2023 62nd IEEE Conference on Decision and Control (CDC)
December 13-15, 2023. Marina Bay Sands, Singapore

Tuning Convergence Rate via Non-Bayesian Social Learning: A Trade-Off between Internal Belief and External Information

Dongyan Sui, Chun Guan, Zhongxue Gan, Wei Lin, Siyang Leng

Abstract— **Social learning strategies have been recently developed for multi-agents to learn progressively an underlying state of nature by information communications and evolutions. Existing works define algorithms mainly by swapping the Bayesian update and belief aggregation steps and/or discovering diverse underlying network structures. Inspired by the diversity of agents when they are exposed to new information, this work designs a non-Bayesian learning strategy, named as Parametric Social Learning, by introducing an agent stubbornness parameter to trade-off the significance between its internal belief and external information. This strategy thus allows for tuning the convergence rate by adjusting the introduced parameter, which is consistent highly with the sociological intuition. Theoretical analyses and numerical examples are provided to illustrate several sociological insights. Our work therefore has appealing potential in practical tasks such as dispersed information aggregation and distributed parameter estimation.**

I. Introduction

Collective consensus is generally reached through local exchange of information. This topic attracts extensive studies on distributed learning over the last few decades. Typical scenarios include the dispersed information aggregation problem in multi-agent systems and distributed parameter estimation task in data communication networks. In social networks, modeling and regulating the opinion formation are essential topics [1]–[3], motivating the development of social learning strategies for agents to learn progressively an underlying state of nature by information communications and evolutions.

Banerjee [4] and Bikhchandani et al. [5] seminally formulate the social learning paradigm in a fully Bayesian manner, and Smith and Sørensen [6] further provide a comprehensive analysis of this environment, introducing the important concepts of bounded and unbounded beliefs. Because of the requirement for *a priori* information and the computational burden, Bayesian learning is prohibitive even in simple networks [7]. Non-Bayesian algorithms [8]–[11], consisting of a *belief aggregation* step and a *Bayesian update* step, are then introduced into social learning and prosperously developed following the celebrated work of Jadbabaie et al. in [12]–[14]. In [15], a distributed parameter estimation model is presented using logarithmic aggregation, with its convergence and asymptotic normality being proved. Kar et al. [16] design a set of distributed parameter estimation algorithms by combining a consensus step and an innovation step in the update rule and apply them in sensor networks. Shahrampour et al. [17] and Nedić et al. [18] consider and prove the convergence of similar learning rules, under the assumption of bounded ratios of likelihood functions, while the latter further analyze the learning rule for time-varying graphs [19]. Convergence result of non-Bayesian learning algorithm for fixed graphs is provided in [20] and large deviation convergence rates are given, proving the existence of a random time after which the beliefs will concentrate exponentially fast. Authors in [21] consider the learning rules on weakly-connected graphs, and social learning with time-varying weights is studied in [22]. Recently, Bordignon et al. [23] propose a novel learning strategy, called adaptive social learning, addressing the poor performance under nonstationary conditions. A comprehensive review can be referred to [24].

However, existing works extend the field of social learning mainly through swapping the *update* and *aggregation* steps and/or discovering diverse underlying network structures [25], neglecting the diversity of agents. In practical terms, the stubbornness or openness varies from individual to individual, hence agents will assign weights to trade-off the significance between their internal beliefs and external information. This inspires us to design a new learning strategy, named as Parametric Social Learning (PSL), allowing for tuning the convergence rate by adjusting the introduced parameter describing the stubbornness of the agents. The Bayesian update step of the proposed strategy arises naturally from an optimization problem balancing the internal belief and external information. We demonstrate the importance of being able to tune the convergence rate through presenting the fact that too quick convergence will even result in false learning. Our analysis reveals the sociological phenomenon that open-mind contributes to faster reach of consensus. Numerical examples also provide interesting insights that in a group open agent with higher social influence and/or being more informative accelerates the learning time.

The remaining part of this paper is organized as follows: Section II provides a full description of the problem settings, reviewing necessary algorithms and proposing our PSL strategies. Section III presents sufficient assumptions/lemmas and proves the convergence of the proposed algorithms. Section IV provides extensive numerical examples illustrating the theoretical results and revealing sociological insights. The findings are concluded in Section V with possible future works directed.

