Non-Bayesian Social Learning with Multiview Observations

Dongyan Sui, Weichen Cao, Stefan Vlaski, Chun Guan, Siyang Leng

Abstract— Non-Bayesian social learning enables multiple a-gents to conduct networked signal and information processing through observing environmental signals and information aggregating. Traditional non-Bayesian social learning models only consider single signals, limiting their applications in scenarios where multiple viewpoints of information are available. In this work, we exploit, in the information aggregation step, the independently learned results from observations taken from multiple viewpoints and propose a novel non-Bayesian social learning model for scenarios with multiview observations. We prove the convergence of the model under traditional assumptions and provide convergence conditions for the algorithm in the presence of misleading signals. Through theoretical analyses and numerical experiments, we validate the strong reliability and robustness of the proposed algorithm, showcasing its potential for real-world applications.

I. INTRODUCTION

Networked signal and information processing [1]–[5] refers to the collaborative processing of information and signals among a network of distributed agents. This approach leverages the collective capabilities of interconnected devices to perform tasks like decision-making, inference, and learning more efficiently than isolated systems. The significance of networked information processing lies in its ability to enhance performance through cooperation, offering advantages such as improved scalability, resilience, and resource efficiency. It is particularly relevant in applications like sensor networks, distributed control systems, and collaborative robotics.

Non-Bayesian social learning [6]–[20] offers a novel framework for networked signal and information processing, enabling a distributed way for agents with limited rationality and diverse sensing capabilities to infer collectively over a network. Agents process streams of incomplete data based on the underlying true state of the world, using network communications to form beliefs about various possible hypotheses and make an estimate of the underlying true state. This collaborative mechanism, which integrates neighbors' insights with fresh individual data, fosters a scalable method of learning without prior knowledge of network structure or historical data.

Various social learning models have been proposed, including aggregation methods such as linear averages [6], [7], geometric averages [8], [9], and the minimum operator [10]. These rules can be applied to different network structures, including undirected/directed, time-varying [11], [12], weakly-connected graphs [13], and higher-order topology [21], as well as to agents with growing self-confidence [14] and heterogeneous stubbornness parameters [15], disparate hypothesis [16], under inferential attacks [17], and in adversarial conditions [18]. Research has also explored learning under uncertain likelihood models and performance against malicious agents [19], [20]. All of these models offer theoretical assurances that, over time, agents can collectively learn the underlying true state of the world.

In previous non-Bayesian social learning models, the agents receive signals from the environment and cooperatively infer the underlying state based on their *a priori* knowledge of the signals. However, in practice, the group could perceive different features of the environment from various perspectives. For example, individuals could judge the species of trees based on the characteristics of both leaves and trunks; customers always infer the quality of a target product by observing the quality of other products from the same brand. Traditional methods cope with such situations by integrating these multiple viewpoints into one single signal, making it challenging to determine the likelihood functions of agents with the integrated signal due to the requirement of substantial data to assess the independence or correlation among multiple signals.

In this work, we propose a novel non-Bayesian social learning algorithm based on multiview observations. Our algorithm allows the group to learn independently from multiple signals and achieves interaction among multiview observations during the information aggregation process. Similarly to previous methods, we prove the correct convergence of our proposed algorithm under traditional assumptions. Additionally, we provide convergence conditions based on the presence of misleading signals. Numerical experiments validate the effectiveness of our theoretical analysis, and we showcase the robust fault-tolerance capability of our proposed algorithm in the task of multi-agent collaborative localization.

The remaining part of this paper is organized as follows: Section II provides a full description of the problem settings and introduces our learning strategies. Section III presents sufficient assumptions/lemmas and proves the convergence of the proposed algorithm. Section IV provides extensive numerical examples illustrating the theoretical results and demonstrating the effectiveness and applicability of the algorithm. The findings are concluded in Section V with possible future works.

D.Y.S., W.C.C., C.G., and S.Y.L. are with Academy for Engineering and Technology, Fudan University, Shanghai 200433, China. S.Y.L. is with Research Institute of Intelligent Complex Systems, Fudan University, Shanghai 200433, China. S.V. is with Department of Electrical and Electronic Engineering, Imperial College London, UK. Corresponding e-mails: {chunguan, syleng}@fudan.edu.cn.

II. PRELIMINARIES AND THE MODEL

A. Problem formulation

Consider a group of n agents, collectively trying to reveal the underlying true state of nature, denoted as θ^*, from a finite set of hypotheses $\Theta = \{\theta_1, \theta_2, \cdots, \theta_m\}$. At each time step $t = 1, 2, \cdots$, agent i obtains p types of observations $\{s_{i,t}^l\}_{l=1}^p$, which may come from multiple perspectives or represent different features of the true state. Each element $s_{i,t}^l$ is the realization of an environmental random variable $\boldsymbol{S}_{i,t}^l$. The set $\boldsymbol{s}_t^l = \{s_{1,t}^l, s_{2,t}^l, \cdots, s_{n,t}^l\}$ represents the actual observations made by all agents from signal type l at time t, generated according to the likelihood function $\boldsymbol{f}^l(\cdot)$ associated with the underlying true state θ^*. The set $\tilde{\boldsymbol{s}}_t = \{\boldsymbol{s}_t^1, \cdots, \boldsymbol{s}_t^p\}$ and $\tilde{\boldsymbol{f}} = \boldsymbol{f}^1 \times \cdots \times \boldsymbol{f}^p$. Each $\boldsymbol{S}_{i,t}^l$ has its individual observation space S_i^l and is i.i.d. with respect to t.

The signal structure for agent i with signal type l and possible state θ is described by a probability distribution $\ell_i^l(\cdot|\theta)$. In these settings, $\ell_i^l(s_{i,t}^l|\theta)$ indicates the likelihood of agent i observing type l signal $s_{i,t}^l$ at time t when it believes θ is the underlying true state.

