

SOCIAL LEARNING UNDER INFERENCE ATTACKS

Konstantinos Ntemos Virginia Bordignon Stefan Vlaski Ali H. Sayed

School of Engineering, Ecole Polytechnique Fédérale de Lausanne (EPFL)

ABSTRACT

A common assumption in the social learning literature is that agents exchange information in an unselfish manner. In this work, we consider the scenario where a subset of agents aims at driving the network beliefs to the wrong hypothesis. The adversaries are unaware of the true hypothesis. However, they will “blend in” by behaving similarly to the other agents and will manipulate the likelihood functions used in the belief update process to launch *inferential attacks*. We will characterize the conditions under which the network is misled. Then, we will explain that it is possible for such attacks to succeed by showing that strategies exist that can be adopted by the malicious agents for this purpose. We examine both situations in which the agents have minimal or no information about the network model.

Index Terms— social learning, malicious agents, information diffusion, inferential attacks.

1. INTRODUCTION AND RELATED WORK

In the *social learning* paradigm [1–9] agents aim at *learning* an underlying system state through their own observations as well as information provided by other agents. The communication among the agents is dictated by an underlying graph topology where each agent directly communicates only with its neighbors. The standard assumption in social learning literature is that agents exchange honestly their beliefs in a cooperative fashion. However, in many cases agents may exhibit intentional *misbehavior* or operate in a *faulty* manner. Many works studied the robustness of distributed processing systems against *Byzantine attacks* [10, 11], where adversaries can deviate from the system protocol in an *arbitrary* way. Another related line of research includes *robust detection* [12–15], where adversaries aim at driving a fusion center to the wrong decision.

More recently, non-Bayesian social learning in the presence of malicious/faulty agents was studied in [16–19]. This problem poses new challenges, as each agent tries to learn the system state in a decentralized fashion. The authors in [17, 18] studied social learning in the presence of Byzantine attackers and showed that under certain assumptions on the structure of the communication graph, the normal agents successfully identify the underlying state. Other works studied more specific attack scenarios. One such scenario, which is the subject of interest in this paper, is driving the network beliefs to the wrong state. This cannot be guaranteed by letting adversaries send arbitrary information, since it can result in undesired non-convergent behavior. One pattern of adversarial behavior that may mislead the network is the use of *corrupted* likelihood functions when adversaries update their beliefs.

This kind of adversarial behavior was considered in [16, 19]. In [16] the authors devised a detection scheme where each agent

reuses some of its most recent observations to update its beliefs along with the use of a convex combination of the beliefs of multiple social networks. The case where the agents’ models are unknown and malicious agents are also present in the network was considered in [19]. The authors discuss different attack scenarios and discriminate between the so-called *weak* and *strong* malicious agents. Weak malicious agents manipulate only the likelihood functions in the belief update rule, while strong malicious agents can additionally filter out information sent from honest agents. The focus of these works is on detection of adversaries. Instead, our work is focused on investigating adversarial strategies, which is not addressed by [16, 19].

Our work addresses the case where the adversaries do not have knowledge about the true state and aim at forcing the network beliefs to the wrong hypothesis. We assume that adversaries participate in the information diffusion process as dictated by the social learning protocol, but disseminate falsified beliefs, which are produced by the use of corrupted likelihood functions. The fact that adversaries incorporate information from the network allows them to “blend in” by appearing to follow the social learning protocol. We refer to this type of attacks as “inferential attacks” due to the fact that adversaries have no knowledge of the true state and try to drive the network beliefs to the wrong state by manipulating their inference model (i.e. likelihood functions). We contribute to the related literature by investigating adversarial strategies.

More specifically, in this work, we answer the following questions. First, under what conditions is an unguarded network (no detection mechanism exists) misled under inferential attacks? Second, if adversaries do not know the true state, is there a way to construct fake likelihood functions that drive the normal agents’ beliefs to the wrong state? Finally, in scenarios of incomplete information, when adversaries do not have any knowledge about the network, how should they manipulate the observation models?

In the aforementioned setup, we characterize the conditions under which the network is misled. We prove that it depends on the agents’ observation models, malicious and benign agents’ centrality, and attack strategies. In this way, we reveal an interplay between the network topology, which captures the diffusion of information, and injection of mis-information in the social learning paradigm. Then, we prove that if an adversary knows certain network characteristics, then there is always an attack strategy that misleads the network. Finally, we study the scenario when adversaries have no knowledge about the network properties, propose an attack strategy, and illustrate its impact on the learning performance.