D.Y.S., C.G., Z.X.G., and S.Y.L. are with Academy for Engineering and Technology, Fudan University, Shanghai 200433, China. W.L. and S.Y.L. are with Research Institute of Intelligent Complex Systems, Fudan University, Shanghai 200433, China. Corresponding e-mails: {ganzhongxue, syleng}@fudan.edu.cn.

S.Y.L. is supported by the National Natural Science Foundation of China (No. 12101133) and Shanghai Sailing Program (No. 21YF1402300). This work is also supported by Shanghai Municipal Science and Technology Major Project (No. 2021SHZDZX0103).

II. Preliminaries and models

A. Problem formulation

Consider a group of n agents, trying to reveal the underlying true state of nature θ^* from a finite set of hypotheses $\Theta = \{\theta_1, \theta_2, \cdots, \theta_m\}$. At each discrete time step $t = 1, 2, \cdots$, each agent i obtains an observation $s_{i,t}$ of an environmental random process, where $\boldsymbol{s}_t = (s_{1,t}, s_{2,t}, \cdots, s_{n,t})^\top$ is generated according to likelihood function $\boldsymbol{\ell}(\cdot|\theta^*)$. The corresponding random variable of agent i's observation at time t is denoted as $\boldsymbol{S}_{i,t}$ and $\boldsymbol{S}_t = (\boldsymbol{S}_{1,t}, \cdots, \boldsymbol{S}_{n,t})^\top$. Here each $\boldsymbol{S}_{i,t}$ has its individual observation space S_i and is i.i.d. with respect to t.

The signal structure of agent i for state θ is described by a probability distribution $\ell_i(\cdot|\theta)$. In these settings $\ell_i(s_{i,t}|\theta)$ characterizes the probability that signal $s_{i,t}$ can be observed by agent i at time t when he/she believes θ is the true state. It is required that $\ell_i(\cdot|\theta^*)$ coincides with the i-th marginal distribution of $\boldsymbol{\ell}(\cdot|\theta^*)$, which thus also describes the probability distribution of $\boldsymbol{S}_{i,t}$.

The agents interact in a networked fashion [26]–[28], 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 *belief* of agent i at time t is denoted as $\mu_{i,t}$, which is a probability distribution over the set of states Θ, i.e., $\sum_{k=1}^{m} \mu_{i,t}(\theta_k) = 1, \forall i = 1, \cdots, n, \forall t = 0, 1, \cdots$. Here $\mu_{i,0}$ represents the *initial belief* of agent i.

B. Parametric Social Learning

Traditional social learning usually consists of two steps at each time for agents to update their beliefs, i.e., the Bayesian update step and the step of aggregating neighbors' beliefs. Taking different orders of these two steps leads to the two basic social learning strategies, called *LoAB* (Logarithmic Aggregation and Bayesian update) and *BLoA* (Bayesian update and Logarithmic Aggregation) respectively [24].

In the Bayesian update step, at time $t+1$, every agent i combines its prior belief $\mu_{i,t}$ with observation $s_{i,t+1}$ from environment to form its posterior belief $\tilde{\mu}_{i,t+1}$. This process is usually described as solving an optimization problem:

$$\tilde{\mu}_{i,t+1} = \underset{f\in\mathbb{P}(\Theta)}{\arg\min}\{D_{\mathrm{KL}}(f \parallel \mu_{i,t}) - \sum_{\theta\in\Theta} f(\theta)\log(\ell_i(s_{i,t+1}|\theta))\}, \tag{1}$$

where $D_{\mathrm{KL}}(\cdot \parallel \cdot)$ is the Kullback-Leibler divergence (KL-divergence) between two probability distributions. Specifically, $D_{\mathrm{KL}}(f \parallel \mu_{i,t}) = \sum_{\theta\in\Theta} f(\theta)\log\frac{f(\theta)}{\mu_{i,t}(\theta)}$, which can be interpreted as the difference between the prior belief and the posterior belief. The second term on r.h.s. of (1) describes the maximum likelihood estimation given the latest observation $s_{i,t+1}$. In real-world scenarios, different agents may assign different weights to balance the two terms. For example, some people are stubborn or unwilling to accept new information, while in contrast some others are sensitive to external information and easily make changes. Considering the cases, we improve the update rule (1) by introducing an individual parameter $\delta_i \in (0,1)$ for agent i to trade-off the two terms:

