2019 American Control Conference (ACC)
Philadelphia, PA, USA, July 10-12, 2019

A New Approach for Distributed Hypothesis Testing with Extensions to Byzantine-Resilience

Aritra Mitra, John A. Richards and Shreyas Sundaram

***Abstract*— We study a setting where a group of agents, each receiving partially informative private observations, seek to collaboratively learn the true state (among a set of hypotheses) that explains their joint observation profiles over time. To solve this problem, we propose a distributed learning rule that differs fundamentally from existing approaches, in the sense that it does not employ any form of "belief-averaging". Specifically, every agent maintains a local belief on each hypothesis that is updated in a Bayesian manner without any network influence, and an actual belief that is updated (up to normalization) as the minimum of its own local belief and the actual beliefs of its neighbors. Under minimal requirements on the signal structures of the agents and the underlying communication graph, we establish consistency of the proposed belief update rule, i.e., we show that the actual beliefs of the agents asymptotically concentrate on the true state almost surely. As one of the key benefits of our approach, we show that our learning rule can be extended to scenarios that capture misbehavior on the part of certain agents in the network, modeled via the Byzantine adversary model. In particular, we prove that each non-adversarial agent can asymptotically learn the true state of the world almost surely, under appropriate conditions on the observation model and the network topology.**

I. Introduction

Various distributed learning problems arising in social networks (such as opinion formation and spreading), and in engineered systems (such as target recognition by a group of aerial robots) can be studied under the formal framework of distributed hypothesis testing. Within this framework, a group of agents repeatedly observe certain private signals, and aim to infer the "true state of the world" that explains their joint observations. While much of the earlier work on this topic assumed the existence of a centralized fusion center for performing computational tasks [1], more recent endeavors focus on a distributed setting where interactions among agents are captured by a communication graph [2]–[10]. Our work here falls in the latter class. A typical belief update rule in the distributed setting combines a local Bayesian update with a consensus-based opinion pooling of neighboring beliefs. Specifically, linear opinion pooling is studied in [2]–[4], whereas the log-linear form of belief aggregation is studied in the context of distributed hypothesis testing in [5]–[8], and distributed parameter estimation in [9], [10]. Notably, exponential convergence rates are achieved in [3], [5]–[7], while a finite-time analysis is presented in [8]. Extensions to time-varying graphs are studied in [4], [5].

In [5, Section III], the authors explain that the commonly studied linear and log-linear forms of belief aggregation are specific instances of a more general class of opinion pooling known as g-Quasi-Linear Opinion pools (g-QLOP), introduced in [11]. The main contribution of our paper is the development of a novel belief update rule that deviates fundamentally from the broad family of g-QLOP learning rules discussed above. Specifically, the learning algorithm that we propose in Section III-A does not rely on any linear consensus-based belief aggregation protocol. Instead, each agent maintains two sets of beliefs: a local belief that is updated in a Bayesian manner based on the private observations (without neighbor interactions), and an actual belief that is updated (up to normalization) as the minimum of the agent's own local belief and the actual beliefs of its neighbors. In Section V, we establish that under minimal requirements on the agents' signal structures and the communication graph, the actual beliefs of the agents asymptotically concentrate on the true state almost surely. In Section IV, we argue that our approach works under graph-theoretic conditions that are milder than the standard assumption of strong-connectivity.

In addition to the above contribution to the distributed hypothesis testing problem, we also show in this paper that our approach is capable of handling agents that do not follow the prescribed learning algorithm. Such agents may represent stubborn individuals or ideological extremists in the context of a social network, or model faults (either benign or malicious) in a networked control system. We ask: *In the presence of such misbehaving entities, how should the remaining agents process their private observations and the beliefs of their neighbors to eventually learn the truth?* To answer this question, we model misbehaving agents via the classical Byzantine adversary model, and develop a provably correct, resilient version of our proposed learning rule in Section III-B. The only related work (that we are aware of) in this regard is reported in [7]. As we discuss in Section III-B, our proposed approach is significantly less computationally intensive relative to those in [7]. We identify conditions on the observation model and the network structure that guarantee applicability of our Byzantine-resilient learning rule, and argue in Section IV that such conditions can be checked in polynomial time.

A. Mitra and S. Sundaram are with the School of Electrical and Computer Engineering at Purdue University. J. A. Richards is with Sandia National Laboratories. Email: {mitra14, sundara2}@purdue.edu, jaricha@sandia.gov. This work was supported in part by NSF CAREER award 1653648, and by a grant from Sandia National Laboratories. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. The views expressed in the article do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

II. Model and Problem Formulation

Network Model: We consider a group of agents $\mathcal{V} = \{1, 2, \ldots, n\}$ interacting over a time-invariant, directed communication graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. An edge $(i, j) \in \mathcal{E}$ indicates that agent i can directly transmit information to agent j. If $(i, j) \in \mathcal{E}$, then agent i will be called a neighbor of agent j, and agent j will be called an out-neighbor of agent i. The set of all neighbors of agent i will be denoted $\mathcal{N}_i$. Given two disjoint sets $\mathcal{C}_1, \mathcal{C}_2 \subseteq \mathcal{V}$, we say that $\mathcal{C}_2$ is reachable from $\mathcal{C}_1$ if for every $i \in \mathcal{C}_2$, there exists a directed path from some $j \in \mathcal{C}_1$ to agent i (note that j will in general be a function of i). We will use $|\mathcal{C}|$ to denote the cardinality of a set $\mathcal{C}$.

Observation Model: Let $\Theta = \{\theta_1, \theta_2, \ldots, \theta_m\}$ denote m possible states of the world; each $\theta_i \in \Theta$ will be called a hypothesis. Let $\mathbb{N}$ and $\mathbb{N}_+$ denote the set of non-negative integers and positive integers, respectively. Then at each time-step $t \in \mathbb{N}_+$, every agent $i \in \mathcal{V}$ privately observes a signal $s_{i,t} \in \mathcal{S}_i$, where $\mathcal{S}_i$ denotes the signal space of agent i. The joint observation profile so generated across the network is denoted $s_t = (s_{1,t}, s_{2,t}, \ldots, s_{n,t})$, where $s_t \in \mathcal{S}$, and $\mathcal{S} = \mathcal{S}_1 \times \mathcal{S}_2 \times \cdots \times \mathcal{S}_n$. The signal s_t is generated based on a conditional likelihood function $l(\cdot|\theta^\star)$, governed by the true state of the world $\theta^\star \in \Theta$. Let $l_i(\cdot|\theta^\star), i \in \mathcal{V}$ denote the i-th marginal of $l(\cdot|\theta^\star)$. The signal structure of each agent $i \in \mathcal{V}$ is then characterized by a family of parameterized marginals $\{l_i(w_i|\theta) : \theta \in \Theta, w_i \in \mathcal{S}_i\}$.

