,\theta_k)t \qquad (17)\end{aligned}$$

where we applied Lemma 2 to derive the last step. Using (4), we simplify above to get

$$\mathbb{E}\left[\Phi_{i,t}(k) - \Phi_{i,t}(1)\right] \leq \frac{8B\log n}{1-\lambda_{\max}(W)} - \mathcal{I}(\theta_1,\theta_2)t \qquad (18)$$

for any $k = 2, \ldots, m$. Plugging (18) into (16) and combining with (15), we have

$$\begin{aligned}
&\|\mu_{i,t} - \mathbf{e}_1\|_{TV} \\
&\leq \sum_{k=2}^{m} \exp\left\{-\eta\mathcal{I}(\theta_1,\theta_2)t + \eta\sqrt{2B^2 t \log\frac{m}{\delta}} + \frac{8\eta B \log n}{1-\lambda_{\max}(W)}\right\} \\
&\leq m \exp\left\{-\eta\mathcal{I}(\theta_1,\theta_2)t + \eta\sqrt{2B^2 t \log\frac{m}{\delta}} + \frac{8\eta B \log n}{1-\lambda_{\max}(W)}\right\}
\end{aligned}$$

with probability at least $1-\delta$, and thereby completing the proof of the first part. Letting $\delta = 1/t^2$ in above and applying Borel-Cantelli lemma, the almost sure convergence follows immediately. ■

Proof of Lemma 4: We recall from the statement of the lemma that $q_{i,t}(k) = \phi_{i,t}(k) - \phi_t(k)$, and calculate the ratio $\mu_{i,t}(k)/\mu_t(k)$ for any $k \in [m]$ as follows:

$$\begin{aligned}
\frac{\mu_{i,t}(k)}{\mu_t(k)} &= \exp\{\eta q_{i,t}(k)\} \frac{\mathbb{E}_{\mu_0}[\exp\{\eta\phi_t\}]}{\mathbb{E}_{\mu_0}[\exp\{\eta\phi_{i,t}\}]} \\
&= \exp\{\eta q_{i,t}(k)\} \frac{\mathbb{E}_{\mu_0}[\exp\{\eta\phi_t\}]}{\mathbb{E}_{\mu_0}[\exp\{\eta\phi_t\}\exp\{\eta q_{i,t}\}]} \\
&= \exp\{\eta q_{i,t}(k)\} \frac{1}{\mathbb{E}_{\mu_0}\left[\frac{\exp\{\eta\phi_t\}}{\mathbb{E}_{\mu_0}[\exp\{\eta\phi_t\}]}\exp\{\eta q_{i,t}\}\right]} \\
&= \exp\{\eta q_{i,t}(k)\} \frac{1}{\mathbb{E}_{\mu_0}\left[\frac{\mu_t}{\mu_0}\exp\{\eta q_{i,t}\}\right]} \\
&= \exp\{\eta q_{i,t}(k)\} \frac{1}{\mathbb{E}_{\mu_t}[\exp\{\eta q_{i,t}\}]}.
\end{aligned}$$

This entails

$$\begin{aligned}
\frac{1}{\eta}\mathbb{E}_{\mu_{i,t}}\left[\log\frac{\mu_{i,t}}{\mu_t}\right] &= \mathbb{E}_{\mu_{i,t}}[q_{i,t}] - \frac{1}{\eta}\log\mathbb{E}_{\mu_t}[\exp\{\eta q_{i,t}\}] \\
&\leq \mathbb{E}_{\mu_{i,t}}[q_{i,t}] - \mathbb{E}_{\mu_t}[q_{i,t}]
\end{aligned}$$

where we used Jensen's inequality on the convex function $-\log(\cdot)$. Setting the expectation measures in the right hand side of above to μ_t, and recalling the ratio $\mu_{i,t}/\mu_t$ from above, we conclude that