$$\begin{aligned}\tilde{\mu}_{i,t+1} = \underset{f\in\mathbb{P}(\Theta)}{\arg\min}\{&\delta_i D_{\mathrm{KL}}(f \parallel \mu_{i,t}) \\ &- (1-\delta_i)\sum_{\theta\in\Theta} f(\theta)\log(\ell_i(s_{i,t+1}|\theta))\}.\end{aligned} \tag{2}$$

Notice that large values of δ_i corresponds to stubborn agent i, whereas small values of δ_i means agent i is sensitive to new information. Specifically, when $\delta_i = 1/2$, equal weight is assigned to both terms, and the case degenerates to traditional social learning (1). We thus name the parameter δ_i as the stubbornness of agent i. By directly solving the optimization problem (2), we obtain the update formula for the posterior belief as follows:

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

where $\gamma_i = 1/\delta_i - 1$. Further combining the improved Bayesian update step and the aggregation step in different orders, we obtain two new learning algorithms, i.e.,

a) PSL-LoAB (Parametric Social Learning-Logarithmic Aggregation and Bayesian update):

$$\mu_{i,t+1}(\theta) = \frac{\prod_{j=1}^{n}\mu_{j,t}^{a_{ij}}(\theta)\ell_i^{\gamma_i}(s_{i,t+1}|\theta)}{\sum_{\theta'\in\Theta}\prod_{j=1}^{n}\mu_{j,t}^{a_{ij}}(\theta')\ell_i^{\gamma_i}(s_{i,t+1}|\theta')}; \tag{3}$$

b) PSL-BLoA (Parametric Social Learning-Bayesian update and Logarithmic Aggregation):

$$\mu_{i,t+1}(\theta) = \frac{\prod_{j=1}^{n}\mu_{j,t}^{a_{ij}}(\theta)\ell_j^{\gamma_j a_{ij}}(s_{j,t+1}|\theta)}{\sum_{\theta'\in\Theta}\prod_{j=1}^{n}\mu_{j,t}^{a_{ij}}(\theta')\ell_j^{\gamma_j a_{ij}}(s_{j,t+1}|\theta')}. \tag{4}$$

We should emphasize that the PSL strategies differ from [23] in which a parameter called step-size is introduced to improve the model's adaptive performance. Moreover, our PSL strategies are more sociologically explicable as their Bayesian update step is derived by directly solving (2), instead of modifying the update rules implicitly.

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 our PSL 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 positive diagonal entries.

Here Assumption 1b) describes the fact that all agents can at least receive information from themselves. Moreover, Assumption 1 guarantees that A is the transition matrix of an irreducible, aperiodic Markov chain of finite states. We recall the following lemma [29]:

Lemma 1: If a Markov chain of 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 $1 \leqslant i, j \leqslant n$.

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

Assumption 2 (Belief and signal structure): For all agents $i = 1, 2, \cdots, n$,

a) they have positive initial beliefs on all states, i.e., $\mu_{i,0}(\theta) > 0$ for all $\theta \in \Theta$;

b) they have positive signal structures, i.e., $\ell_i(s_i|\theta) > 0$ for all $s_i \in S_i$ and $\theta \in \Theta$.

Notice that if the initial belief of agent i on state θ is zero, following our learning algorithms, its belief remains zero all the time. In this circumstance θ is meaningless for agent i, and we thus eliminate the situation by imposing Assumption 2a). For the signal structures and Assumption 2b), the same explanation can be applied.

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

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

Under this assumption, for all $\theta \neq \theta^*$, there exists at least an agent i such that $D_{\mathrm{KL}}(\ell_i(\cdot|\theta^*) \parallel \ell_i(\cdot|\theta))$ is strictly positive.