The agents interact in a networked fashion, which is usually modelled by a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. $\mathcal{V} = \{1, 2, \cdots, n\}$ is the set of vertices representing the n agents, and $\mathcal{E} = \{(i,j)|\text{agent } j \text{ can receive information from agent } i\}$ is the set of directed edges. We denote $A = (a_{ij})_{n \times n}$ as the weight matrix of $\mathcal{G}$, which is assumed to be row-stochastic, i.e., $\sum_{j=1}^n a_{ij} = 1, \forall i = 1, \cdots, n$, and $a_{ij} > 0$ if $(j, i) \in \mathcal{E}$. The row-stochastic condition of A ensures that all agents assign normalized weights to the information, i.e., proportions of the total, that they receive from neighbors.

The *belief* of agent i at time t with signal type l is denoted as $\mu_{i,t}^l$, which is a probability distribution over the set of states Θ, i.e., $\sum_{k=1}^m \mu_{i,t}^l(\theta_k) = 1$, $\forall i = 1, \cdots, n, \forall l = 1, \cdots, p$, and $\forall t = 0, 1, \cdots$. Here $\mu_{i,0}^l$ represents the *initial belief* of agent i with signal type l.

Define a probability triple $(\Omega, \mathcal{F}, \mathbb{P}^*)$, where $\Omega = \{\omega|\omega = (\tilde{\boldsymbol{s}}_1, \tilde{\boldsymbol{s}}_2, \cdots)\}$, $\mathcal{F}$ is the σ-algebra generated by the observations, and $\mathbb{P}^*$ is the probability measure induced by paths in Ω, i.e., $\mathbb{P}^* = \prod_{t=1}^\infty \tilde{\boldsymbol{f}}$. We use $\mathbb{E}^*[\cdot]$ to denote the expectation operator associated with measure $\mathbb{P}^*$.

In this work, we consider the following two different circumstances:

Circumstance 1: For every agent i and every signal type l, the signal structure $\ell_i^l(\cdot|\theta^*)$ aligns with the i-th marginal distribution of $\boldsymbol{f}^l(\cdot)$ for all $l = 1, \cdots, p$, thereby characterizing the probability distribution of $\boldsymbol{S}_{i,t}^l$. In this case, all agents' *a priori* knowledge is accurate, and none of the signal types are misleading.

Circumstance 2: The condition in Circumstance 1 is not satisfied, and there may exist a signal type l such that $\boldsymbol{\ell}^l(\cdot|\theta^*)$ is not the best match of the real distribution $\boldsymbol{f}^l(\cdot)$ from the group's perspective. In this case, the group may experience false learning solely based on the type l signal.

The second circumstance could be quite common in practical applications, often arising from faults in signal perception or incorrect prior information due to a lack of training data.

B. Social Learning Strategies with Multiple Signals

Non-Bayesian social learning typically involves two steps for agents to update their beliefs at each time, i.e., the Bayesian update step and the aggregation of neighbors' beliefs [22]–[24]. In the belief aggregation step, every agent shares its current belief with its one-hop neighbors, whereas in the Bayesian update step, every agent combines its prior belief with observations from the environment to form its posterior belief.

Traditionally, when dealing with tasks involving multiview observations, social learning algorithms integrate these diverse signals into a single signal and design a joint likelihood function as the signal structure. This approach demands a thorough understanding of the correlations among different viewpoints of observations, often making it challenging to achieve in practical tasks. In our work, however, we allow agents to independently perform Bayesian inference for each signal type and integrate information from multiview observations during the information aggregation process.

The algorithm we propose can be described in the following two steps:

1) Information aggregation. For each agent i and signal type $l = 1, \cdots, p$, we calculate the updated belief using the formula:

$$\tilde{\mu}_{i,t+1}^l(\theta) =$$

$$\frac{\exp\left(\gamma_l \sum_{j=1}^n a_{ij} \log \mu_{j,t}^l(\theta) + \sum_{k \neq l}^p \gamma_k \log \mu_{i,t}^k(\theta)\right)}{\sum_{\theta' \in \Theta} \exp\left(\gamma_l \sum_{j=1}^n a_{ij} \log \mu_{j,t}^l(\theta') + \sum_{k \neq l}^p \gamma_k \log \mu_{i,t}^k(\theta')\right)},$$

where the assigned parameter $\gamma_l \in (0, 1)$ for all $l = 1, \cdots, p$, and $\sum_{l=1}^p \gamma_l = 1$.

2) Bayesian update. For each agent i and signal type $l = 1, \cdots, p$, the posterior belief is given by:

$$\mu_{i,t+1}^l(\theta) = \frac{\tilde{\mu}_{i,t+1}^l(\theta)\ell_i^l(s_{i,t+1}^l|\theta)}{\sum_{\theta' \in \Theta} \tilde{\mu}_{i,t+1}^l(\theta')\ell_i^l(s_{i,t+1}^l|\theta')}.$$

III. ASSUMPTIONS AND RESULTS

As widely discussed in previous works of social learning, we care about the convergence of the algorithms as well as the rate of convergence. The following assumptions are required to ensure the convergence of the proposed social learning strategies:

Assumption 1 (Communication network): The graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and its weight matrix A satisfy that:
  a) The graph is strongly-connected;
  b) A has at least one positive diagonal entry.

Assumption 1 ensures that A is the transition matrix of an irreducible, aperiodic Markov chain with finite states. We recall the following lemma [25]:

Lemma 1: If a Markov chain with finite states is irreducible, then it has a unique stationary distribution π. Let A be the transition matrix of the Markov chain and further suppose it is aperiodic, then we have $\lim_{k \to \infty} [A^k]_{ij} = \pi_j$, for $1 \leq i, j \leq n$.

The stationary distribution π can be interpreted as the normalized left eigenvector of A corresponding to eigenvalue 1, known as the *eigenvector centrality* in related literature. The Perron-Frobenius theorem ensures that all components of π are strictly positive.