2. SYSTEM MODEL

We assume a set $\mathcal{N} = \mathcal{N}^n \cup \mathcal{N}^m$ of agents, where \mathcal{N}^n and \mathcal{N}^m denote the sets of normal and malicious agents, respectively. The types of the agents (i.e., normal or malicious) are unknown. The network is represented by an undirected graph $\mathcal{G} = \langle \mathcal{N}, \mathcal{E} \rangle$, where \mathcal{E} includes bidirectional links between agents. The set of neighbors

This work was supported in part by the Swiss National Science Foundation grant 205121-184999. E-mails: konstantinos.ntemos@epfl.ch, virginia.bordignon@epfl.ch, stefan.vlaski@epfl.ch, ali.sayed@epfl.ch.

of an agent $k \in \mathcal{N}$ is denoted by \mathcal{N}_k (including agent k).

We consider an adversarial setting where the normal agents aim at learning the true state $\theta^* \in \Theta = \{\theta_1, \theta_2\}$, while malicious agents try to impede the normal agents by forcing their beliefs towards the wrong state. All agents are unaware of the true state θ^* .

We assume that each agent k has access to observations $\zeta_{k,i} \in \mathcal{Z}_k$ at every time $i \geq 1$. Agent k also has access to the likelihood functions $L_k(\zeta_{k,i}|\theta)$, $\theta \in \Theta$. The signals $\zeta_{k,i}$ are independent and identically distributed (i.i.d.) over time. The sets \mathcal{Z}_k are assumed to be finite with $|\mathcal{Z}_k| \geq 2$ for all $k \in \mathcal{N}$. We will use the notation $L_k(\theta)$ instead of $L_k(\zeta_{k,i}|\theta)$ whenever it is clear from the context.

Assumption 1. (Finiteness of KL divergences). For any agent $k \in \mathcal{N}$ and for any $\theta \neq \theta^*$, $D_{KL}(L_k(\theta^*)||L_k(\theta))$ is finite. ■

At each time i , agent k keeps a *belief vector* $\mu_{k,i}$, which is a probability distribution over the possible system states. The belief components $\mu_{k,i}(\theta)$, $\theta \in \Theta$, quantify the confidence of agent k that each θ is the true state. We assume that all agents, both normal and malicious, are unaware of the true state. Thus, we impose the following assumption on initial beliefs.

Assumption 2. (Positive initial beliefs). $\mu_{k,0}(\theta) > 0, \forall \theta \in \Theta, k \in \mathcal{N}$. ■

3. SOCIAL LEARNING WITH ADVERSARIES

Each normal agent k uses the acquired observations $\zeta_{k,i}$, along with the likelihood function $L_k(\zeta_{k,i}|\theta)$, to update their belief vector using Bayes' rule. Agents communicate with each other and exchange information. We consider the log-linear social learning rule [7,9,20]. The normal agents update their beliefs in the following manner:

$$\psi_{k,i}(\theta) = \frac{L_k(\zeta_{k,i}|\theta)\mu_{k,i-1}(\theta)}{\sum_{\theta'} L_k(\zeta_{k,i}|\theta')\mu_{k,i-1}(\theta')}, \quad k \in \mathcal{N}^n \quad (1)$$

$$\mu_{k,i}(\theta) = \frac{\prod_{\ell \in \mathcal{N}_k} \psi_{\ell,i}^{a_{\ell k}}(\theta)}{\sum_{\theta'} \prod_{\ell \in \mathcal{N}_k} \psi_{\ell,i}^{a_{\ell k}}(\theta')}, \quad k \in \mathcal{N}^n \quad (2)$$

where $a_{\ell k}$ denotes the *combination weight* assigned by agent k to neighboring agent ℓ , satisfying $1 \geq a_{\ell k} > 0$, for all $\ell \in \mathcal{N}_k$, $a_{\ell k} = 0$ for all $\ell \notin \mathcal{N}_k$ and $\sum_{\ell \in \mathcal{N}_k} a_{\ell k} = 1$. Let A denote the *combination matrix* which consists of all agents' combination weights. We impose the following assumption on the network topology.