$$\begin{aligned}
\mathbb{E}_{\mu_{i,t}}\left[\log\frac{\mu_{i,t}}{\mu_t}\right] &\leq \mathbb{E}_{\mu_t}\left[\frac{\mu_{i,t}}{\mu_t}\eta q_{i,t}\right] - \mathbb{E}_{\mu_t}[\eta q_{i,t}] \\
&= \mathbb{E}_{\mu_t}\left[\left(\frac{\exp\{\eta q_{i,t}\}}{\mathbb{E}_{\mu_t}[\exp\{\eta q_{i,t}\}]} - 1\right)\eta q_{i,t}\right] \\
&= \sum_{k=1}^{m}\mu_t(k)\eta q_{i,t}(k)\left(\frac{\exp\{\eta q_{i,t}(k)\}}{\mathbb{E}_{\mu_t}[\exp\{\eta q_{i,t}\}]} - 1\right) \\
&= \sum_{k=1}^{m}\mu_t(k)\eta q_{i,t}(k)\frac{\langle\mathbf{e}_k - \mu_t, \exp\{\eta q_{i,t}\}\rangle}{\langle\mu_t, \exp\{\eta q_{i,t}\}\rangle} \\
&\leq \frac{\exp\left\{\frac{1}{4}\right\}}{4}\sum_{k=1}^{m}\mu_t(k)\left|\langle\mathbf{e}_k - \mu_t, \exp\{\eta q_{i,t}\}\rangle\right|
\end{aligned}$$

where we used the condition $\eta\|q_{i,t}\|_\infty \leq 1/4$ to obtain the last line. We now apply Hölder's inequality for primal-dual norm pairs and use $\eta\|q_{i,t}\|_\infty \leq 1/4$ again to simplify above as

$$\begin{aligned}
&\mathbb{E}_{\mu_{i,t}}\left[\log\frac{\mu_{i,t}}{\mu_t}\right] \\
&\leq \frac{\exp\left\{\frac{1}{4}\right\}}{4}\sum_{k=1}^{m}\mu_t(k)\|\mathbf{e}_k - \mu_t\|_1 \|\exp\{\eta q_{i,t}\}\|_\infty \\
&\leq \frac{\exp\left\{\frac{1}{2}\right\}}{4}\sum_{k=1}^{m}\mu_t(k)\|\mathbf{e}_k - \mu_t\|_1 \\
&\leq \frac{\exp\left\{\frac{1}{2}\right\}}{4}\left(\|\mathbf{e}_1 - \mu_t\|_1 + 2\sum_{k=2}^{m}\mu_t(k)\right) \qquad (19)
\end{aligned}$$

where the last step follows from the fact that $\|\mathbf{e}_k - \mu_t\|_1 \leq 2$ for any $k \in [m]$. Recalling

$$\begin{aligned}
\frac{1}{2}\|\mathbf{e}_1 - \mu_t\|_1 &= \frac{1}{2}\left(1 - \mu_t(1) + \sum_{k=2}^{m}\mu_t(k)\right) \\
&= \frac{1}{2}\left(\sum_{k=1}^{m}\mu_t(k) - \mu_t(1) + \sum_{k=2}^{m}\mu_t(k)\right) = \sum_{k=2}^{m}\mu_t(k)
\end{aligned}$$

as well as the fact $\|\mathbf{e}_1 - \mu_t\|_{TV} = (1/2)\|\mathbf{e}_1 - \mu_t\|_1$, we simplify (19) to get

$$\mathbb{E}_{\mu_{i,t}}\left[\log\frac{\mu_{i,t}}{\mu_t}\right] \leq \exp\left\{\frac{1}{2}\right\}\|\mathbf{e}_1 - \mu_t\|_{TV} \leq 2\|\mathbf{e}_1 - \mu_t\|_{TV} \qquad (20)$$

and thereby completing the proof. ■

Proof of Theorem 5: We recall that $q_{i,t}$ in the statement of Lemma 4 satisfies

$$\begin{aligned}
\|q_{i,t}\|_\infty &= \left\|\sum_{\tau=1}^{t}\sum_{j=1}^{n}\left([W^{t-\tau}]_{ij} - \pi(j)\right)\psi_{j,t}\right\|_\infty \\
&\leq B\sum_{\tau=1}^{t}\sum_{j=1}^{n}\left|[W^{t-\tau}]_{ij} - \pi(j)\right| \leq \frac{4B\log n}{1-\lambda_{\max}(W)}
\end{aligned}$$

due to Lemma 2 and assumption **A1**. Therefore, the choice of $\eta = (1-\lambda_{\max}(W))/16B\log n$ guarantees that $q_{i,t}$ satisfies $\eta\|q_{i,t}\|_\infty \leq 1/4$ for all $t \in [T]$. Let us follow exactly the same steps in the proof of Lemma 3, and note that the centralized update can be recovered using $W = \mathbb{1}\pi^\top$. It can be verified from (17) that for any $t \in [T]$, we only remain with

$$\mathbb{E}[\Phi_t(k) - \Phi_t(1)] \leq -\mathcal{I}(\theta_1,\theta_2)t$$

which yields

$$\frac{1}{\eta}\log\|\mu_t - \mathbf{e}_1\|_{TV} \leq -\mathcal{I}(\theta_1,\theta_2)t + \sqrt{2B^2 t\log\frac{m}{\delta_t}} + \frac{\log m}{\eta} \qquad (21)$$