Assumption 2 (Belief and signal structure): Every agent $i = 1, \cdots, n$ in the group satisfies:

a) It has positive initial beliefs on all states regarding all types of signals, i.e., $\mu_{i,0}^l(\theta) > 0$ for all $l = 1, \cdots, p$ and $\theta \in \Theta$;

b) The logarithms of its signal structures are integrable, i.e., $\mathbb{E}^* \left[|\log \ell_i^l(s_i^l|\theta)| \right] < \infty$ for all $l = 1, \cdots, p$, $s_i^l \in S_i^l$, and $\theta \in \Theta$.

Assumption 2a) is imposed to ensure the well-definedness of $\log \mu_{i,t}^l(\cdot)$. Meanwhile, Assumption 2b) guarantees that $\log \ell_i^l(s_i^l|\theta)$ is real-valued almost surely [26]. In practical scenarios where the signal structures of the agents are Gaussian, Assumption 2b) holds naturally since Gaussian random variables are square integrable.

Two states, θ_j and θ_k, are called *observationally equivalent* with signal type l for agent i if $\ell_i^l(s_i^l|\theta_j) = \ell_i^l(s_i^l|\theta_k)$ for all $s_i^l \in S_i^l$, in which case the agent can not distinguish between these states using its own information obtained from type l signal. The true state is called *globally identifiable* if the set $\hat{\Theta} = \bigcap_{l=1}^{p} \bigcap_{i=1}^{n} \hat{\Theta}_i^l$ has only one element θ^*, where $\hat{\Theta}_i^l = \{\theta \in \Theta | \ell_i^l(s_i|\theta) = \ell_i^l(s_i^l|\theta^*), \forall s_i^l \in S_i^l\}$. Intuitively, if a state θ' is observationally equivalent to θ^* with all types of signals for all agents, i.e., $\hat{\Theta} = \{\theta^*, \theta'\}$, then the two states are indistinguishable from the view of all agents, and they can not collectively learn the underlying true state.

To ensure the convergence of groups' beliefs on the true state, we introduce the following assumption:

Assumption 3 (Globally identifiable): The true state θ^* is globally identifiable.

Under this assumption, for all $\theta \neq \theta^*$, there exists at least one agent i and a signal type l such that $D_{\text{KL}}(\ell_i^l(\cdot|\theta^*) \parallel \ell_i^l(\cdot|\theta))$ is strictly positive, where $D_{\text{KL}}(P \parallel Q)$ represents the Kullback-Leibler divergence between two probability distributions P and Q.

Denote in the following that $K_i^l(\theta^*, \theta) = D_{\text{KL}}(f_i^l(\cdot) \parallel \ell_i^l(\cdot|\theta^*)) - D_{\text{KL}}(f_i^l(\cdot) \parallel \ell_i^l(\cdot|\theta))$. Its positivity or negativity depends on whether, from the perspective of agent i, state θ or θ^* is more likely to be the underlying true state. Notice that under Circumstance 1, $K_i^l(\theta^*, \theta) = -D_{\text{KL}}(\ell_i^l(\cdot|\theta^*) \parallel \ell_i^l(\cdot|\theta))$. Now we can state the main results describing the correct convergence of the proposed strategy.

Theorem 1: Under Circumstance 1 and Assumptions 1, 2 and 3, the proposed social learning strategy satisfies:

$$\lim_{t \to \infty} \mu_{i,t}^l(\theta^*) = 1, \quad \mathbb{P}^* - \text{a.s.}, \quad \forall 1 \leq i \leq n, 1 \leq l \leq p.$$

Proof: For each agent i, signal type l, and $\theta \neq \theta^*$, we have

$$\log \frac{\mu_{i,t+1}^l(\theta)}{\mu_{i,t+1}^l(\theta^*)} = \gamma_l \sum_{j=1}^{n} a_{ij} \log \frac{\mu_{j,t}^l(\theta)}{\mu_{j,t}^l(\theta^*)} + \sum_{k \neq l}^{p} \gamma_k \log \frac{\mu_{i,t}^k(\theta)}{\mu_{i,t}^k(\theta^*)}$$
$$+ \log \frac{\ell_i^l(s_{i,t+1}^l|\theta)}{\ell_i^l(s_{i,t+1}^l|\theta^*)}.$$

By denoting $\nu_{i,t+1}^l(\theta) = \log \frac{\mu_{i,t+1}^l(\theta)}{\mu_{i,t+1}^l(\theta^*)}$ and $L_{i,t+1}^l(\theta) = \log \frac{\ell_i^l(s_{i,t+1}^l|\theta)}{\ell_i^l(s_{i,t+1}^l|\theta^*)}$, the above equation simplifies to

$$\nu_{i,t+1}^l(\theta) = \gamma_l \sum_{j=1}^{n} a_{ij} \nu_{j,t}^l(\theta) + \sum_{k \neq l}^{p} \gamma_k \nu_{i,t}^k(\theta) + L_{i,t+1}^l(\theta). \tag{1}$$

Define the n-dimensional column vector $\boldsymbol{\nu}_t^l(\theta) = \left(\nu_{1,t}^l(\theta), \cdots, \nu_{n,t}^l(\theta)\right)^\top$ for each $l = 1, \cdots, p$ and the np-dimensional column vector $\tilde{\boldsymbol{\nu}}_t(\theta) = \left(\boldsymbol{\nu}_t^1(\theta)^\top, \cdots, \boldsymbol{\nu}_t^p(\theta)^\top\right)^\top$. Similarly, $\boldsymbol{L}_t^l(\theta) = \left(L_{1,t}^l(\theta), \cdots, L_{n,t}^l(\theta)\right)^\top$ and $\tilde{\boldsymbol{L}}_t(\theta) = \left(\boldsymbol{L}_t^1(\theta)^\top, \cdots, \boldsymbol{L}_t^p(\theta)^\top\right)^\top$. Additionally, denote the matrix

$$\tilde{A} = \begin{bmatrix} \gamma_1 A & \gamma_2 I & \cdots & \gamma_p I \\ \gamma_1 I & \gamma_2 A & \cdots & \gamma_p I \\ \vdots & \vdots & \ddots & \vdots \\ \gamma_1 I & \gamma_2 I & \cdots & \gamma_p A \end{bmatrix},$$

where I is the identity matrix. It is evident that $\tilde{A}$ is a row-stochastic matrix, we further demonstrate that $\tilde{A}$ serves as the transition matrix for an irreducible, aperiodic Markov chain with finite states. Given that A has at least one positive diagonal element, $\tilde{A}$ contains a minimum of n positive diagonal elements, making it aperiodic.