Denote in the following that $K_i(\theta^*, \theta) = D_{\mathrm{KL}}(\ell_i(\cdot|\theta^*) \parallel \ell_i(\cdot|\theta))$, and define a probability triple $(\Omega, \mathcal{F}, \mathbb{P}^*)$, where $\Omega = \{\omega | \omega = (\boldsymbol{s}_1, \boldsymbol{s}_2, \cdots)\}$, $\mathcal{F}$ is the σ-algebra generated by the observations, and $\mathbb{P}^*$ is the probability measure induced by paths in Ω, i.e., $\mathbb{P}^* = \prod_{t=1}^{\infty} \boldsymbol{\ell}(\cdot|\theta^*)$. We use $\mathbb{E}^*[\cdot]$ to denote the expectation operator associated with measure $\mathbb{P}^*$. Now we can state the main results describing the convergence of the PSL strategies.

Theorem 1: Under Assumptions 1, 2 and 3, the update rules (3) and (4) satisfy the following properties:

$$\lim_{t\to\infty} \frac{1}{t} \log \frac{\mu_{i,t}(\theta)}{\mu_{i,t}(\theta^*)} = -\sum_{j=1}^{n} \pi_j \gamma_j K_j(\theta^*, \theta), \quad \forall \theta \neq \theta^*, \tag{5}$$

and

$$\lim_{t\to\infty} \mu_{i,t}(\theta^*) = 1, \quad \mathbb{P}^*\text{−a.s.}, \quad \forall i = 1, \cdots, n. \tag{6}$$

Proof: We consider the update rule (3) first. For each agent i and $\theta \neq \theta^*$, we have

$$\log \frac{\mu_{i,t+1}(\theta)}{\mu_{i,t+1}(\theta^*)} = \sum_{j=1}^{n} a_{ij} \log \frac{\mu_{j,t}(\theta)}{\mu_{j,t}(\theta^*)} + \gamma_i \log \frac{\ell_i(s_{i,t+1}|\theta)}{\ell_i(s_{i,t+1}|\theta^*)}.$$

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

$$\nu_{i,t+1}(\theta) = \sum_{j=1}^{n} a_{ij} \nu_{j,t}(\theta) + \gamma_i L_{i,t+1}(\theta). \tag{7}$$

Rewrite (7) in matrix form:

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

where $\Gamma = \mathrm{diag}(\gamma_1, \cdots, \gamma_n)$. Now it follows that

$$\begin{aligned} \frac{1}{t}\boldsymbol{\nu}_{t+1}(\theta) &= \frac{1}{t} A\boldsymbol{\nu}_t(\theta) + \frac{1}{t}\Gamma \boldsymbol{L}_{t+1}(\theta) = \cdots \\ &= \frac{1}{t} A^{t+1} \boldsymbol{\nu}_0(\theta) + \frac{1}{t}\sum_{k=1}^{t} A^k \Gamma \boldsymbol{L}_{t+1-k}(\theta) + \frac{1}{t}\Gamma \boldsymbol{L}_{t+1}(\theta). \end{aligned} \tag{8}$$

The assumptions admit that the first and the third terms on r.h.s. of (8) go to zero as $t \to \infty$. The second term can be deformed as

$$\begin{aligned} \frac{1}{t}\sum_{k=1}^{t} A^k \Gamma \boldsymbol{L}_{t+1-k}(\theta) &= \frac{1}{t}\sum_{k=1}^{t} (A^k - \mathbf{1}_n \pi) \Gamma \boldsymbol{L}_{t+1-k}(\theta) \\ &\quad + \frac{1}{t}\sum_{k=1}^{t} \mathbf{1}_n \pi \Gamma (\boldsymbol{L}_{t+1-k}(\theta) + \boldsymbol{K}(\theta^*, \theta)) \\ &\quad - \frac{1}{t}\sum_{k=1}^{t} \mathbf{1}_n \pi \Gamma \boldsymbol{K}(\theta^*, \theta), \end{aligned} \tag{9}$$

where $\mathbf{1}_n$ is an n-dimensional column vector of ones. Lemma 1 admits that $\lim_{k\to\infty} A^k = \mathbf{1}_n \pi$. Noticing that all elements of $A^k (k = 1, 2, \cdots)$ are bounded, the first term on r.h.s. of (9) converges to zero as $t \to \infty$. Moreover,

$$\begin{aligned} \mathbb{E}^*[L_{i,t}(\theta)] &= \mathbb{E}^*[\log \frac{\ell_i(s_{i,t}|\theta)}{\ell_i(s_{i,t}|\theta^*)}] \\ &= \int_{s\in S_i} \ell_i(s|\theta^*) \log \frac{\ell_i(s|\theta)}{\ell_i(s|\theta^*)} \mathrm{d}s \\ &= -D_{\mathrm{KL}}(\ell_i(\cdot|\theta^*) \parallel \ell_i(\cdot|\theta)) = -K_i(\theta^*, \theta). \end{aligned} \tag{10}$$