Assumption 3. (Strongly-connected network). The communication graph is strongly connected (i.e., there always exists a path with positive weights linking any two agents and at least one agent has a self-loop, meaning that there is at least one agent $k \in \mathcal{N}$ with $a_{kk} > 0$). ■

For a strongly connected network, the limiting behavior of A^T is given by $\lim_{i \rightarrow \infty} (A^T)^i = \mathbb{1}u^T$, where u is the Perron eigenvector [21]. The eigenvector u is associated with the eigenvalue at 1, all its entries are positive and are normalized to add up to 1. Moreover, its k -th entry u_k expresses a measure of influence of agent k on the network and it is also called the *centrality* of agent k .

We consider the scenario where adversaries aim at misleading the network to accept the wrong hypothesis by modifying the way they use their observations. More specifically, we assume that malicious agents deviate in step (1) by using a fake likelihood function, denoted by $\widehat{L}_k(\cdot)$ instead of $L_k(\cdot)$ to update their beliefs, while they

follow (2) without deviation. Inferential attacks are therefore modeled by assuming that adversaries follow the following update rule:

$$\psi_{k,i}(\theta) = \frac{\widehat{L}_k(\zeta_{k,i}|\theta)\mu_{k,i-1}(\theta)}{\sum_{\theta'} \widehat{L}_k(\zeta_{k,i}|\theta')\mu_{k,i-1}(\theta')}, \quad k \in \mathcal{N}^m. \quad (3)$$

We also impose the following technical assumption on the distorted likelihood functions.

Assumption 4. (Distorted likelihood functions with full support). For every agent $k \in \mathcal{N}^m$, the distorted likelihood function satisfies $\epsilon \leq \widehat{L}_k(\zeta_{k,i}|\theta)$ for all $\zeta_{k,i} \in \mathcal{Z}_k$, $\theta \in \Theta$, where $0 < \epsilon \ll 1$ is a small positive real constant that satisfies $\epsilon < \min_k \frac{1}{|\mathcal{Z}_k|}$. ■

We say that an agent k 's belief converges *almost surely* (a.s.) to the true state if $\mu_{k,i}(\theta^*) \rightarrow 1$ as $i \rightarrow \infty$ with probability 1. Conversely, agent k 's belief converges a.s. to the wrong state if $\mu_{k,i}(\theta^*) \rightarrow 0$ as $i \rightarrow \infty$ with probability 1. The following result characterizes the asymptotic learning behavior of the network; the proof, as well as the proofs for subsequent results, are omitted due to space limitations.

Theorem 1. (Belief convergence with adversaries). Under Assumptions 1, 2, 3, 4 the following are true:

1. The agents' beliefs converge a.s. to the wrong state if

$$\sum_{k \in \mathcal{N}^n} u_k \mathbb{E} \left\{ \log \frac{L_k(\zeta_k|\theta^*)}{L_k(\zeta_k|\theta)} \right\} < \sum_{k \in \mathcal{N}^m} u_k \mathbb{E} \left\{ \log \frac{\widehat{L}_k(\zeta_k|\theta)}{\widehat{L}_k(\zeta_k|\theta^*)} \right\}. \quad (4)$$

2. The agents' beliefs converge a.s. to the true state if

$$\sum_{k \in \mathcal{N}^n} u_k \mathbb{E} \left\{ \log \frac{L_k(\zeta_k|\theta^*)}{L_k(\zeta_k|\theta)} \right\} > \sum_{k \in \mathcal{N}^m} u_k \mathbb{E} \left\{ \log \frac{\widehat{L}_k(\zeta_k|\theta)}{\widehat{L}_k(\zeta_k|\theta^*)} \right\} \quad (5)$$

where $\theta^*, \theta \in \Theta$, $\theta^* \neq \theta$. ■

The Theorem characterizes under what condition the agents in the graph can be misled, namely, when condition (4) holds. Thus, malicious agents would strive to construct their distorted likelihood functions to satisfy (4). The expectation in (4) and (5) is taken w.r.t. the true likelihood distributions, $L_k(\zeta_k|\theta^*)$. Since $\zeta_{k,i}$ are i.i.d. over time, we omit the time index i . The threshold rule (4)-(5) fully characterizes the convergence of network beliefs. Note that whether or not the agents' beliefs will converge to the true system state depends on agents' observation models (informativeness of the signals), on the distorted likelihood functions, and on network topology (agents' centrality).