We then examine the strong connectivity of the corresponding graph of the $np \times np$ matrix $\tilde{A}$ to prove its irreducibility. For any node with index $i = n(l-1) + i_0$, where $1 \leq i_0 \leq n$ and $1 \leq l \leq p$, it can establish a path to any node with index j within the range $n(l-1)+1 \leq j \leq nl$ due to the irreducibility of A. We need to further find a path from node i to any node j, where $j = n(k-1) + j_0$, $1 \leq j_0 \leq n$, and $k \neq l$. Since $\tilde{a}_{n(k-1)+i_0,i} = [\gamma_l I]_{i_0 i_0} = \gamma_l > 0$, a path P_1 from node i to node $n(k-1) + i_0$ exists. Additionally, node $n(k-1)+i_0$ can find a path P_2 to node j as previously demonstrated. Combining these paths as $P = P_1 \cup P_2$ establishes a path from node i to node j. Thus, we can conclude that every node i has a path to any node j in the graph corresponding to matrix $\tilde{A}$, confirming its irreducibility.

Subsequently we can rewrite (1) in matrix form:

$$\tilde{\boldsymbol{\nu}}_{t+1}(\theta) = \tilde{A} \tilde{\boldsymbol{\nu}}_t(\theta) + \tilde{\boldsymbol{L}}_{t+1}(\theta).$$

Now it follows that

$$\frac{1}{t}\tilde{\boldsymbol{\nu}}_{t+1}(\theta) = \frac{1}{t}\tilde{A}\tilde{\boldsymbol{\nu}}_t(\theta) + \frac{1}{t}\tilde{\boldsymbol{L}}_{t+1}(\theta) = \cdots$$

$$= \frac{1}{t}\tilde{A}^{t+1}\tilde{\boldsymbol{\nu}}_0(\theta) + \frac{1}{t}\sum_{k=1}^{t}\tilde{A}^k\tilde{\boldsymbol{L}}_{t+1-k}(\theta) + \frac{1}{t}\tilde{\boldsymbol{L}}_{t+1}(\theta). \tag{2}$$

The assumptions admit that the first and the third terms on r.h.s. of (2) go to zero as $t \to \infty$. The second term can be deformed as

$$\frac{1}{t}\sum_{k=1}^{t}\tilde{A}^k\tilde{\boldsymbol{L}}_{t+1-k}(\theta) = \frac{1}{t}\sum_{k=1}^{t}(\tilde{A}^k - \mathbf{1}_{np}\tilde{\pi})\tilde{\boldsymbol{L}}_{t+1-k}(\theta)$$

$$+ \frac{1}{t}\sum_{k=1}^{t}\mathbf{1}_{np}\tilde{\pi}(\tilde{\boldsymbol{L}}_{t+1-k}(\theta) - \tilde{\boldsymbol{K}}(\theta^*, \theta))$$

$$+ \frac{1}{t}\sum_{k=1}^{t}\mathbf{1}_{np}\tilde{\pi}\tilde{\boldsymbol{K}}(\theta^*, \theta), \tag{3}$$

where $\mathbf{1}_{np}$ is an np-dimensional column vector of ones, $\tilde{\pi}$ is the eigenvector centrality corresponding to matrix $\tilde{A}$ and is a row vector, $\tilde{\boldsymbol{K}}(\theta^*, \theta) = \left(\boldsymbol{K}^1(\theta^*, \theta)^\top, \cdots, \boldsymbol{K}^p(\theta^*, \theta)^\top\right)^\top$, and $\boldsymbol{K}^l(\theta^*, \theta) = \left(K_1^l(\theta^*, \theta), \cdots, K_n^l(\theta^*, \theta)\right)^\top$. Lemma 1 admits that $\lim_{k\to\infty}\tilde{A}^k = \mathbf{1}_{np}\tilde{\pi}$. Noticing that all elements of $\tilde{A}^k(k = 1, 2, \cdots)$ are bounded, the first term on r.h.s. of (3) converges to zero as $t \to \infty$. Moreover, under Circumstance 1, for all $l = 1, \cdots, p$ we have

$$\mathbb{E}^*[L_{i,t}^l(\theta)] = \mathbb{E}^*\left[\log\frac{\ell_i^l(s_{i,t}^l|\theta)}{\ell_i^l(s_{i,t}^l|\theta^*)}\right]$$

$$= \int_{s^l \in S_i^l}\ell_i^l(s^l|\theta^*)\log\frac{\ell_i^l(s^l|\theta)}{\ell_i^l(s^l|\theta^*)}\mathrm{d}s^l$$

$$= -D_{\mathrm{KL}}(\ell_i^l(\cdot|\theta^*)\,\|\,\ell_i^l(\cdot|\theta)) = K_i^l(\theta^*, \theta).$$