The Kolmogorov's strong law of large numbers gives that

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

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

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

Now (9) gives that

$$\lim_{t\to\infty}\frac{1}{t}\sum_{k=1}^{t}A^k\Gamma\boldsymbol{L}_{t+1-k}(\theta)=-\mathbf{1}_n\pi\Gamma\boldsymbol{K}(\theta^*,\theta),\quad \mathbb{P}^*\text{−a.s.}. \tag{11}$$

Therefore, property (5) can be directly induced from (8) and (11). It follows that with probability one, for any $\epsilon > 0$, there exists an integer T such that $\forall t > T$ and $\forall\theta \neq \theta^*$,

$$\left|\frac{1}{t}\log\frac{\mu_{i,t}(\theta)}{\mu_{i,t}(\theta^*)}+\sum_{j=1}^{n}\pi_j\gamma_jK_j(\theta^*,\theta)\right|<\epsilon.$$

Noticing that $\sum_{\theta\neq\theta^*}\mu_{i,t}(\theta)=1-\mu_{i,t}(\theta^*)$, we have

$$\frac{1}{1+\sum_{\theta\neq\theta^*}\exp\left(\left(\epsilon-\sum_{j=1}^{n}\pi_j\gamma_jK_j(\theta^*,\theta)\right)t\right)}<\mu_{i,t}(\theta^*)\leqslant 1. \tag{12}$$

Letting $t\to\infty$, property (6) is then proved because of the arbitrariness of ϵ.

For update rule (4), analogous to (7), we have

$$\nu_{i,t+1}(\theta)=\sum_{j=1}^{n}a_{ij}(\nu_{j,t}(\theta)+\gamma_jL_{j,t+1}(\theta)). \tag{13}$$

Rewriting (13) in matrix form and by recursion we have

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

The same analysis gives that (5) and (6) also hold for update rule (4). ∎

Theorem 1 indicates that all agents will eventually learn the underlying true state as long as the assumptions are satisfied. Moreover, the convergence rate is closely related to the stubbornness δ_i of each agent i. Specifically, the smaller stubbornness all agents have, hence larger γ_i, the faster the group's beliefs converge. This result is consistent with our intuition that a group with open-minded agents is more willing to adapt to new environment, leading to faster reach of consensus. We will illustrate these insights through numerical examples in the next section. From Theorem 1, the following result characterizing the convergence rate can also be obtained:

Corollary 1: Under Assumptions 1, 2 and 3, the update rules (3) and (4) satisfy that for all $i=1,2,\cdots,n$ and all $\theta\neq\theta^*$,

$$\lim_{t\to\infty}\mu_{i,t}(\theta)\leqslant\exp(-\alpha_\theta t),\quad \mathbb{P}^*\text{−a.s.},$$

where $\alpha_\theta=\sum_{j=1}^{n}\pi_j\gamma_jK_j(\theta^*,\theta)$.

Consider a special case of $\delta_i = 1$ for an agent i, from (2) it can be interpreted that agent i is completely stubborn or blocked from environment. In this circumstance, $\gamma_i = 0$ and both update rules (3) and (4) become

$$\mu_{i,t+1}(\theta)=\frac{\prod_{j=1}^{n}\mu_{j,t}^{a_{ij}}(\theta)}{\sum_{\theta'\in\Theta}\prod_{j=1}^{n}\mu_{j,t}^{a_{ij}}(\theta')},$$

indicating agent i does not accept any external information at the Bayesian update step. Practically, to minimize the influence on learning performance brought by the stubbornness of some agent, we can characterize its location by solving $\arg\min_{i=1,\cdots,n}\sum_{j\neq k}\pi_iK_i(\theta_j,\theta_k)$. Particularly, if all agents have equivalent information on all states, i.e., $K_i(\theta_j,\theta_k)\equiv K$, for all $i=1,\cdots,n$ and $j\neq k$, then the optimal location is the one with lowest eigenvector centrality (or remotely located in the network). If all agents have identical eigenvector centrality, i.e., $\pi_i\equiv\pi^*$ for all $i=1,\cdots,n$, the optimal one should be the least informative agent. These results will be demonstrated in numerical examples in Section IV.