We make the following standard assumptions [2]–[8]: (i) The signal space of each agent i, namely $\mathcal{S}_i$, is finite. (ii) Each agent i has knowledge of its local likelihood functions $\{l_i(\cdot|\theta_p)\}_{p=1}^m$, and it holds that $l_i(w_i|\theta) > 0, \forall w_i \in \mathcal{S}_i$, and $\forall \theta \in \Theta$. (iii) The observation sequence of each agent is described by an i.i.d. random process over time; however, at any given time-step, the observations of different agents may potentially be correlated. (iv) There exists a fixed true state of the world $\theta^\star \in \Theta$ (unknown to the agents) that generates the observations of all the agents. Finally, we define a probability triple $(\Omega, \mathcal{F}, \mathbb{P}^{\theta^\star})$, where $\Omega \triangleq \{\omega : \omega = (s_1, s_2, \ldots), \forall s_t \in \mathcal{S}, \forall t \in \mathbb{N}_+\}$, $\mathcal{F}$ is the σ-algebra generated by the observation profiles, and $\mathbb{P}^{\theta^\star}$ is the probability measure induced by sample paths in Ω. Specifically, $\mathbb{P}^{\theta^\star} = \prod_{t=1}^{\infty} l(\cdot|\theta^\star)$. We will use the abbreviation a.s. to indicate that an event occurs almost surely w.r.t. the probability measure $\mathbb{P}^{\theta^\star}$.

Given the above setup, the goal of each agent in the network is to discern the true state of the world $\theta^\star$. The challenge associated with such a task stems from the fact that the private signal structure of any given agent is in general only partially informative. To make this notion precise, define $\Theta_i^{\theta^\star} \triangleq \{\theta \in \Theta : l_i(w_i|\theta) = l_i(w_i|\theta^\star), \forall w_i \in \mathcal{S}_i\}$. In words, $\Theta_i^{\theta^\star}$ represents the set of hypotheses that are *observationally equivalent* to the true state $\theta^\star$ from the perspective of agent i. In general, for any agent $i \in \mathcal{V}$, we may have $|\Theta_i^{\theta^\star}| > 1$, necessitating collaboration among agents. While inter-agent collaboration is implicitly assumed in the related literature, in this paper we will also allow misbehavior on the part of certain agents, modeled as follows.

Adversary Model: We assume that a certain fraction of the agents are adversarial, and model their behavior based on the Byzantine fault model [12]. In particular, Byzantine agents possess complete knowledge of the observation model, the network model, the algorithms being used, the information being exchanged, and the true state of the world. Leveraging such information, adversarial agents can behave arbitrarily and in a coordinated manner, and can in particular, send incorrect, potentially inconsistent information to their out-neighbors. We will consider an f-local adversarial model, i.e., we assume that there are at most f adversaries in the neighborhood of any non-adversarial agent. As is fairly standard in the distributed fault-tolerant literature [13]–[17], we only assume that non-adversarial agents know the upper bound f on the number of adversaries in their neighborhood, but are otherwise unaware of their identities. The adversarial set will be denoted by $\mathcal{A} \subset \mathcal{V}$, and the remaining agents $\mathcal{R} = \mathcal{V} \setminus \mathcal{A}$ will be called the regular agents.

Our **objective** in this paper will be to design a distributed learning rule that allows each regular agent $i \in \mathcal{R}$ to identify the true state of the world almost surely, despite (i) the partially informative signal structures of the agents, and (ii) the actions of any f-local Byzantine adversarial set. To this end, we introduce the following notion of *source agents*.

Definition 1. *(**Source agents**) An agent i is said to be a source agent for a pair of distinct hypotheses $\theta_p, \theta_q \in \Theta$, if $D(l_i(\cdot|\theta_p)||l_i(\cdot|\theta_q)) > 0$, where $D(l_i(\cdot|\theta_p)||l_i(\cdot|\theta_q))$ represents the KL-divergence between the distributions $l_i(\cdot|\theta_p)$ and $l_i(\cdot|\theta_q)$. The set of all source agents for the pair θ_p, θ_q is denoted by $\mathcal{S}(\theta_p, \theta_q)$.* □

In words, a source agent for a pair $\theta_p, \theta_q \in \Theta$ is an agent that can distinguish between the pair of hypotheses θ_p, θ_q based on its private signal structure. In our developments, we will require the following two definitions.

Definition 2. *(r-**reachable set**) [14] For a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, a set $\mathcal{C} \subseteq \mathcal{V}$, and an integer $r \in \mathbb{N}_+$, $\mathcal{C}$ is an r-reachable set if there exists an $i \in \mathcal{C}$ such that $|\mathcal{N}_i \setminus \mathcal{C}| \geq r$.* □

Definition 3. *(**strongly r-robust graph** w.r.t. $\mathcal{S}(\theta_p, \theta_q)$) For $r \in \mathbb{N}_+$ and $\theta_p, \theta_q \in \Theta$, a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is strongly r-robust w.r.t. the set of source agents $\mathcal{S}(\theta_p, \theta_q)$, if for every non-empty subset $\mathcal{C} \subseteq \mathcal{V} \setminus \mathcal{S}(\theta_p, \theta_q)$, $\mathcal{C}$ is r-reachable.* □

III. Proposed Learning Rules

A. A Novel Belief Update Rule

In this section, we propose a novel belief update rule and discuss the intuition behind it. To introduce the key ideas underlying our basic approach, we first consider a scenario where all agents are regular, i.e., $\mathcal{R} = \mathcal{V}$. Every agent i maintains and updates (at every time-step) two separate sets of belief vectors, namely, $\boldsymbol{\pi}_{i,t}$ and $\boldsymbol{\mu}_{i,t}$. Each of these vectors are probability distributions over the hypothesis set Θ. We will refer to $\boldsymbol{\pi}_{i,t}$ and $\boldsymbol{\mu}_{i,t}$ as the "local" belief vector (for reasons that will soon become obvious), and the "actual" belief vector, respectively, maintained by agent i. The **goal** of each agent $i \in \mathcal{V}$ in the network will be to use its own private signals, and the information available from its neighbors, to

update $\boldsymbol{\mu}_{i,t}$ sequentially so that $\lim_{t\to\infty}\mu_{i,t}(\theta^*)=1$ almost surely. To do so, for each $\theta\in\Theta$, and at each time-step $t+1, t\in\mathbb{N}$, agent i first generates $\pi_{i,t+1}(\theta)$ via a local Bayesian update rule that incorporates the private observation $s_{i,t+1}$ using $\pi_{i,t}(\theta)$ as a prior. Having generated $\pi_{i,t+1}(\theta)$, agent i updates $\mu_{i,t+1}(\theta)$ (up to normalization) by setting it to be the minimum of its locally generated belief $\pi_{i,t+1}(\theta)$, and the actual beliefs $\mu_{j,t}(\theta), j\in\mathcal{N}_i$ of its neighbors at the previous time-step. It then reports its actual belief $\mu_{i,t+1}(\theta)$ to each of its out-neighbors.[1] The belief vectors are initialized as $\mu_{i,0}(\theta)>0, \pi_{i,0}(\theta)>0, \forall\theta\in\Theta, \forall i\in\mathcal{V}$. Subsequently, these vectors are updated at each time-step $t+1, t\in\mathbb{N}$ as:

- **Step 1: Update of the local beliefs**:

$$\pi_{i,t+1}(\theta)=\frac{l_i(s_{i,t+1}|\theta)\pi_{i,t}(\theta)}{\sum_{p=1}^{m} l_i(s_{i,t+1}|\theta_p)\pi_{i,t}(\theta_p)}. \quad (1)$$

- **Step 2: Update of the actual beliefs**:

$$\mu_{i,t+1}(\theta)=\frac{\min\{\{\mu_{j,t}(\theta)\}_{j\in\mathcal{N}_i},\pi_{i,t+1}(\theta)\}}{\sum_{p=1}^{m}\min\{\{\mu_{j,t}(\theta_p)\}_{j\in\mathcal{N}_i},\pi_{i,t+1}(\theta_p)\}}. \quad (2)$$

Intuition behind the learning rule: Consider the set of source agents $\mathcal{S}(\theta^*,\theta)$ who can differentiate between a certain false hypothesis θ and the true state $\theta^\star$. We ask: how do the agents in the set $\mathcal{S}(\theta^\star,\theta)$ contribute to the process of collaborative learning? To answer this question, we note that the signal structures of such agents are rich enough for them to be able to eliminate θ on their own, i.e., without the support of their neighbors. Thus, the agents in $\mathcal{S}(\theta^\star,\theta)$ should contribute towards driving the actual beliefs of their out-neighbors (and eventually, of all the agents in the set $\mathcal{V}\backslash\mathcal{S}(\theta^\star,\theta)$) on the hypothesis θ to zero. To achieve the above objective, we are especially interested in devising a rule that ensures that the capability of the source agents $\mathcal{S}(\theta^\star,\theta)$ to eliminate θ is not diminished due to neighbor interactions. It is precisely these considerations that motivate us to employ (i) an auxiliary belief vector $\boldsymbol{\pi}_{i,t+1}$ generated via local processing (i.e., without any network influence) of the private signals, and (ii) a min-rule of the form (2). Specifically, if $i\in\mathcal{S}(\theta^\star,\theta)$, then the sequence of local beliefs $\pi_{i,t+1}(\theta)$ will almost surely converge to 0 based on the update rule (1). Hence, for a source agent $i\in\mathcal{S}(\theta^\star,\theta)$, $\pi_{i,t+1}(\theta)$ will play the key role of an external network-independent input in the min-rule (2) that triggers a process of belief reduction on the hypothesis θ originating at the source set $\mathcal{S}(\theta^\star,\theta)$, and eventually propagating via the proposed min-rule to each agent in the network reachable from $\mathcal{S}(\theta^\star,\theta)$. The above discussion will be made precise in Section V.

Remark 1. *Note that our proposed algorithm does not employ any form of "belief-averaging" unlike existing approaches to distributed hypothesis testing that rely either on linear opinion pooling [2]–[4], or log-linear opinion pooling [5]–[10]. As such, the lack of linearity in our belief update rule precludes (direct or indirect) adaptation of existing analysis techniques to suit our needs. Consequently, we develop a novel sample path based proof technique in Section V to establish consistency of the proposed learning rule.* □

[1]Note that based on our algorithm, agents only exchange their actual beliefs, and not their local beliefs.

B. A Byzantine-Resilient Belief Update Rule

As pointed out in the Introduction, a key benefit of our approach is that it can be extended to account for the worst-case Byzantine adversarial model described in Section II. A standard way to analyze the impact of such adversarial agents while designing resilient distributed consensus-based protocols (for applications in consensus [13], [14], optimization [15], [16] and hypothesis testing [7]) is to to express the iterates of a regular agent as a convex combination of the iterates of its regular neighbors, based on appropriate filtering techniques, and under certain assumptions on the network structure. While this can indeed be achieved efficiently for scalar consensus problems, for problems requiring consensus on vectors (like the belief vectors in our setting), such an approach becomes computationally prohibitive [7]. To bypass such heavy computations, and yet accommodate Byzantine attacks, we now develop a resilient version of the learning rule introduced in Section III-A, as follows. Each agent $i\in\mathcal{R}$ acts as follows at every time-step $t+1$ (where $t\in\mathbb{N}$).

- **Step 1: Update of the local beliefs**: The local belief $\pi_{i,t+1}(\theta)$ is updated as before, based on (1).

- **Step 2: Filtering extreme beliefs**: If $|\mathcal{N}_i|\geq(2f+1)$, then agent i performs a filtering operation as follows. It collects the actual beliefs $\mu_{j,t}(\theta)$ from each neighbor $j\in\mathcal{N}_i$ and sorts them from highest to lowest. It rejects the highest f and the lowest f of such beliefs (i.e., it throws away $2f$ beliefs in all). In other words, for each hypothesis, a regular agent retains only the moderate beliefs received from its neighbors.

- **Step 3: Update of the actual beliefs**: If $|\mathcal{N}_i|\geq(2f+1)$, then agent i updates $\mu_{i,t+1}(\theta)$ as follows. Let the set of neighbors whose beliefs on θ are not rejected by agent i (based on the previous filtering step) be denoted by $\mathcal{M}_{i,t}^{\theta}\subset\mathcal{N}_i$. The actual belief $\mu_{i,t+1}(\theta)$ is then updated as follows:

$$\mu_{i,t+1}(\theta)=\frac{\min\{\{\mu_{j,t}(\theta)\}_{j\in\mathcal{M}_{i,t}^{\theta}},\pi_{i,t+1}(\theta)\}}{\sum_{p=1}^{m}\min\{\{\mu_{j,t}(\theta_p)\}_{j\in\mathcal{M}_{i,t}^{\theta_p}},\pi_{i,t+1}(\theta_p)\}}. \quad (3)$$

If $|\mathcal{N}_i|<(2f+1)$, then agent i updates $\mu_{i,t+1}(\theta)$ as follows:

$$\mu_{i,t+1}(\theta)=\pi_{i,t+1}(\theta). \quad (4)$$

Agent i transmits $\mu_{i,t+1}(\theta)$ to each of its out-neighbors on completion of the above steps. We will refer to the above sequence of actions as the Local-Filtering based Resilient Hypothesis Elimination (LFRHE) algorithm.