with probability at least $1-\delta_t$. To have the above work for all $t \in [T]$ (simultaneously) with probability at least $1-\delta$, we need to take a union bound over any $t \in [T]$. Therefore, we have to choose $\{\delta_t\}_{t=1}^T$ such that $\sum_{t=1}^T \delta_t \leq \delta$. Letting $\delta_t := \delta \exp\{-t^{1/3}\}/6$, we have

$$\sum_{t=1}^{T} \delta_t \leq \frac{\delta}{6} \int_0^\infty \exp\left\{-t^{\frac{1}{3}}\right\} d_t$$

$$= \frac{\delta}{6} \int_0^\infty 3u^2 \exp\{-u\}\, d_u = \frac{\delta}{6} 3! = \delta. \tag{22}$$

Let us avoid notational clutter, by defining $a := \mathcal{I}(\theta_1, \theta_2)$, $b := (2B^2 \log(6m/\delta))^{1/2}$ and $c := \sqrt{2}B$, respectively. Also, define

$$t_1 := \max\left\{\left(\frac{3b}{a}\right)^2, \left(\frac{3c}{a}\right)^3\right\} \text{ and } t_2 := \frac{3}{a\eta}\log m.$$

Then, in view of (21) and Lemma 4, with probability at least $1-\delta_t$ we have

$$\begin{aligned} D_{KL}(\mu_{i,t} \| \mu_t) &\leq 2\|\mathbf{e}_1 - \mu_t\|_{TV} \\ &\leq 2m \exp\left\{\eta\left(-at + bt^{\frac{1}{2}} + ct^{\frac{2}{3}}\right)\right\} \\ &\leq 2m \exp\left\{-\frac{a}{3}\eta t\right\} \quad \text{for } t \geq t_1 \\ &\leq 2, \quad \text{for } t \geq t_2. \end{aligned}$$

Let $t_0 = \max\{t_1, t_2\}$, note all the inequalities above together, and observe the fact that $\|\mathbf{e}_1 - \mu_t\|_{TV} \leq 1$ for any $t \in [T]$. Also, recall the proper choice of δ_t for (22) to bound $\mathbf{Cost}_{i,T}$ as

$$\begin{aligned} \sum_{t=1}^{T} D_{KL}(\mu_{i,t} \| \mu_t) &\leq 2\sum_{t=1}^{t_0} \|\mathbf{e}_1 - \mu_t\|_{TV} \\ &\quad + 2\sum_{t=t_0+1}^{T} m \exp\left\{-\frac{a}{3}\eta t\right\} \\ &\leq 2t_0 + 2\sum_{t=t_2+1}^{T} m \exp\left\{-\frac{a}{3}\eta t\right\} \\ &\leq 2t_0 + 2\int_{t_2}^{\infty} m \exp\left\{-\frac{a}{3}\eta t\right\} d_t \\ &= 2t_0 + \frac{6}{a\eta} \end{aligned}$$

with probability at least $1-\delta$. Plugging our choice of η into above completes the proof. ∎

REFERENCES

[1] R. R. Tenney and N. R. Sandell, Jr., "Detection with distributed sensors," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 17, pp. 501–510, 1981.

[2] J. N. Tsitsiklis *et al.*, "Decentralized detection," *Adv. Stat. Signal Processing*, vol. 2, no. 2, pp. 297–344, 1993.

[3] V. Borkar and P. P. Varaiya, "Asymptotic agreement in distributed estimation," *IEEE Trans. Autom. Control*, vol. 27, no. 3, pp. 650–655, 1982.

[4] S. Kar, J. M. Moura, and K. Ramanan, "Distributed parameter estimation in sensor networks: Nonlinear observation models and imperfect communication," *IEEE Trans. Inform. Theory*, vol. 58, no. 6, pp. 3575–3605, 2012.

[5] O. Dekel, R. Gilad-Bachrach, O. Shamir, and L. Xiao, "Optimal distributed online prediction using mini-batches," *J. Machine Learning Res.*, vol. 13, no. 1, pp. 165–202, 2012.

[6] A. Nedic and A. Ozdaglar, "Distributed subgradient methods for multi-agent optimization," *IEEE Trans. Autom. Control*, vol. 54, no. 1, pp. 48–61, 2009.