3.1. Attack strategies with known network divergence

In this section, we answer the question of whether an adversary k can construct $\widehat{L}_k(\theta_1), \widehat{L}_k(\theta_2)$ in such a way that the network will always be driven to the wrong hypothesis. Note that the state is unknown to adversaries. Thus, an adversary should select $\widehat{L}_k(\theta_1), \widehat{L}_k(\theta_2)$ such that (4) is satisfied for both $\theta^* = \theta_1$ and $\theta^* = \theta_2$ to ensure that the network will converge to the wrong hypothesis no matter what the true state is. Let S_j denote the term on the left-hand side of (4) for $\theta^* = \theta_j, j = 1, 2$. Then, inequality (4) is equivalent to the following two conditions:

$$S_j < u_k \sum_{\zeta_k} L_k(\zeta_k|\theta_j) \log \frac{\widehat{L}_k(\zeta_k|\theta_j)}{\widehat{L}_k(\zeta_k|\theta_j)}, \quad k \in \mathcal{N}^m \quad (6)$$

where $j, j' \in \{1, 2\}, j \neq j'$. We call S_j *network divergence*, or simply *divergence*, of the normal subnetwork for $\theta^* = \theta_j$.

Identifying a set of probability mass functions (PMFs) $\widehat{L}_k(\theta_1), \widehat{L}_k(\theta_2)$ that mislead the network for both $\theta^* = \theta_1$ and $\theta^* = \theta_2$ requires solving the system of inequalities (6) w.r.t. $\widehat{L}_k(\zeta_k|\theta_1), \widehat{L}_k(\zeta_k|\theta_2), \zeta_k \in \mathcal{Z}_k$. In the following, we present one construction that captures such a family of PMFs. Before we present the result, we note that PMFs are *uninformative* if $L_k(\zeta_k|\theta_1) = L_k(\zeta_k|\theta_2)$ for all $\zeta_k \in \mathcal{Z}_k$, otherwise the PMFs are *informative*. Further, let us introduce the following quantities:

$$n_j \triangleq L_\ell(\zeta_\ell^j|\theta_2)S_1 + L_\ell(\zeta_\ell^j|\theta_1)S_2, \quad j = 1, 2 \quad (7)$$

$$d_\ell \triangleq L_\ell(\zeta_\ell^2|\theta_2)L_\ell(\zeta_\ell^1|\theta_1) - L_\ell(\zeta_\ell^2|\theta_1)L_\ell(\zeta_\ell^1|\theta_2) \quad (8)$$

Theorem 2. (Distorted PMFs with known divergences). *The following construction drives the network to the wrong hypothesis for any $\theta^* \in \Theta$, given that there exists at least one adversary with informative PMFs, for sufficiently small ϵ . Every adversary $\ell \in \mathcal{N}^m$ with informative PMFs uses the following construction.*

$$\widehat{L}_\ell(\zeta_\ell|\theta_j) = \begin{cases} \epsilon_{j'}, & \text{if } \zeta_\ell = \zeta_\ell^{j'} \\ \alpha - \epsilon_{j'}, & \text{if } \zeta_\ell = \zeta_\ell^j \\ \epsilon, & \text{otherwise} \end{cases} \quad (9)$$

where $\alpha = 1 - (|\mathcal{Z}_\ell| - 2)\epsilon$, $j, j' \in \{1, 2\}, j \neq j'$, $\zeta_\ell^1, \zeta_\ell^2 \in \mathcal{Z}_\ell$ are such that $L_\ell(\zeta_\ell^1|\theta_1)L_\ell(\zeta_\ell^2|\theta_2) \neq L_\ell(\zeta_\ell^1|\theta_2)L_\ell(\zeta_\ell^2|\theta_1)$ and

$$\epsilon_1 = \frac{e^{x_1}\alpha(e^{x_2} - 1)}{e^{x_2} - e^{x_1}} \quad (10)$$

$$\epsilon_2 = \frac{\alpha(1 - e^{x_1})}{e^{x_2} - e^{x_1}} \quad (11)$$

where $x^+ > x_1 > x'_{\ell,1}$ if $d_\ell > 0$ and $x^- < x_1 < x'_{\ell,1}$ if $d_\ell < 0$, $x^- = \log(\epsilon/(\alpha - \epsilon))$, $x^+ = -x^-$, $x'_{\ell,1} = n_2/(u_\ell d_\ell)$, $x'_{\ell,2} = n_1/(u_\ell d_\ell)$ and

$$x_2 = \beta_\ell(x_1 - x'_{\ell,1}) + x'_{\ell,2}, \quad |x_2| < x^+ \quad (12)$$

with β_ℓ such that

$$\min_j \left\{ -\frac{L_\ell(\zeta_\ell^1|\theta_j)}{L_\ell(\zeta_\ell^2|\theta_j)} \right\} < \beta_\ell < \max_j \left\{ -\frac{L_\ell(\zeta_\ell^1|\theta_j)}{L_\ell(\zeta_\ell^2|\theta_j)} \right\} \quad (13)$$

where $j = 1, 2$. If an adversary $\ell \in \mathcal{N}^m$ has uninformative PMFs, then it sets $\widehat{L}_\ell(\theta_1) = \widehat{L}_\ell(\theta_2) = L_\ell(\theta_1) = L_\ell(\theta_2)$. ■