IV. Main Results and Discussion

The main results of the paper are as follows.

Theorem 1. *Suppose $\mathcal{R} = \mathcal{V}$, and that the following hold:*

(i) *For every pair of hypotheses $\theta_p, \theta_q \in \Theta$, the corresponding source set $\mathcal{S}(\theta_p, \theta_q)$ is non-empty.*

(ii) *For every pair of hypotheses $\theta_p, \theta_q \in \Theta$, $\mathcal{V} \setminus \mathcal{S}(\theta_p, \theta_q)$ is reachable from the source set $\mathcal{S}(\theta_p, \theta_q)$.*

(iii) *Agents have non-zero prior beliefs on each hypothesis, i.e., $\pi_{i,0}(\theta) > 0, \mu_{i,0}(\theta) > 0, \forall i \in \mathcal{V}, \forall \theta \in \Theta$.*

Then, the learning rule described by equations (1) and (2) guarantees that $\mu_{i,t}(\theta^\star) \to 1$ a.s. $\forall i \in \mathcal{V}$. □

Theorem 2. *Suppose the following are true:*

(i) *For every pair of hypotheses $\theta_p, \theta_q \in \Theta$, the graph $\mathcal{G}$ is strongly $(2f+1)$-robust w.r.t. the corresponding source set $\mathcal{S}(\theta_p, \theta_q)$.*

(ii) *Each regular agent $i \in \mathcal{R}$ has a non-zero prior belief on each hypothesis, i.e., $\pi_{i,0}(\theta) > 0, \mu_{i,0}(\theta) > 0$ for all $i \in \mathcal{R}$, and for all $\theta \in \Theta$.*

Then, the LFRHE algorithm described by equations (1), (3) and (4) guarantees that $\mu_{i,t}(\theta^\star) \to 1$ a.s. $\forall i \in \mathcal{R}$, despite the actions of any f-local set of Byzantine adversaries. □

Remark 2. *For any pair $\theta_p, \theta_q \in \Theta$, notice that condition (i) of Theorem 2 (together with the definition of strong-robustness in Def. 3) requires $|\mathcal{S}(\theta_p, \theta_q)| \geq (2f+1)$, if $\mathcal{V} \setminus \mathcal{S}(\theta_p, \theta_q)$ is non-empty.* □

Remark 3. *While the first condition in Theorem 1 is a basic global identifiability condition, the second condition on the network structure is in general weaker than the standard assumption of strong-connectivity made in [2], [3], [6], [8]–[10]. To see this, consider a scenario where $\Theta = \{\theta_1, \theta_2\}$. Clearly, any agent $i \in \mathcal{S}(\theta_1, \theta_2)$ can discern the true state without neighbor interactions, precluding the need for incoming edges to such agents.* □

Remark 4. *The first condition in Theorem 2 blends requirements on the signal structures of the agents with those on the communication graph. To gain intuition about this condition, suppose $\Theta = \{\theta_1, \theta_2\}$, and consider an agent $i \in \mathcal{V} \setminus \mathcal{S}(\theta_1, \theta_2)$. To enable i to learn the truth despite potential adversaries in its neighborhood, one requires (i) redundancy in the signal structures of the agents (see Remark 2), and (ii) redundancy in the network structure to ensure reliable information flow from $\mathcal{S}(\theta_1, \theta_2)$ to agent i. These requirements are captured by condition (i). For a fixed source set $\mathcal{S}(\theta_p, \theta_q)$, checking whether $\mathcal{G}$ is strongly $(2f+1)$-robust w.r.t. $\mathcal{S}(\theta_p, \theta_q)$ can be done in polynomial time [17]. Since the source sets for each pair $\theta_p, \theta_q \in \Theta$ can also be computed in polynomial time via a simple inspection of the agents' signal structures, it follows that condition (i) in Theorem 2 can be checked in polynomial time.* □

V. Proofs of the Main Results

We start with the following simple lemma that characterizes the asymptotic behavior of the local belief sequences generated based on (1); for a proof, see [18].

Lemma 1. *Consider an agent $i \in \mathcal{S}(\theta^\star, \theta) \cap \mathcal{R}$. Suppose $\pi_{i,0}(\theta^\star) > 0$. Then, the update rule (1) ensures that (i) $\pi_{i,t}(\theta) \to 0$ a.s., and (ii) $\pi_{i,\infty}(\theta^\star) \triangleq \lim_{t\to\infty} \pi_{i,t}(\theta^\star)$ exists a.s., and satisfies $\pi_{i,\infty}(\theta^\star) \geq \pi_{i,0}(\theta^\star)$.* □

We now sketch the proofs of Theorems 1 and 2; details can be found in [18].

A. Proof of Theorem 1 (Sketch)

Proof. Let $\bar{\Omega} \subseteq \Omega$ denote the set of sample paths along which for each agent $i \in \mathcal{V}$, the following hold: (i) for each $\theta \in \Theta \setminus \Theta_i^{\theta^\star}$, $\pi_{i,t}(\theta) \to 0$, and (ii) $\pi_{i,\infty}(\theta^\star) \triangleq \lim_{t\to\infty} \pi_{i,t}(\theta^\star)$ exists, and satisfies $\pi_{i,\infty}(\theta^\star) \geq \pi_{i,0}(\theta^\star)$. Based on condition (iii) in Theorem 1, and Lemma 1, we infer that $\bar{\Omega}$ has measure 1. Thus, to prove the desired result, it suffices to confine our attention to $\bar{\Omega}$. Specifically, fix any sample path $\omega \in \bar{\Omega}$, and pick any $\epsilon > 0$. Our goal will be to establish that along the sample path ω, there exists $t(\omega, \epsilon)$ such that for all $t \geq t(\omega, \epsilon)$, $\mu_{i,t}(\theta) < \epsilon, \forall i \in \mathcal{V}, \forall \theta \neq \theta^\star$. We complete the proof in the following two steps.