[7] A. Nedic, A. Olshevsky, A. Ozdaglar, and J. N. Tsitsiklis, "On distributed averaging algorithms and quantization effects," *IEEE Trans. Autom. Control*, vol. 54, no. 11, pp. 2506–2517, 2009.

[8] J.-F. Chamberland and V. V. Veeravalli, "Decentralized detection in sensor networks," *IEEE Trans. Signal Processing*, vol. 51, no. 2, pp. 407–416, 2003.

[9] F. Bullo, J. Cortés, and S. Martinez, *Distributed Control of Robotic Networks: A Mathematical Approach to Motion Coordination Algorithms*. Princeton, NJ, USA: Princeton Univ. Press, 2009.

[10] N. A. Atanasov, J. Le Ny, and G. J. Pappas, "Distributed algorithms for stochastic source seeking with mobile robot networks," *J. Dyn. Syst., Meas., Control*, 2014. [Online]. Available: http://www.georgejpappas.org/papers/ASME.pdf.

[11] S. Shahrampour, S. Rakhlin, and A. Jadbabaie, "Online learning of dynamic parameters in social networks," *Advances in Neural Information Processing Systems*, 2013. [Online]. Available: http://machinelearning.wustl.edu/mlpapers/paper_files/NIPS2013_4976.pdf.

[12] J. N. Tsitsiklis, "Problems in Decentralized Decision Making and Computation," DTIC Document Tech. Rep., 1984.

[13] A. Jadbabaie, J. Lin, and A. S. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Trans. Autom. Control*, vol. 48, no. 6, pp. 988–1001, 2003.

[14] R. Olfati-Saber and R. M. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Trans. Autom. Control*, vol. 49, no. 9, pp. 1520–1533, 2004.

[15] F. S. Cattivelli and A. H. Sayed, "Distributed detection over adaptive networks using diffusion adaptation," *IEEE Trans. Signal Processing*, vol. 59, no. 5, pp. 1917–1932, 2011.

[16] D. Jakovetic, J. M. Moura, and J. Xavier, "Distributed detection over noisy networks: Large deviations analysis," *IEEE Trans. Signal Processing*, vol. 60, no. 8, pp. 4306–4320, 2012.

[17] D. Bajovic, D. Jakovetic, J. M. Moura, J. Xavier, and B. Sinopoli, "Large deviations performance of consensus+ innovations distributed detection with non-gaussian observations," *IEEE Trans. Signal Processing*, vol. 60, no. 11, pp. 5987–6002, 2012.

[18] A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi, "Non-Bayesian social learning," *Games Economic Behav.*, vol. 76, no. 1, pp. 210–225, 2012.

[19] S. Shahrampour and A. Jadbabaie, "Exponentially fast parameter estimation in networks using distributed dual averaging," in *Proc. IEEE Conf. Decision Control (CDC)*, 2013, pp. 6196–6201.

[20] S. Shahrampour, M. A. Rahimian, and A. Jadbabaie, "Switching to learn," in *Proc. Amer. Control Conf. (ACC)*, Jul. 2015, pp. 2918–2923.

[21] A. Lalitha, A. Sarwate, and T. Javidi, "Social learning and distributed hypothesis testing," in *Proc. Int. Symp. Inform. Theory (ISIT)*, 2014, pp. 551–555.

[22] K. Rahnama Rad and A. Tahbaz-Salehi, "Distributed parameter estimation in networks," in *Proc. IEEE Conf. Decision Control (CDC)*, 2010, pp. 5050–5055.

[23] J. C. Duchi, A. Agarwal, and M. J. Wainwright, "Dual averaging for distributed optimization: Convergence analysis and network scaling," *IEEE Trans. Autom. Control*, vol. 57, no. 3, pp. 592–606, 2012.

[24] A. Jadbabaie, P. Molavi, and A. Tahbaz-Salehi, "Information heterogeneity and the speed of learning in social networks," *Columbia Business School Research Paper*, no. 13–28, 2013. [Online]. Available: http://www.columbia.edu/~at2761/LearningRate.pdf.