The Kolmogorov's strong law of large numbers gives that $\forall l = 1, \cdots, p$,

$$\frac{1}{t}\sum_{k=1}^{t}\boldsymbol{L}_{t+1-k}^l(\theta) - \frac{1}{t}\sum_{k=1}^{t}\mathbb{E}^*[\boldsymbol{L}_{t+1-k}^l(\theta)] \to \mathbf{0}, \quad \mathbb{P}^*-\text{a.s.},$$

as $t \to \infty$, which leads to

$$\lim_{t\to\infty}\frac{1}{t}\sum_{k=1}^{t}\mathbf{1}_{np}\tilde{\pi}(\tilde{\boldsymbol{L}}_{t+1-k}(\theta) - \tilde{\boldsymbol{K}}(\theta^*, \theta)) = \mathbf{0}, \quad \mathbb{P}^*-\text{a.s.}.$$

Now (3) gives that

$$\lim_{t\to\infty}\frac{1}{t}\sum_{k=1}^{t}\tilde{A}^k\tilde{\boldsymbol{L}}_{t+1-k}(\theta) = \mathbf{1}_{np}\tilde{\pi}\tilde{\boldsymbol{K}}(\theta^*, \theta), \quad \mathbb{P}^*-\text{a.s.}.$$

According to Assumption 3 and Lemma 1, for all $\theta \neq \theta^*$ we have

$$\lim_{t\to\infty}\frac{1}{t}\tilde{\boldsymbol{\nu}}_{t+1}(\theta) < \mathbf{0}, \quad \mathbb{P}^*-\text{a.s.}. \tag{4}$$

Thus $\nu_{i,t+1}^l(\theta) \to -\infty$ almost surely for all agents $i = 1, \cdots, n$ and signal types $l = 1, \cdots, p$. This implies
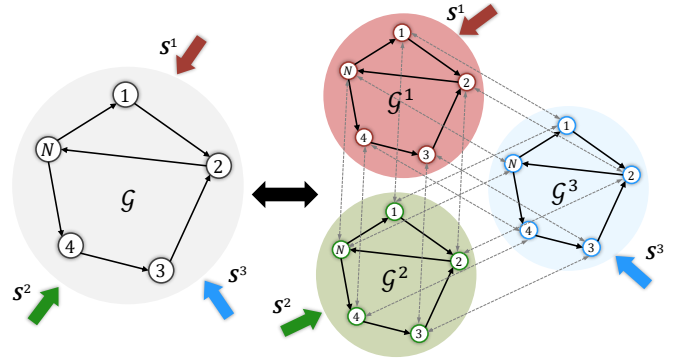

Fig. 1. An intuitive illustration for an understanding of the proposed algorithm.

$\mu_{i,t}^l(\theta) \to 0$ for all $i = 1, \cdots, n$ and $l = 1, \cdots, p$ almost surely. ■

In the proof of Theorem 1, it is noteworthy that we construct a new row-stochastic matrix $\tilde{A}$ and demonstrate its primitivity. Hence, our proposed algorithm can be viewed as duplicating the network $\mathcal{G}$ of agents into p identical subnetworks $\mathcal{G}^1, \cdots, \mathcal{G}^p$, establishing bidirectional links between each node and its duplicate, and assigning a distinct signal to each subnetwork for classic non-Bayesian social learning with geometric averaging. The weight matrix corresponding to the augmented network is exactly $\tilde{A}$. An illustration is shown in Fig. 1.

Theorem 1 guarantees that all agents will eventually learn the underlying true state with our learning strategy as long as some fundamental assumptions are satisfied. Notably, our algorithm does not require every type of signal to be informative for the group. As long as, for every pair of states θ and θ^*, there exist an agent capable of distinguishing between them with a certain type of signal, the entire group can achieve correct learning. Subsequently, we will demonstrate that, in certain scenarios, our method is capable of learning correctly even in the presence of misleading signals.

From the proof of Theorem 1, as long as (4) is satisfied, all agents will learn the underlying true state correctly. Therefore, it is important to figure out $\tilde{\pi}$, leading to the following lemma.

Lemma 2: Let $\tilde{\pi}$ be the normalized left eigenvector of matrix $\tilde{A}$ associated with eigenvalue 1, and π be the normalized left eigenvector of matrix A associated with eigenvalue 1. Then, we have $\tilde{\pi} = (\gamma_1\pi, \cdots, \gamma_p\pi)$.

Proof:

$$\tilde{\pi}\tilde{A} = (\gamma_1(\gamma_1\pi A + \gamma_2\pi + \cdots + \gamma_p\pi), \cdots,$$
$$\gamma_p(\gamma_1\pi + \gamma_2\pi + \cdots + \gamma_p\pi A))$$
$$= (\gamma_1\pi, \cdots, \gamma_p\pi) = \tilde{\pi}.$$

The second equality follows from $\pi A = \pi$ and $\sum_{l=1}^{p}\gamma_l = 1$. ■

Theorem 2: Under Circumstance 2 and Assumptions 1 and 2, all agents will correctly learn the underlying true state,

i.e.,

$$\lim_{t \to \infty} \mu_{i,t}^l(\theta^*) = 1, \quad \mathbb{P}^*-\text{a.s.}, \quad \forall 1 \le i \le n, 1 \le l \le p,$$

if and only if

$$\sum_{l=1}^{p} \gamma_l \sum_{i=1}^{n} \pi_i K_i^l(\theta^*, \theta) < 0.$$

Proof: Firstly, similar to the proof of Theorem 1, under Circumstance 2 we have

$$
\begin{aligned}
\mathbb{E}^* \left[L_{i,t}^l(\theta) \right] &= \mathbb{E}^* \left[\log \frac{\ell_i^l(s_{i,t}^l|\theta)}{\ell_i^l(s_{i,t}|\theta^*)} \right] \\
&= \int_{s^l \in S_i^l} f_i^l(s^l) \log \frac{\ell_i^l(s^l|\theta)}{\ell_i^l(s^l|\theta^*)} \mathrm{d}s^l \\
&= \int_{s^l \in S_i^l} f_i^l(s^l) \left(\log \frac{f_i^l(s^l)}{\ell_i^l(s^l|\theta^*)} - \log \frac{f_i^l(s^l)}{\ell_i^l(s^l|\theta)} \right) \mathrm{d}s^l \\
&= D_{\mathrm{KL}}(f_i^l(\cdot) || \ell_i^l(\cdot|\theta^*)) - D_{\mathrm{KL}}(f_i^l(\cdot) || \ell_i^l(\cdot|\theta)) \\
&= K_i^l(\theta^*, \theta).
\end{aligned}
$$