The following example shows that practically it is not always better to accelerate the convergence. Too quick convergence will result in false learning of the true state, due to limited communication precision in real scenarios.

Example 1: Consider a group of two agents, and assume $\Theta=\{\theta_1,\theta_2\}$ with θ_1 being the underlying true state θ^*. Both agents have positive initial beliefs over Θ, i.e., $\mu_{i,0}(\theta)>0,\forall i=1,2$ and $\forall\theta\in\Theta$. Signals are assumed to be generated at each time from sets $S_1=S_2=\{\mathrm{H},\mathrm{T}\}$, and according to the probability distribution of $\ell(\mathrm{H}|\theta^*)=0.8$ and $\ell(\mathrm{T}|\theta^*)=0.2$. Moreover, signal structures are set as $\ell_i(\mathrm{H}|\theta_1)=\ell_i(\mathrm{T}|\theta_2)=0.8$, $\ell_i(\mathrm{H}|\theta_2)=\ell_i(\mathrm{T}|\theta_1)=0.2$, for $i=1,2$. The two agents can receive information from each other, thus the network is strongly-connected and has all positive elements in the weight matrix. PSL-BLoA algorithm is performed on the group. Firstly, the Bayesian update step admits that for $i=1,2$, the private beliefs on true state θ_1 at time step 1 satisfy:

$$\tilde{\mu}_{i,1}(\theta_1)=\frac{\mu_{i,0}(\theta_1)\ell_i^{\gamma_i}(s_{i,1}|\theta_1)}{\mu_{i,0}(\theta_1)\ell_i^{\gamma_i}(s_{i,1}|\theta_1)+\mu_{i,0}(\theta_2)\ell_i^{\gamma_i}(s_{i,1}|\theta_2)}.$$

If at time 1 signal T is generated, with probability 0.2, then we have

$$\tilde{\mu}_{i,1}(\theta_1)=\frac{\mu_{i,0}(\theta_1)}{\mu_{i,0}(\theta_1)+(1-\mu_{i,0}(\theta_1))4^{\gamma_i}}.$$

Here we consider a special but common in reality case that the network is subject to some communication constraint, i.e., the communication between the agents is quantized. In this circumstance, the messages sent by the agents are reduced to $[D\tilde{\mu}_{i,1}(\theta)]$, where $D\in\mathbb{Z}^+$ is predefined and

$$[x]=\begin{cases}\lfloor x\rfloor, & \text{if } x\leqslant\lfloor x\rfloor+0.5,\\ \lfloor x\rfloor+1, & \text{if } x>\lfloor x\rfloor+0.5.\end{cases}$$

Therefore, if the stubbornness of any agent i is small enough that $\gamma_i>\log_4\frac{(2D-1)\mu_{i,0}(\theta_1)}{1-\mu_{i,0}(\theta_1)}$, then $\tilde{\mu}_{i,1}(\theta_1)<\frac{1}{2D}$ and $[D\tilde{\mu}_{i,1}(\theta_1)]=0$, which leads to $\mu_{i,1}(\theta_1)=0$ for $i=1,2$ after aggregation step. In this case, no agent in the network can learn the true state θ_1. This schematic example illustrates that small stubbornness, hence fast convergence, does not always lead to good learning performance, signifying the importance of the capability to tune the convergence rate.