Step 1: *Lower bounding the actual beliefs on the true state*: Define $\gamma_1 \triangleq \min_{i\in\mathcal{V}} \pi_{i,0}(\theta^\star)$ and notice that $\gamma_1 > 0$ based on condition (iii) of the theorem. Given the choice of the sample path ω, we notice that $\pi_{i,\infty}(\theta^\star)$ exists for each $i \in \mathcal{V}$, and that $\pi_{i,\infty}(\theta^\star) \geq \gamma_1$. Pick a small number $\delta > 0$ such that $\delta < \gamma_1$. The following statement is then immediate. There exists a time-step $\bar{t}_1(\omega, \delta)$ such that for all $t \geq \bar{t}_1(\omega, \delta)$, $\pi_{i,t}(\theta^\star) \geq \gamma_1 - \delta > 0, \forall i \in \mathcal{V}$. Now define $\gamma_2(\omega) \triangleq \min_{i\in\mathcal{V}}\{\mu_{i,\bar{t}_1(\omega,\delta)}(\theta^\star)\}$. We claim $\gamma_2(\omega) > 0$. To see this, observe that given the assumption of non-zero prior beliefs on the true state, and the structure of the proposed min-rule (2), $\gamma_2(\omega)$ can be 0 if and only if there exists some time-step $t'(\omega) \leq \bar{t}_1(\omega, \delta)$ such that $\pi_{i,t'(\omega)}(\theta^\star) = 0$, for some $i \in \mathcal{V}$. However, given the structure of the local Bayesian update rule (1), we would then have $\pi_{i,t}(\theta^\star) = 0$, for all $t \geq t'(\omega)$, contradicting the previously established fact that $\pi_{i,t}(\theta^\star) \geq \gamma_1 - \delta > 0, \forall t \geq \bar{t}_1(\omega, \delta) \geq t'(\omega), \forall i \in \mathcal{V}$. Having thus established that $\gamma_2(\omega) > 0$, define $\eta(\omega) \triangleq \min\{\gamma_1 - \delta, \gamma_2(\omega)\} > 0$. In other words, $\eta(\omega)$ lower-bounds the lowest belief (considering both local and actual beliefs) on the true state $\theta^\star$ held by an agent at time-step $\bar{t}_1(\omega, \delta)$. We claim the following:

$$\mu_{i,t}(\theta^\star) \geq \eta(\omega), \forall t \geq \bar{t}_1(\omega, \delta), \forall i \in \mathcal{V}. \tag{5}$$

To see why (5) is true, fix an agent $i \in \mathcal{V}$, and observe that based on (2):

$$\mu_{i,\bar{t}_1(\omega,\delta)+1}(\theta^\star) \overset{(a)}{\geq} \frac{\eta(\omega)}{\sum_{p=1}^{m} \min\{\{\mu_{j,\bar{t}_1(\omega,\delta)}(\theta_p)\}_{j\in\mathcal{N}_i}, \pi_{i,\bar{t}_1(\omega,\delta)+1}(\theta_p)\}} \geq \frac{\eta(\omega)}{\sum_{p=1}^{m} \pi_{i,\bar{t}_1(\omega,\delta)+1}(\theta_p)} \overset{(b)}{=} \eta(\omega), \tag{6}$$

where (a) follows from the way $\eta(\omega)$ is defined and by noting that $\pi_{i,t}(\theta^\star) \geq \eta(\omega), \forall t \geq \bar{t}_1(\omega, \delta), \forall i \in \mathcal{V}$, and (b) follows by noting that the local belief vectors generated via (1) (at each time-step) are valid probability distributions

over Θ, and hence $\sum_{p=1}^{m} \pi_{i,\bar{t}_1(\omega,\delta)+1}(\theta_p) = 1$. Since the above reasoning applies to each $i \in \mathcal{V}$, (5) follows via induction.

Step 2: *Upper bounding the actual beliefs on each false hypothesis*: Given an $\epsilon > 0$, pick a small $\bar{\epsilon}(\omega) > 0$ such that $\bar{\epsilon}(\omega) < \min\{\eta(\omega), \epsilon\}$. Fix a false hypothesis $\theta \neq \theta^\star$. By virtue of condition (i) of the theorem, we know that $|\mathcal{S}(\theta^\star, \theta)| > 0$. Let $q = d(\mathcal{G}) + 2$, where $d(\mathcal{G})$ represents the diameter of the graph $\mathcal{G}$. Then, based on Lemma 1, for each $i \in \mathcal{S}(\theta^\star, \theta)$, there exists $t_i^\theta(\omega, \bar{\epsilon}(\omega))$ such that for all $t \geq t_i^\theta(\omega, \bar{\epsilon}(\omega))$, $\pi_{i,t}(\theta) \leq \bar{\epsilon}^q(\omega)$. Define

$$\bar{t}_2^\theta(\omega, \delta, \bar{\epsilon}(\omega)) \triangleq \max\{\bar{t}_1(\omega, \delta), \max_{i \in \mathcal{S}(\theta^\star,\theta)} \{t_i^\theta(\omega, \bar{\epsilon}(\omega))\}\}. \tag{7}$$

Throughout the rest of the proof, we suppress the dependence of $\bar{t}_2$ on θ, ω, δ and $\bar{\epsilon}(\omega)$ to avoid cluttering the exposition. For any agent $i \in \mathcal{S}(\theta^\star, \theta)$, (2) yields:

$$\begin{aligned} \mu_{i,\bar{t}_2+1}(\theta) &\overset{(a)}{\leq} \frac{\bar{\epsilon}^q(\omega)}{\min\{\{\mu_{j,\bar{t}_2}(\theta^\star)\}_{j \in \mathcal{N}_i}, \pi_{i,\bar{t}_2+1}(\theta^\star)\}} \\ &\overset{(b)}{\leq} \frac{\bar{\epsilon}^q(\omega)}{\eta(\omega)} \overset{(c)}{<} \bar{\epsilon}^{(q-1)}(\omega) \leq \bar{\epsilon}(\omega) < \epsilon, \end{aligned} \tag{8}$$

where (a) follows from the fact that for each $i \in \mathcal{S}(\theta^\star, \theta)$, we have $\pi_{i,t}(\theta) \leq \bar{\epsilon}^q(\omega), \forall t \geq \bar{t}_2$, (b) follows from (5) and (7), and (c) follows from the way $\bar{\epsilon}(\omega)$ has been chosen. The chain of reasoning used to arrive at (8) applies to subsequent time-steps as well, thereby yielding:

$$\mu_{i,t}(\theta) < \bar{\epsilon}^{(q-1)}(\omega), \forall t \geq \bar{t}_2 + 1, \forall i \in \mathcal{S}(\theta^\star, \theta). \tag{9}$$