The above result states that even one adversary with informative likelihood functions can always construct fake PMFs that mislead the network, given that it has access to divergences S_1, S_2 of the normal subnetwork.

3.2. Attack strategies with unknown network divergence

In general, it is not realistic to assume knowledge of network divergences S_1, S_2 is available to the adversaries. Thus, in this section, we investigate what is the best that adversaries can do when they do not know the characteristics of the normal subnetwork. Rearranging (4), we can define the following cost function.

$$\begin{aligned} \mathcal{C}(\theta^*) &= \sum_{k \in \mathcal{N}^n} u_k D_{KL}(L_k(\theta^*) || L_k(\theta)) \\ &+ \sum_{\ell \in \mathcal{N}^m} u_\ell \sum_{\zeta_\ell} L_\ell(\zeta_\ell|\theta^*) \log \frac{\widehat{L}_\ell(\zeta_\ell|\theta^*)}{\widehat{L}_\ell(\zeta_\ell|\theta)} \end{aligned} \quad (14)$$

where $\theta^*, \theta \in \Theta, \theta_1 \neq \theta_2$. We observe that the second term in (14) is under malicious agents' control. Thus, one option for the adversaries is to minimize (14) over $\widehat{L}_\ell(\theta_1), \widehat{L}_\ell(\theta_2)$. However, θ^* is unknown as well. A viable alternative is to treat the true state θ^* as a random variable, assuming some prior distribution over the states $\pi = (\pi_{\theta_1}, \pi_{\theta_2})$. We assume that malicious agents share a common prior. Thus, taking expectation over the true state θ^* in (14) leads to the following minimization problem for the malicious agents:

$$\begin{aligned} \min_{\widehat{L}_\ell(\theta_1), \widehat{L}_\ell(\theta_2)} \sum_{\theta \in \Theta} \pi_\theta C(\theta^* = \theta), \quad \ell \in \mathcal{N}^m \quad (15) \\ \text{s.t. } \widehat{L}_\ell(\zeta_\ell|\theta) \geq \epsilon, \quad \forall \zeta_\ell \in \mathcal{Z}_\ell, \theta \in \Theta, \\ \sum_{\zeta_\ell \in \mathcal{Z}_\ell} \widehat{L}_\ell(\zeta_\ell|\theta) = 1, \quad \forall \theta \in \Theta \end{aligned}$$

It should be noted that the solution to the minimization problem above, denoted by $\widehat{L}_\ell^*(\theta_1), \widehat{L}_\ell^*(\theta_2)$, minimizes the weighted average of (14) for $\theta^* = \theta_1$ and $\theta^* = \theta_2$. This means that (4) is not necessarily satisfied if adversaries utilize $\widehat{L}_\ell^*(\theta_1), \widehat{L}_\ell^*(\theta_2)$. We also highlight that the solution depends on the prior distribution of the true state π . The optimization problem is decomposable across agents $\ell \in \mathcal{N}^m$ and thus (15) reduces to the following for each agent $\ell \in \mathcal{N}^m$:

$$\min_{\widehat{L}_\ell(\theta_1)} \sum_{\zeta_\ell} Z_\ell(\zeta_\ell) \log \widehat{L}_\ell(\zeta_\ell|\theta_1) - \max_{\widehat{L}_\ell(\theta_2)} \sum_{\zeta_\ell} Z_\ell(\zeta_\ell) \log \widehat{L}_\ell(\zeta_\ell|\theta_2) \quad (16)$$

$$\text{s.t. } \widehat{L}_\ell(\zeta_\ell|\theta_1) \geq \epsilon, \quad \widehat{L}_\ell(\zeta_\ell|\theta_2) \geq \epsilon$$

$$\sum_{\zeta_\ell \in \mathcal{Z}_\ell} \widehat{L}_\ell(\zeta_\ell|\theta_1) = 1, \quad \sum_{\zeta_\ell \in \mathcal{Z}_\ell} \widehat{L}_\ell(\zeta_\ell|\theta_2) = 1$$

where

$$Z_\ell(\zeta_\ell) \triangleq \pi_{\theta_1} L_\ell(\zeta_\ell|\theta_1) - \pi_{\theta_2} L_\ell(\zeta_\ell|\theta_2), \quad \zeta_\ell \in \mathcal{Z}_\ell. \quad (17)$$