and

$$
\begin{aligned}
\lim_{t \to \infty} \frac{1}{t} \tilde{\boldsymbol{\nu}}_{t+1}(\theta) &= \mathbf{1}_{np} \tilde{\pi} \tilde{K}(\theta^*, \theta) \\
&= \mathbf{1}_p \otimes \sum_{l=1}^{p} \gamma_l \mathbf{1}_n \boldsymbol{\pi} \boldsymbol{K}^l(\theta^*, \theta),
\end{aligned}
\tag{5}
$$

$$\mathbf{1}_n \boldsymbol{\pi} \boldsymbol{K}^l(\theta^*, \theta) = \sum_{i=1}^{n} \pi_i K_i^l(\theta^*, \theta) \mathbf{1}_n,$$

hence the condition $\sum_{l=1}^{p} \gamma_l \mathbf{1}_n \boldsymbol{\pi} \boldsymbol{K}^l(\theta^*, \theta) < \mathbf{0}$ holds if and only if $\sum_{l=1}^{p} \gamma_l \sum_{i=1}^{n} \pi_i K_i^l(\theta^*, \theta) < 0$. ∎

Theorem 2 demonstrates the robustness of our algorithm when dealing with multiple signals. Even if some types of signals might be misleading, the likelihood of collective mis-learning can be reduced by adjusting the assigned parameter γ_l. Additionally, as can be seen from (5), assigning a higher weight γ_l to signals that are more instructive, i.e., better able to help the group distinguish between correct and incorrect states, can accelerate the convergence rate of the algorithm.

IV. NUMERICAL EXAMPLES

A. Learning with multiview observations

We first demonstrate that our proposed multiview observations algorithm can address the observational equivalence issue present when only a single viewpoint of information is available.

Example 1: Consider a strongly-connected network consisting of two agents. The corresponding weight matrix is

$$A = \begin{bmatrix} 0 & 1 \\ 0.7 & 0.3 \end{bmatrix}.$$

The two agents are engaged in the task of localizing a target situated within a 4×4 grid. They receive Gaussian signals (type 1) with mean values corresponding to the distances
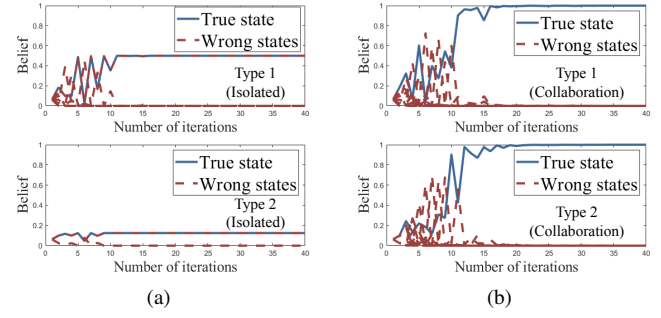

Fig. 2. The evolution of beliefs of Agent 1 on different states. (a) The two agents are unable to identify the underlying true state with a single type of signal. (b) The two agents achieve correct learning by combining the information from two types of signals.
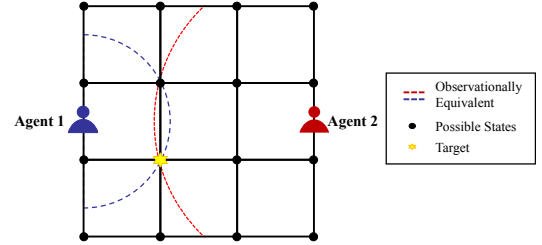

Fig. 3. Illustration of the scenario in Example 1. In this example, the two agents, due to the observational equivalence problem, cannot achieve correct learning relying solely on a single type of signal.

from the target, which could be at any of the 16 grid points. The signal structure of the two agents with respect to θ also follows a Gaussian distribution, with the mean value equal to the true distance between the agent and θ. As shown in Fig. 3, the two agents struggle to distinguish between two states due to the overlapping circles centered on these states, where each circle has a radius equal to the distance from the target.

At the same time, both agents can receive signals (type 2) regarding whether the target is above or below them. If the target is above an agent, at each moment, there is a 0.8 probability of receiving signal U and a 0.2 probability of receiving signal D. The signal structure is set as $\ell_i^2(U|\theta) = 0.8$ if θ is located above i and $\ell_i^2(U|\theta) = 0.2$ if θ is located below i. It is obvious that both of the two agents can not identify the underlying true state based solely on type 2 signal.

The initial beliefs about both types of signal are uniform distributions over all possible states. Under our settings, the signal structures of all agents and all types of signal about the true state θ^* align with its actual distribution, which satisfy the conditions in Circumstance 1.