We now wish to investigate how the effect of (9) propagates through the rest of the network. If $\mathcal{V} \setminus \mathcal{S}(\theta^\star, \theta)$ is empty, then we have reached the desired conclusion w.r.t. the false hypothesis θ. If not, define

$$\mathcal{L}_1^{(\theta^\star,\theta)} \triangleq \{i \in \{\mathcal{V} \setminus \mathcal{S}(\theta^\star, \theta)\} : |\mathcal{N}_i \cap \mathcal{S}(\theta^\star, \theta)| > 0\} \tag{10}$$

as the set of out-neighbors of the source set $\mathcal{S}(\theta^\star, \theta)$. Condition (ii) of the theorem implies that if $\mathcal{V} \setminus \mathcal{S}(\theta^\star, \theta)$ is non-empty, then so is $\mathcal{L}_1^{(\theta^\star,\theta)}$. Consider any agent $i \in \mathcal{L}_1^{(\theta^\star,\theta)}$. By definition, i has a neighbor in $\mathcal{S}(\theta^\star, \theta)$ satisfying (9). This observation, coupled with equations (5), (7), and arguments similar to those used to arrive at (8), yields:

$$\mu_{i,t}(\theta) < \bar{\epsilon}^{(q-2)}(\omega), \forall t \geq \bar{t}_2 + 2, \forall i \in \mathcal{L}_1^{(\theta^\star,\theta)}. \tag{11}$$

With $\mathcal{L}_0^{(\theta^\star,\theta)} \triangleq \mathcal{S}(\theta^\star, \theta)$, define the sets $\mathcal{L}_r^{(\theta^\star,\theta)}, 1 \leq r \leq d(\mathcal{G})$ recursively as follows:

$$\mathcal{L}_r^{(\theta^\star,\theta)} \triangleq \{i \in \mathcal{V} \setminus \{\bigcup_{c=0}^{r-1} \mathcal{L}_c^{(\theta^\star,\theta)}\} : |\mathcal{N}_i \cap \{\bigcup_{c=0}^{r-1} \mathcal{L}_c^{(\theta^\star,\theta)}\}| > 0\}. \tag{12}$$

Whenever $\mathcal{V} \setminus \{\bigcup_{c=0}^{r-1} \mathcal{L}_c^{(\theta^\star,\theta)}\}$ is non-empty, condition (ii) of the theorem implies that $\mathcal{L}_r^{(\theta^\star,\theta)}$ is also non-empty. One can then easily verify via induction on r that:

$$\mu_{i,t}(\theta) < \bar{\epsilon}^{(q-(r+1))}(\omega), \forall t \geq \bar{t}_2 + (r+1), \forall i \in \mathcal{L}_r^{(\theta^\star,\theta)}, \tag{13}$$

where $1 \leq r \leq d(\mathcal{G})$. Noting that $q = d(\mathcal{G}) + 2$, we conclude $\mu_{i,t}(\theta) < \bar{\epsilon}(\omega) < \epsilon$, $\forall t \geq \bar{t}_2 + d(\mathcal{G}) + 1, \forall i \in \mathcal{V}$. An identical argument applies to every false hypothesis $\theta \neq \theta^\star$. □

B. Proof of Theorem 2 (Sketch)

Proof. Consider an f-local adversarial set $\mathcal{A} \subset \mathcal{V}$, and let $\mathcal{R} = \mathcal{V} \setminus \mathcal{A}$. We study two separate cases.

Case 1: Consider a regular agent $i \in \mathcal{R}$ such that $|\mathcal{N}_i| < (2f+1)$. One can show that condition (i) of the theorem implies $i \in \mathcal{S}(\theta_p, \theta_q)$, for every pair $\theta_p, \theta_q \in \Theta$. It then follows from Lemma 1 that such an agent can learn the true state $\theta^\star$ by simply updating its beliefs based on (1) and (4).

Case 2: We now focus only on regular agents i satisfying $|\mathcal{N}_i| \geq (2f+1)$. A key property of the proposed LFRHE algorithm that will be used throughout the proof is as follows. For any $i \in \mathcal{R}$, and any $\theta \in \Theta$, the filtering operation of the LFRHE algorithm ensures that at each time-step $t \in \mathbb{N}$:

$$\mu_{j,t}(\theta) \in Conv(\Psi_{i,t}^\theta), \forall j \in \mathcal{M}_{i,t}^\theta, \tag{14}$$

where

$$\Psi_{i,t}^\theta \triangleq \{\mu_{j,t}(\theta) : j \in \mathcal{N}_i \cap \mathcal{R}\}, \tag{15}$$

and $Conv(\Psi_{i,t}^\theta)$ is used to denote the convex hull formed by the points in the set $\Psi_{i,t}^\theta$. To see why (14) is true, partition the neighbor set $\mathcal{N}_i$ of a regular agent into three sets $\mathcal{U}_{i,t}^\theta, \mathcal{M}_{i,t}^\theta$, and $\mathcal{J}_{i,t}^\theta$ as follows. Sets $\mathcal{U}_{i,t}^\theta$ and $\mathcal{J}_{i,t}^\theta$ are each of cardinality f, and contain neighbors of agent i that transmit the highest f and the lowest f actual beliefs respectively, on the hypothesis θ, to agent i at time-step t. The set $\mathcal{M}_{i,t}^\theta$ contains the remaining neighbors of agent i, and is non-empty at every time-step since $|\mathcal{N}_i| \geq (2f+1)$. If $\mathcal{M}_{i,t}^\theta \cap \mathcal{A} = \emptyset$, then (14) holds trivially. Thus, consider the case when there are adversaries in the set $\mathcal{M}_{i,t}^\theta$, i.e., $\mathcal{M}_{i,t}^\theta \cap \mathcal{A} \neq \emptyset$. Given the f-locality of the adversarial model, and the nature of the filtering operation in the LFRHE algorithm, we infer that for each $j \in \mathcal{M}_{i,t}^\theta \cap \mathcal{A}$, there exist regular agents $u, v \in \mathcal{N}_i \cap \mathcal{R}$, such that $u \in \mathcal{U}_{i,t}^\theta$, $v \in \mathcal{J}_{i,t}^\theta$, and $\mu_{v,t}(\theta) \leq \mu_{j,t}(\theta) \leq \mu_{u,t}(\theta)$. This establishes our claim regarding equation (14).