Note that each coefficient $Z_\ell(\zeta_\ell)$ expresses the *relative confidence* that ζ_ℓ resulted from state θ_1 instead of θ_2 . We define the set

$$\mathcal{D}_\ell^1 \triangleq \{\zeta_\ell : Z(\zeta_\ell) \geq 0, \quad \ell \in \mathcal{N}^m\} \quad (18)$$

which is comprised of all observations $\zeta_\ell \in \mathcal{Z}_\ell$ for which the level of confidence that these observations were generated by state θ_1 is greater than the confidence that they are generated by θ_2 . The set $\mathcal{D}_\ell^2 = \mathcal{Z}_\ell \setminus \mathcal{D}_\ell^1, \ell \in \mathcal{N}^m$ is comprised of the observations for which the confidence that they are generated from θ_2 is greater compared to θ_1 . The solution depends on whether these sets are both non-empty or not. Next, we examine the two scenarios.

3.2.1. Mixed Confidence

We study first the scenario where both sets $\mathcal{D}_\ell^1, \mathcal{D}_\ell^2$ are non-empty, which means that some observations are more likely to have been generated by θ_1 , while some others by θ_2 . In this case, the solution to (16) is given by the following result.

Theorem 3. (Distorted PMFs with unknown divergences and mixed confidence). *If both $\mathcal{D}_\ell^1, \mathcal{D}_\ell^2$ are non-empty sets, then the attack strategy optimizing (16) for an agent $\ell \in \mathcal{N}^m$ is given by*

$$\widehat{L}_\ell(\zeta_\ell|\theta_j) = \begin{cases} \epsilon, & \text{if } \zeta_\ell \in \mathcal{D}_\ell^j, \\ \frac{Z_\ell(\zeta_\ell)(1 - |\mathcal{D}_\ell^j|\epsilon)}{\sum_{\zeta_\ell \notin \mathcal{D}_\ell^j} Z_\ell(\zeta_\ell)}, & \text{if } \zeta_\ell \notin \mathcal{D}_\ell^j \end{cases} \quad (19)$$

where $j \in \{1, 2\}$. ■

The intuition behind the attack strategy is the following. We focus on the construction for $\widehat{L}_\ell(\theta_1)$ and the rationale is the same for $\widehat{L}_\ell(\theta_2)$. The constructed PMF is such that the least possible probability mass (i.e., ϵ) is assigned to every observation ζ_ℓ that is more likely to be generated from state θ_1 (i.e., for all $\zeta_\ell \in \mathcal{D}_\ell^1$). For the remaining observations that are more likely to be generated from θ_2 (i.e., $\zeta_\ell \in \mathcal{D}_\ell^2$) the probability mass placed on every $\zeta_\ell \in \mathcal{D}_\ell^2$ is in proportion to the difference in confidence that ζ_ℓ is generated from θ_2 instead of θ_1 . The more likely it is for ζ_ℓ to be generated from θ_2 , the more probability mass is placed on $\widehat{L}(\zeta_\ell|\theta_1)$.

3.2.2. Pure Confidence

In the scenario where all observations are more likely to be generated from either θ_1 or θ_2 , the solution is different, but the intuition remains similar. The solution to (16) is the following.

Theorem 4. (Distorted PMFs with unknown divergences and pure confidence). Let $\mathcal{D}_\ell^1 = \emptyset$ or $\mathcal{D}_\ell^2 = \emptyset$. Then, the attack strategy optimizing (16) for an agent $\ell \in \mathcal{N}^m$ is given by

$$\widehat{L}_\ell(\zeta_\ell|\theta_j) = \begin{cases} 1 - (|\mathcal{Z}_\ell| - 1)\epsilon, & \text{if } \mathcal{D}_\ell^j = \mathcal{Z}_\ell, \zeta_\ell = \zeta_{\min} \\ \epsilon, & \text{if } \mathcal{D}_\ell^j = \mathcal{Z}_\ell \text{ and } \zeta_\ell \neq \zeta_{\min} \\ \frac{Z_\ell(\zeta_\ell)}{\sum_{\zeta_\ell \in \mathcal{Z}_\ell} Z_\ell(\zeta_\ell)}, & \text{if } \mathcal{D}_\ell^j = \emptyset \end{cases} \quad (20)$$

where $j \in \{1, 2\}$ and $\zeta_{\min} = \arg \min_{\zeta_\ell} \{Z_\ell(\zeta_\ell)\}$. ■