The experimental results indicate that the agents are unable to collectively learn the true state according to one type of signal solely, as shown in Fig. 2 (a). However, by combining the information provided by both types of signals, the two agents can successfully collaborate to identify the underlying true state due to the fact that Assumption 3 is satisfied, as shown in Fig. 4 (b). This example demonstrates that learning from multiple signals can resolve the issue of observation-
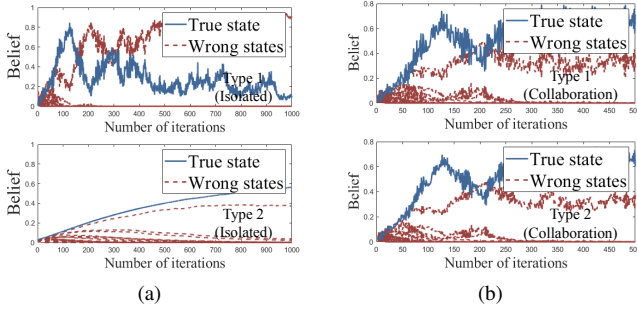
Fig. 4. The evolution of beliefs of Agent 1 on all possible states in the first scenario of distributed cooperative localization task. (a) The agent can identify the optimal state solely based on azimuth information, but using only distance information results in erroneous learning. (b) By employing our algorithm to integrate the information from both types of signals, the beliefs of the agent converge to the true state.
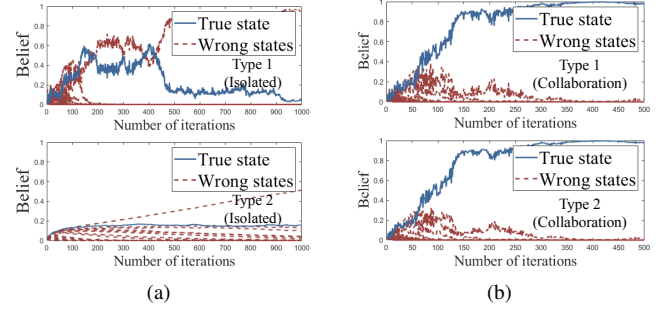

Fig. 5. The evolution of beliefs of Agent 1 on all possible states in the second scenario of distributed cooperative localization task. (a) The agents cannot achieve correct learning solely relying on distance or azimuth information. (b) By employing our algorithm to integrate the information from both types of signals, the agents can learn the underlying true state asymptotically.

ally equivalence present in single-signal scenarios, thereby offering more tolerant conditions for successful learning.

B. Learning with misleading signal

In the following example, we will demonstrate how our proposed algorithm addresses the issue of erroneous learning that may occur with a single viewpoint of signal by aggregating information from multiview observations, thereby validating the enhanced robustness of our algorithm.

Example 2: Consider a scenario where a group of N sensors is randomly distributed in a two-dimensional square, denoted as $[0,1]^2$. Each sensor receives two viewpoints of observations at every time step, one is related to distance and the other related to orientation. The first type is a Gaussian signal with noise, representing the distance from the sensor to the target. Specifically, $\boldsymbol{S}_{i,t}^1 \sim \mathcal{N}(d_i^*, \sigma_1^2)$ for all $i = 1, \cdots, N$ and $t = 1, 2, \cdots$, where d_i^* denotes the distance from the i-th sensor to the target. The second type of signal also contains Gaussian noise and pertains to the azimuth between the agent and the target, and $\boldsymbol{S}_{i,t}^2 \sim \mathcal{N}(\alpha_i^*, \sigma_1^2)$, where α_i^* denotes the azimuth. There are M possible positions θ_m ($m = 1, \cdots, M$) uniformly distributed in the square space. Let d_i^m and α_i^m denote the distance and azimuth from sensor i to a possible state θ_m respectively. The collective aim of these sensors is to find the state that has a position closest to the target, and it is conceived as underlying true state θ^*. For every agent i, we set $\ell_i^1(\cdot|\theta_m)$ to align with $\mathcal{N}(d_i^m, \sigma_1^2)$, and since the orientation serves as auxiliary information, we let $\ell_i^2(\cdot|\theta_m) \sim \mathcal{N}(\alpha_i^m, 10\sigma_2^2)$.

We set $N = 2$, $M = 36$, $\sigma_1 = \sigma_2 = 0.5$, and $\gamma_1 = \gamma_2 = 0.5$. The weight matrix is the same as that in Example 1, indicating that the corresponding network is strongly-connected. The initial beliefs about both types of signal are uniform distributions over all possible states. In situations where the target does not coincide with any possible state, our experimental setup clearly meets the conditions of Circumstance 2. In this case, based on the proof of Theorem 2 and through calculation, the collective learning
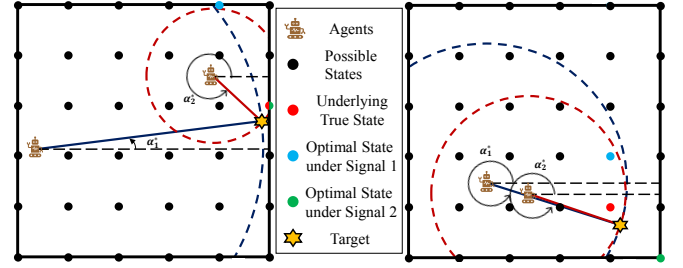

Fig. 6. Illustrations of two typical scenarios. In the first scenario, the group incorrectly selects the optimal state based on distance information. In the second scenario, due to the similar orientation information of the two agents, no type of signals leads to correct learning.

outcome $\hat{\theta}$ satisfies:

$$\hat{\theta} = \theta_{j_0},$$

$$j_0 = \arg\min_{j=1,\cdots,m} \sum_{i=1}^2 \pi_i \left((d_i^* - d_i^j)^2 + \frac{1}{10}(\alpha_i^* - \alpha_i^j)^2 \right),$$

where $\boldsymbol{\pi} = (\pi_1, \pi_2)$ is the eigenvector centrality corresponding to the weight matrix.

We fix the position of the target and allow agents to be distributed in different locations to study the collaborative learning outcomes with two types of signals separately and jointly. In 1000 experiments, our proposed method, which combines distance and azimuth information, successfully achieves localization 853 times, while relying solely on distance or azimuth information leads to successful localization 550 and 582 times, respectively. This clearly demonstrates the benefits of integrating both types of signal information.