Our goal will be to now establish each of the two steps in the proof of Theorem 1. To this end, let $\bar{\Omega} \subseteq \Omega$ denote the set of sample paths along which for each agent $i \in \mathcal{R}$, the following hold: (i) for each $\theta \in \Theta \setminus \Theta_i^{\theta^\star}$, $\pi_{i,t}(\theta) \to 0$, and (ii) $\pi_{i,\infty}(\theta^\star) \triangleq \lim_{t\to\infty} \pi_{i,t}(\theta^\star)$ exists, and satisfies $\pi_{i,\infty}(\theta^\star) \geq \pi_{i,0}(\theta^\star)$. Based on condition (ii) of the theorem, and Lemma 1, we infer that $\mathbb{P}^{\theta^\star}(\bar{\Omega}) = 1$. Now fix a sample path $\omega \in \bar{\Omega}$, and pick $\epsilon > 0$. Define $\gamma_1 = \min_{i \in \mathcal{R}} \pi_{i,0}(\theta^\star)$, pick a small number $\delta > 0$ satisfying $\delta < \gamma_1$, and observe that there exists $\bar{t}_1(\omega, \delta)$ such that $\pi_{i,t}(\theta^\star) \geq \gamma_1 - \delta > 0, \forall t \geq \bar{t}_1(\omega, \delta), \forall i \in \mathcal{R}$. Define $\gamma_2(\omega) \triangleq \min_{i \in \mathcal{R}} \{\mu_{i,\bar{t}_1(\omega,\delta)}(\theta^\star)\}$. As before, we claim $\gamma_2(\omega) > 0$. To see this, suppose by way of contradiction that there exists a time-step $t'(\omega)$ satisfying:

$$t'(\omega) = \min\{t \in \mathbb{N} : \exists i \in \mathcal{R} \text{ with } \mu_{i,t}(\theta^\star) = 0\}. \tag{16}$$

Clearly, $t'(\omega) \neq 0$ based on condition (ii) of the theorem. Suppose $t'(\omega)$ is some positive integer, and focus on how agent i updates $\mu_{i,t'(\omega)}(\theta^\star)$ based on (3). Following similar arguments as in the proof of Theorem 1, we know that $\pi_{i,t}(\theta^\star) > 0, \forall t \in \mathbb{N}, \forall i \in \mathcal{R}$. At the same time, every belief featuring in the set $\Psi_{i,t'(\omega)-1}^{\theta^\star}$ (as defined in equation (15)) is strictly positive based on the way $t'(\omega)$ is defined. The

above arguments coupled with (14), (15) readily imply that $\mu_{i,t'(\omega)}(\theta^\star) > 0$, yielding the desired contradiction.[2] With $\eta(\omega) \triangleq \min\{\gamma_1 - \delta, \gamma_2(\omega)\} > 0$, equations (14), (15), and arguments similar to those used to arrive at (6) yield

$$\mu_{i,t}(\theta^\star) \geq \eta(\omega), \forall t \geq \bar{t}_1(\omega, \delta), \forall i \in \mathcal{R}. \quad (17)$$

This completes Step 1. To proceed with Step 2 (i.e., upper-bounding the actual beliefs on each false hypothesis), given an $\epsilon > 0$, pick a small $\bar{\epsilon}(\omega) > 0$ such that $\bar{\epsilon}(\omega) < \min\{\eta(\omega), \epsilon\}$. Fix a hypothesis $\theta \neq \theta^\star$, let $q = n + 1$, and note that based on Lemma 1, for each $i \in \mathcal{S}(\theta^\star, \theta) \cap \mathcal{R}$, there exists $t_i^\theta(\omega, \bar{\epsilon}(\omega))$ such that for all $t \geq t_i^\theta(\omega, \bar{\epsilon}(\omega))$, $\pi_{i,t}(\theta) \leq \bar{\epsilon}^q(\omega)$. Define $\bar{t}_2 \triangleq \max\{\bar{t}_1(\omega, \delta), \max_{i \in \mathcal{S}(\theta^\star, \theta) \cap \mathcal{R}} \{t_i^\theta(\omega, \bar{\epsilon}(\omega))\}\}$. For any agent $i \in \mathcal{S}(\theta^\star, \theta) \cap \mathcal{R}$, observe that $\min\{\{\mu_{j,\bar{t}_2}(\theta^\star)\}_{j \in \mathcal{M}^{\theta^\star}_{i,\bar{t}_2}}, \pi_{i,\bar{t}_2+1}(\theta^\star)\} \geq \eta(\omega)$. Combining the above with arguments used to arrive at (8), we obtain:

$$\mu_{i,t}(\theta) < \bar{\epsilon}^{(q-1)}(\omega), \forall t \geq \bar{t}_2 + 1, \forall i \in \mathcal{S}(\theta^\star, \theta) \cap \mathcal{R}. \quad (18)$$

If $\mathcal{V} \setminus \mathcal{S}(\theta^\star, \theta)$ is empty, then we are done. Else, define

$$\mathcal{L}_1^{(\theta^\star, \theta)} \triangleq \{i \in \{\mathcal{V} \setminus \mathcal{S}(\theta^\star, \theta)\} : |\mathcal{N}_i \cap \mathcal{S}(\theta^\star, \theta)| \geq (2f+1)\}. \quad (19)$$

Whenever $\mathcal{V} \setminus \mathcal{S}(\theta^\star, \theta)$ is non-empty, we claim that $\mathcal{L}_1^{(\theta^\star, \theta)}$ (as defined above) is also non-empty based on condition (i) of the theorem. To see this, note that if $\mathcal{L}_1^{(\theta^\star, \theta)}$ is empty, then $\mathcal{C} = \mathcal{V} \setminus \mathcal{S}(\theta^\star, \theta)$ is not $(2f+1)$-reachable, violating the fact that $\mathcal{G}$ is strongly $(2f+1)$-robust w.r.t. $\mathcal{S}(\theta^\star, \theta)$. We claim

$$\min_{j \in \mathcal{M}^\theta_{i,\bar{t}_2+1}} \mu_{j,\bar{t}_2+1}(\theta) < \bar{\epsilon}^{(q-1)}(\omega), \forall i \in \mathcal{L}_1^{(\theta^\star, \theta)} \cap \mathcal{R}. \quad (20)$$