The intuition behind the result is the following. If $\mathcal{D}_\ell^2 = \emptyset$, then it is more likely that all the observations are generated by state θ_1 . Thus, the PMF $\widehat{L}_\ell(\zeta_\ell|\theta_1)$ is generated according to the following rationale. The maximum possible probability mass is placed on the observation that is the least likely to be generated from state θ_1 , while the minimum possible probability mass ϵ is placed on the rest of the observations. Regarding the PMF $\widehat{L}_\ell(\zeta_\ell|\theta_2)$, the probability mass of every ζ_ℓ , is in proportion to the difference in confidence that ζ_ℓ is generated from θ_1 instead of θ_2 . The more likely ζ_ℓ is to be generated from θ_1 instead of θ_2 , the more probability mass is placed on $\widehat{L}_\ell(\zeta_\ell|\theta_2)$. The rationale is the same for the case $\mathcal{D}_\ell^1 = \emptyset$.

4. SIMULATIONS

We assume 15 agents, with 11 normal and 4 malicious agents, interacting over a strongly-connected network. Each agent assigns uniform combination weights to its neighbors. The agents observe the state through a binary symmetric channel (i.e., $\mathcal{Z}_k = \{\zeta_1, \zeta_2\}$ for all $k \in \mathcal{N}$) with observation probabilities $L_k(\zeta_1|\theta_1) = L_k(\zeta_2|\theta_2) = p$ and $L_k(\zeta_2|\theta_1) = L_k(\zeta_1|\theta_2) = 1 - p$. We set $\epsilon = 10^{-3}$. In our experiments we want to highlight the impact of adversarial strategies on the learning process in conjunction with the dependence on the informativeness of agents' observation models and network topology. In Fig. 1 the binary symmetric channel is parametrized with $p = 0.8$, while in Fig. 2 agents' observation models are more discriminating between the two states with $p = 0.95$. Apart from the dependence on the observation models, we demonstrate the impact of network topology by considering a random topology (left sub-figures) and star topology (right sub-figures). In the star topology, the central agent is malicious. In each case, we consider two different attack strategies, namely the attack strategy with unknown divergences presented in Theorem 3 for prior distribution $(\pi_{\theta_1}, \pi_{\theta_2}) =$




Fig. 1: Evolution of agents' average belief (i.e., $\bar{\mu}_i(\theta^*) \triangleq \frac{\sum_{k \in \mathcal{N}} \mu_{k,i}(\theta^*)}{|\mathcal{N}|}$) for $p = 0.8$. Left: random topology, Right: star topology. ASUD: Attack Strategy with Unknown Divergences, RAS: Random Attack strategy.




Fig. 2: Evolution of agents' average belief with highly discriminating models ($p = 0.95$). Left: random topology, Right: star topology.

(0.5, 0.5) and a random attack strategy, where the distortion functions $\widehat{L}(\theta_1), \widehat{L}(\theta_2)$ are chosen randomly by the malicious agents. As we observe in Fig. 1, the network is misled under the proposed attack strategy for both network topologies in both cases when the system state is $\theta^* = \theta_1$ and $\theta^* = \theta_2$. The impact of random attack strategy is not sufficient to mislead the network.

The same rationale is followed in the experiments conducted for more discriminating models ($p = 0.95$). As we observe in Fig. 2, the impact of malicious behavior is smaller in this setup, since normal agents are more capable to discriminate between the two hypotheses. More specifically, in the left sub-figure of Fig. 2, the network converges to the true state regardless of the attack type for the random network topology. On the other hand, for the star topology, where the central agent is malicious, the network is misled under the proposed attack strategy, as presented in the right sub-figure. This is because the overall centrality of the malicious agents is bigger in star topology compared to the random network topology.

5. CONCLUSIONS

In this paper, the impact of inferential attacks on social learning was analyzed. We characterized the evolution of agents' beliefs and the adversaries' attack strategies were investigated. Our results are expected to shed light on the study of more elaborate attack schemes as well as on the development of light-weight detection mechanisms based on agents' characteristics (i.e., network centrality and observation models) and provide useful insight to situations where networks compete with each other in a strategic fashion.