We will further present the detailed results of two typical scenarios. In the first scenario, agents relying solely on distance information fail to find the state closest to the target, while exclusive reliance on azimuth information results in successful localization, as shown in Fig. 4 (a). If they combine both types of information, all agents could successfully achieve the task, as shown in Fig. 4 (b). In the second scenario, agents relying solely on either distance or azimuth information are unable to achieve accurate localization, as evidenced in Fig. 5 (a). However, after applying our method,

agents are able to successfully identify the state closest to the target, as demonstrated in Fig. 5 (b). The schematic diagrams of agent and state positions in two scenarios are presented in Fig. 6.

This experiment, serving as a complement to Theorem 2, demonstrates that our proposed multiview observations algorithm can, to a certain extent, alleviate the impact of misleading signals, thereby increasing the fault tolerance of collective learning.

V. CONCLUSION AND FUTURE WORK

In this paper, we extend traditional non-Bayesian social learning algorithms designed for single signal and propose a distributed information processing algorithm that integrates information from multiview observations. Our proposed algorithm enables the group to learn from multiple viewpoints of information independently and achieving interaction among multiview observations during the information aggregation step. By introducing weight parameters for various signal types, we not only ensure the convergence of the algorithm under traditional assumptions, but also, in certain scenarios, correct errors introduced by a single view of signal, significantly enhancing the fault tolerance of collective learning.

Our work not only presents a distributed information processing algorithm capable of handling a more diverse range of task scenarios but also contributes to the theoretical foundation of distributed machine learning. In future endeavors, we plan to explore the application of non-Bayesian social learning algorithms with multiview observations in the design of distributed machine learning methods to address challenges associated with multi-feature or high-dimensional problems. Additionally, we may consider other interaction mechanisms among multiple signals, such as negative feedback, and introduce more parameter settings to enable applications in a wider range of tasks.

REFERENCES

[1] S. Kar and J. M. F. Moura, "Distributed consensus algorithms in sensor networks with imperfect communication: Link failures and channel noise," *IEEE Transactions on Signal Processing*, vol. 57, no. 1, pp. 355–369, 2009.

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

[3] J. Chen and A. H. Sayed, "Diffusion adaptation strategies for distributed optimization and learning over networks," *IEEE Transactions on Signal Processing*, vol. 60, no. 8, pp. 4289–4305, 2012.

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

[5] S. Vlaski, S. Kar, A. H. Sayed, and J. M. Moura, "Networked signal and information processing: Learning by multiagent systems," *IEEE Signal Processing Magazine*, vol. 40, no. 5, pp. 92–105, 2023.

[6] 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.

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

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

[9] M. A. Rahimian, P. Molavi, and A. Jadbabaie, "(non-)bayesian learning without recall," in *53rd IEEE Conference on Decision and Control*, 2014, pp. 5730–5735.

[10] A. Mitra, J. A. Richards, and S. Sundaram, "A new approach for distributed hypothesis testing with extensions to byzantine-resilience," in *2019 American Control Conference (ACC)*, 2019, pp. 261–266.

[11] A. NediÄǦ, A. Olshevsky, and C. A. Uribe, "Fast convergence rates for distributed non-bayesian learning," *IEEE Transactions on Automatic Control*, vol. 62, no. 11, pp. 5538–5553, 2017.

[12] ——, "Nonasymptotic convergence rates for cooperative learning over time-varying directed graphs," in *2015 American Control Conference (ACC)*, 2015, pp. 5884–5889.

[13] H. Salami, B. Ying, and A. H. Sayed, "Social learning over weakly connected graphs," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 3, no. 2, pp. 222–238, 2017.

[14] C. A. Uribe and A. Jadbabaie, "On increasing self-confidence in non-bayesian social learning over time-varying directed graphs," in *2019 American Control Conference (ACC)*, 2019, pp. 3532–3537.

[15] D. Sui, C. Guan, Z. Gan, W. Lin, and S. Leng, "Tuning convergence rate via non-bayesian social learning: A trade-off between internal belief and external information," in *2023 62nd IEEE Conference on Decision and Control (CDC)*, 2023, pp. 1381–1387.

[16] K. Ntemos, V. Bordignon, S. Vlaski, and A. H. Sayed, "Social learning with disparate hypotheses," in *2022 30th European Signal Processing Conference (EUSIPCO)*, 2022, pp. 2171–2175.

[17] ——, "Social learning under inferential attacks," in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 5479–5483.

[18] M. Z. A. Bhotto and W. P. Tay, "Non-bayesian social learning with observation reuse and soft switching," *ACM Transactions on Sensor Networks (TOSN)*, vol. 14, no. 2, pp. 1–21, 2018.

[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 *2019 22th International Conference on Information Fusion (FUSION)*, 2019, pp. 1–8.

[20] J. Z. Hare, C. A. Uribe, L. Kaplan, and A. Jadbabaie, "Non-bayesian social learning with uncertain models," *IEEE Transactions on Signal Processing*, vol. 68, pp. 4178–4193, 2020.

[21] Q. Chen, W. Shi, D. Sui, and S. Leng, "Distributed consensus algorithms in sensor networks with higher-order topology," *Entropy*, vol. 25, no. 8, p. 1200, 2023.

[22] Q. Liu, J. Zhao, and X. Wang, "Distributed detection via bayesian updates and consensus," in *2015 34th Chinese Control Conference (CCC)*. IEEE, 2015, pp. 6992–6997.

[23] A. Lalitha, A. Sarwate, and T. Javidi, "Social learning and distributed hypothesis testing," in *2014 IEEE International Symposium on Information Theory*. IEEE, 2014, pp. 551–555.

[24] A. Nedić, A. Olshevsky, and C. A. Uribe, "Fast convergence rates for distributed non-bayesian learning," *IEEE Transactions on Automatic Control*, vol. 62, no. 11, pp. 5538–5553, 2017.

[25] P. G. Hoel, S. C. Port, and C. J. Stone, *Introduction to Stochastic Processes*. Long Grove, IL, USA: Waveland Press, 1986.

[26] E. Cinlar and E. ðCınlar, *Probability and stochastics*. Springer, 2011, vol. 261.