To verify the above claim, pick any agent $i \in \mathcal{L}_1^{(\theta^\star, \theta)} \cap \mathcal{R}$. When $|\mathcal{M}^\theta_{i,\bar{t}_2+1} \cap \{\mathcal{S}(\theta^\star, \theta) \cap \mathcal{R}\}| > 0$, the claim follows immediately based on (18). Consider the case when $|\mathcal{M}^\theta_{i,\bar{t}_2+1} \cap \{\mathcal{S}(\theta^\star, \theta) \cap \mathcal{R}\}| = 0$. Since $i \in \mathcal{L}_1^{(\theta^\star, \theta)}$, it has at least $(2f+1)$ neighbors in $\mathcal{S}(\theta^\star, \theta)$, out of which at least $f+1$ are regular based on the f-locality of the adversarial model. Since the set $\mathcal{J}^\theta_{i,\bar{t}_2+1}$ has cardinality f, it must then be that $|\mathcal{U}^\theta_{i,\bar{t}_2+1} \cap \{\mathcal{S}(\theta^\star, \theta) \cap \mathcal{R}\}| > 0$. Let $u \in \mathcal{U}^\theta_{i,\bar{t}_2+1} \cap \{\mathcal{S}(\theta^\star, \theta) \cap \mathcal{R}\}$. Based on the way $\mathcal{M}^\theta_{i,\bar{t}_2+1}$ is defined, it must be that $\mu_{j,\bar{t}_2+1}(\theta) \leq \mu_{u,\bar{t}_2+1}(\theta) < \bar{\epsilon}^{(q-1)}(\omega), \forall j \in \mathcal{M}^\theta_{i,\bar{t}_2+1}$, where the last inequality follows from (18). This establishes our claim regarding (20). Consider the update of $\mu_{i,\bar{t}_2+2}(\theta)$ based on (3). The above arguments (that apply to subsequent time-steps as well) imply that the numerator of the fraction on the RHS of (3) is upper-bounded by $\bar{\epsilon}^{(q-1)}(\omega)$, while the denominator is lower-bounded by $\eta(\omega)$, yielding:

$$\mu_{i,t}(\theta) < \bar{\epsilon}^{(q-2)}(\omega), \forall t \geq \bar{t}_2 + 2, \forall i \in \mathcal{L}_1^{(\theta^\star, \theta)} \cap \mathcal{R}. \quad (21)$$

With $\mathcal{L}_0^{(\theta^\star, \theta)} \triangleq \mathcal{S}(\theta^\star, \theta)$, we recursively define the sets $\mathcal{L}_r^{(\theta^\star, \theta)}, 1 \leq r \leq (n-1)$ as follows:

$$\mathcal{L}_r^{(\theta^\star, \theta)} \triangleq \{i \in \mathcal{V} \setminus \{\textstyle\bigcup_{c=0}^{r-1} \mathcal{L}_c^{(\theta^\star, \theta)}\} : |\mathcal{N}_i \cap \{\textstyle\bigcup_{c=0}^{r-1} \mathcal{L}_c^{(\theta^\star, \theta)}\}| \geq (2f+1)\}. \quad (22)$$

[2] In particular, this establishes that based on the LFRHE algorithm, an adversarial agent cannot cause its regular out-neighbors to set their actual beliefs on $\theta^\star$ to be 0 by setting its own actual belief on $\theta^\star$ to be 0.

The proof can then be completed as in Theorem 1 by inducting on r. □

VI. Conclusion

In this paper, we introduced a distributed learning rule that differs fundamentally from those existing in the literature, in the sense that it does not rely on any consensus-based belief aggregation protocol. Using a novel sample path based analysis technique, we established its consistency under minimal requirements on the information structures of the agents and the communication graph. We then showed that a significant benefit of the proposed learning rule is that it can be easily and efficiently modified to account for the presence of misbehaving agents in the network, modeled via the Byzantine adversary model. Ongoing work involves performing a detailed convergence rate analysis to see how such rates compare with those existing in literature.

References

[1] V. V. Veeravalli, T. Basar, and H. V. Poor, "Decentralized sequential detection with a fusion center performing the sequential test," *IEEE Transactions on Information Theory*, vol. 39, no. 2, pp. 433–442, 1993.

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

[3] A. Jadbabaie, P. Molavi, and A. Tahbaz-Salehi, "Information heterogeneity and the speed of learning in social networks," *Columbia Bus. Sch. Res. Paper*, pp. 13–28, 2013.

[4] Q. Liu, A. Fang, L. Wang, and X. Wang, "Social learning with time-varying weights," *Journal of Systems Science and Complexity*, vol. 27, no. 3, pp. 581–593, 2014.

[5] A. Nedić, A. Olshevsky, and C. A. Uribe, "Fast convergence rates for distributed Non-Bayesian learning," *IEEE Trans. on Autom. Control*, vol. 62, no. 11, pp. 5538–5553, 2017.

[6] A. Lalitha, T. Javidi, and A. Sarwate, "Social learning and distributed hypothesis testing," *IEEE Trans. on Information Theory*, vol. 64, no. 9, pp. 6161–6179, 2018.

[7] L. Su and N. H. Vaidya, "Defending Non-Bayesian learning against adversarial attacks," *Distributed Computing*, pp. 1–13, 2016.

[8] S. Shahrampour, A. Rakhlin, and A. Jadbabaie, "Distributed detection: Finite-time analysis and impact of network topology," *IEEE Trans. on Autom. Control*, vol. 61, no. 11, pp. 3256–3268, 2016.

[9] K. R. Rad and A. Tahbaz-Salehi, "Distributed parameter estimation in networks," in *Proceedings of the 49th IEEE Decision and Control Conference*, 2010, pp. 5050–5055.

[10] S. Shahrampour and A. Jadbabaie, "Exponentially fast parameter estimation in networks using distributed dual averaging," in *Proc. of the 52nd Decision and Control Conference*, 2013, pp. 6196–6201.

[11] G. L. Gilardoni and M. K. Clayton, "On reaching a consensus using DeGroot's iterative pooling," *The Annals of Stat.*, pp. 391–401, 1993.

[12] D. Dolev, N. A. Lynch, S. S. Pinter, E. W. Stark, and W. E. Weihl, "Reaching approximate agreement in the presence of faults," *Journal of the ACM (JACM)*, vol. 33, no. 3, pp. 499–516, 1986.

[13] N. H. Vaidya, L. Tseng, and G. Liang, "Iterative approximate Byzantine consensus in arbitrary directed graphs," in *Proc. of the ACM Symp. on Principles of Distributed Computing*, 2012, pp. 365–374.

[14] H. J. LeBlanc, H. Zhang, X. Koutsoukos, and S. Sundaram, "Resilient asymptotic consensus in robust networks," *IEEE Journal on Selected Areas in Communications*, vol. 31, no. 4, pp. 766–781, 2013.

[15] L. Su and N. H. Vaidya, "Fault-tolerant multi-agent optimization: optimal iterative distributed algorithms," in *Proc. of the 2016 ACM Symp. on Principles of Dist. Comp.* ACM, 2016, pp. 425–434.

[16] S. Sundaram and B. Gharesifard, "Distributed optimization under adversarial nodes," *IEEE Trans. on Autom. Control*, vol. 64, no. 3, pp. 1063–1076, 2019.

[17] A. Mitra and S. Sundaram, "Byzantine-resilient distributed observers for LTI systems," *arXiv preprint arXiv:1802.09651*, 2018.

[18] A. Mitra, J. A. Richards, and S. Sundaram, "A new approach for distributed hypothesis testing with extensions to Byzantine-resilience," *arXiv preprint arXiv:1903.05817*, 2019.