6. REFERENCES

- [1] M. H. DeGroot, "Reaching a Consensus," *Journal of American Statistical Association*, vol. 69, no. 345, pp. 118–121, 1974.
- [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] X. Zhao, and A. H. Sayed, "Learning over social networks via diffusion adaptation," in *Proc. Asilomar Conference on Signals, Systems and Computers*, 2012, pp. 709-713.
- [4] H. Salami, B. Ying, and A. H. Sayed, "Social learning over weakly-connected graphs," *IEEE Trans. Signal and Information Processing over Networks*, vol. 3, no. 2, pp. 222-238, June 2017.
- [5] A. Nedić, A. Olshevsky, and C. A. Uribe, "Fast convergence rates for distributed non-Bayesian learning," *IEEE Trans. on Automatic Control*, vol. 62 no. 11, pp. 5538-5553, 2017.
- [6] P. Molavi, A. Tahbaz-Salehi, and A. Jadbabaie, "A theory of non-Bayesian social learning," *Econometrica*, vol. 86, no.2, pp. 445-490, 2018.
- [7] A. Lalitha, A. Sarwate, and T. Javidi, "Social learning and distributed hypothesis testing," in *Proc. International Symposium on Information Theory*, June 2014, pp. 551-555.
- [8] A. Lalitha, T. Javidi, A. D. and Sarwate, "Social learning and distributed hypothesis testing," *IEEE Transactions on Information Theory*, vol. 64, no. 9, pp.6161-6179, 2018.
- [9] V. Bordinon, V. Matta, and A. H. Sayed, "Social learning with partial information sharing," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Barcelona, Spain, May 2020, pp. 5540-5544.
- [10] L. Lamport, R. Shostak, and M. Pease, "The Byzantine generals problem," *ACM Transactions on Programming Languages and Systems*, vol. 4, no. 3, pp. 382–401, 1982.
- [11] D. Dolev, M. J. Fischer, R. Fowler, N. Lynch, and H. R. Strong, "An efficient algorithm for Byzantine agreement without authentication", *Information and Control*, vol. 52, no. 3, pp.257-274, 1983.
- [12] P. J. Huber, "A robust version of the probability ratio test," *Annals of Mathematical Statistics*, vol. 36, no. 6, pp. 1753–1758, 1965.
- [13] S. Marano, V. Matta, and L. Tong, "Distributed detection in the presence of Byzantine attacks," *IEEE Trans. on Signal Processing*, vol. 57 no. 1, pp. 16-29, 2008.
- [14] B. Kailkhura, Y.S. Han, S. Brahma, and P. K. Varshney, "On covert data falsification attacks on distributed detection systems," in *Proc. 13th International Symposium on Communications and Information Technologies (ISCIT)*, September 2013, pp. 412-417.
- [15] A. Vempaty, L. Tong, P. K. and Varshney, "Distributed inference with Byzantine data: State-of-the-art review on data falsification attacks," *IEEE Signal Processing Magazine*, vol. 30, no. 5, pp.65-75, 2013.
- [16] M. Bhotto and W. P. Tay, "Non-Bayesian social learning with observation reuse and soft switching," *ACM Transactions on Sensor Networks*, vol. 14, no. 2, pp. 1-21, 2018.
- [17] L. Su and N. H. Vaidya, "Defending non-Bayesian learning against adversarial attacks," *Distributed Computing*, pp. 1–13, 2018.
- [18] P. Vyavahare, L. Su, and N. H. Vaidya, "Distributed learning with adversarial agents under relaxed network condition," *arXiv preprint arXiv:1901.01943*, 2019.
- [19] J. Z. Hare, C. A. Uribe, L. M. Kaplan, and A. Jadbabaie, "On malicious agents in non-Bayesian social learning with uncertain models," in *Proc. 22th International Conference on Information Fusion (FUSION)*, July 2019, pp. 1-8.
- [20] V. Matta, V. Bordinon, A. Santos, and A. H. Sayed, "Interplay between topology and social learning over weak graphs," to appear in *IEEE Open Journal of Signal Processing*, vol. 1, pp. 99-119, 2020.
- [21] A. H. Sayed, "Adaptation, learning, and optimization over networks," *Foundations and Trends in Machine Learning*, vol. 7, issue 4-5, pp. 311-801, NOW Publishers, Boston-Delft, 2014.