[25] K. Drakopoulos, A. Ozdaglar, and J. N. Tsitsiklis, "On learning with finite memory," *IEEE Trans. Inform. Theory*, vol. 59, no. 10, pp. 6859–6872, 2013.
[26] J. S. Rosenthal, "Convergence rates for Markov chains," *SIAM Rev.*, vol. 37, no. 3, pp. 387–405, 1995.
[27] J. D. Abernethy, E. Hazan, and A. Rakhlin, "Interior-point methods for full-information and bandit online learning," *IEEE Trans. Inform. Theory*, vol. 58, no. 7, pp. 4164–4175, 2012.
[28] A. Nemirovskii and D. Yudin, *Problem Complexity and Method Efficiency in Optimization*. New York, NY, USA: Wiley, 1983.
[29] D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov Chains and Mixing Times*. Providence, RI, USA: American Mathematical Soc., 2009.
[30] S. Boyd, P. Diaconis, and L. Xiao, "Fastest mixing Markov chain on a graph," *SIAM Rev.*, vol. 46, no. 4, pp. 667–689, 2004.
[31] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Randomized gossip algorithms," *IEEE Trans. Inform. Theory*, vol. 52, no. 6, pp. 2508–2530, 2006.
[32] M. G. Rabbat, R. D. Nowak, and J. A. Bucklew, "Generalized consensus computation in networked systems with erasure links," in *Proc. IEEE Workshop Signal Processing Adv. Wireless Commun.*, 2005, pp. 1088–1092.
[33] Y. Hatano and M. Mesbahi, "Agreement over random networks," *IEEE Trans. Autom. Control*, vol. 50, no. 11, pp. 1867–1872, 2005.
[34] F. R. Chung,*Spectral Graph Theory*. Providence, RI, USA: American Mathematical Society, 1997, vol. 92.
[35] C. McDiarmid, "Concentration," in *Probabilistic Methods for Algorithmic Discrete Mathematics*. New York, NY, USA: Springer, 1998, pp. 195–248.

Shahin Shahrampour received the B.Sc. degree from Sharif University of Technology, Tehran, Iran, in 2009 and the M.S.E. degree in electrical engineering, the M.A. degree in statistics, and the Ph.D. degree in electrical and systems engineering from the University of Pennsylvania, Philadelphia, in 2012, 2014, and 2015, respectively.

He is currently a Postdoctoral Fellow in electrical engineering, Harvard University, Cambridge, MA. His research interests are in machine learning, online learning, optimization, and distributed algorithms.

Alexander Rakhlin received the B.A. degrees in mathematics and computer science from Cornell University, Ithaca, NY, in 2000, and the Ph.D. degree in computational neuroscience from MIT, Cambridge, MA, in 2006.

From 2006 to 2009 he was a Postdoctoral Fellow at the Department of Electrical Engineering and Computer Sciences, UC Berkeley. Since 2009, he has been on the faculty in the Department of Statistics, University of Pennsylvania, Philadelphia. His interests include machine learning, statistics, optimization, and game theory.

Ali Jadbabaie (F'14) received the B.S. degree (with high honors) from Sharif University of Technology, Tehran, Iran, the M.S. degree in electrical and computer engineering from the University of New Mexico in Albuquerque, and the Ph.D. degree in control and dynamical systems from the California Institute of Technology (Caltech), Pasadena.

He is currently on leave from the University of Pennsylvania, Philadelphia, at MIT, Cambridge, MA, as a Visiting Professor, where he is also the Interim Director of the Sociotechnical Systems Research Center (SSRC) and the Associate Director of the newly formed MIT Institute for Data, Systems and Society (IDSS). He was a Postdoctoral Scholar at Yale University, New Haven, CT, before joining the faculty of the University of Pennsylvania (Penn) in July 2002, where he is currently the Alfred Fitler Moore Professor of Network Science in the Department of Electrical and Systems Engineering. He holds secondary appointments in the Departments of Computer and Information Science as well as Operations, Information, and Decisions in the Wharton School. A faculty member in Penn's General Robotics, Automation, Sensing & Perception (GRASP) Lab, he is also the Co-Founder and Director of the Raj and Neera Singh Program in Networked & Social Systems Engineering (NETS) at Penn Engineering, a new undergraduate interdisciplinary degree program focused on network science and engineering, operations research, social phenomena and social networks. He is also a faculty member of The Warren Center for Network & Data Sciences at Penn and a faculty affiliate of the Center for Technology, Innovation and Competition at Penn Law. His current research interests include decision and control theory with a focus on distributed optimization and control, collective behavior, network science, and the study of social and economic networks.

Dr. Jadbabaie received the NSF Career Award, an ONR Young Investigator Award, the O. Hugo Schuck Best Paper Award from the American Automatic Control Council, and the George S. Axelby Best Paper Award from the IEEE Control Systems Society. He is the Inaugural Editor-in-Chief of the IEEE TRANSACTIONS ON NETWORK SCIENCE AND ENGINEERING, a new interdisciplinary journal sponsored by several IEEE Societies focused on science and engineering of networks and networked systems.