We state in the following the other main theoretical result of this work:

Theorem 2: Under Assumptions 1, 2 and 3, following the update rules (3) and (4), for all $i = 1, \cdots, n$, $\theta \neq \theta^*$, and all $t \geqslant 0$, we have

$$\mathbb{E}^*[\nu_{i,t+1}(\theta)] \leqslant \beta_\theta - (t+1)\alpha_\theta,$$

where

$$\beta_\theta = \frac{4\gamma^* \log n ||\boldsymbol{K}(\theta^*, \theta)||_\infty}{1 - \lambda_{\max}(A)} + \max_{1 \leqslant i \leqslant n} \log \frac{\mu_{i,0}(\theta)}{\mu_{i,0}(\theta^*)},$$

$\gamma^* = \max_{1 \leqslant i \leqslant n} \gamma_i$, $||\boldsymbol{K}(\theta^*, \theta)||_\infty = \max_{1 \leqslant i \leqslant n} K_i(\theta^*, \theta)$, $\lambda_{\max}(A) = \max_{\lambda_i(A) \neq 1} \{|\lambda_i(A)|\}$, $\lambda_i(A)$ denotes one of A's eigenvalues, and $\lambda_0(A) = 1$. α_θ is as defined in Corollary 1.

The next lemma [17] is required to prove Theorem 2:

Lemma 2: Let the strong connectivity of network hold, and define $\lambda_{\max}(\cdot)$ as in Theorem 2. Then for any t, the stochastic matrix A satisfies

$$\sum_{k=1}^{t} \sum_{j=1}^{n} |[A^k]_{ij} - \pi_j| \leqslant \frac{4 \log n}{1 - \lambda_{\max}(A)}.$$

Proof: We only consider update rule (3) here and proof for the other one is similar. From (9) and (12), it follows that

$$\mathbb{E}^*[\boldsymbol{\nu}_{t+1}(\theta)] = A\mathbb{E}^*[\boldsymbol{\nu}_t(\theta)] - \Gamma \boldsymbol{K}(\theta^*, \theta).$$

By recursion we have

$$\begin{aligned}\mathbb{E}^*[\boldsymbol{\nu}_{t+1}(\theta)] =& A^{t+1}\boldsymbol{\nu}_0(\theta) + \sum_{k=0}^{t} (\mathbf{1}_n \pi - A^k)\Gamma \boldsymbol{K}(\theta^*, \theta) \\ & - \sum_{k=0}^{t} \mathbf{1}_n \pi \Gamma \boldsymbol{K}(\theta^*, \theta).\end{aligned}$$

Now from Lemma 2, we can obtain

$$\begin{aligned}\mathbb{E}^*[\boldsymbol{\nu}_{t+1}(\theta)] \leqslant & ||\boldsymbol{\nu}_0||_\infty \mathbf{1}_n + \frac{4\gamma^* \log n ||\boldsymbol{K}(\theta^*, \theta)||_\infty}{1 - \lambda_{\max}(A)} \mathbf{1}_n \\ & - (t+1)\mathbf{1}_n \pi \Gamma \boldsymbol{K}(\theta^*, \theta).\end{aligned}$$

■

In social learning literatures, $\nu_{i,t+1}(\theta)$ is referred to as the rejection extent of wrong state θ in favor of θ^* for agent i at time $t+1$. Theorem 2 implies that the expectation of the rejection extent of wrong states can be bounded by a linear function of time t. Hence the belief of each agent on any wrong hypothesis will eventually decay exponentially, while there will be a transient due to the existence of β_θ. The transient is influenced by the inhomogeneity of initial beliefs, the stubbornness of the group, and the mixing properties of the graph. It is interesting that increasing γ will raise the value of β_θ, thus prolonging the transient time.

Corollary 2: Assumptions 1, 2 and 3 hold, and further assume that all agents have bounded signal structures, i.e., $\exists 0 < l < L$, such that $l < \ell_i(s_i|\theta) < L$ for all $s_i \in S_i$ and $\theta \in \Theta$. Let $\rho \in (0,1)$ be a given confidence level. Update rules (3) and (4) have the following property: there exists an integer $N(\rho)$ such that with probability $1 - \rho$, for all $t \geqslant N(\rho)$ there holds that for any $\theta \neq \theta^*$,

$$\mu_{i,t}(\theta) \leqslant \exp(-\frac{\alpha^*}{2} t + \beta^*),$$

where $N(\rho) = \frac{8 \log^2 l \log \frac{1}{\rho}}{(\beta^*)^2} + 1$, $\alpha^* = \min_{\theta \neq \theta^*} \alpha_\theta$, $\beta^* = \max_{\theta \neq \theta^*} \beta_\theta$.

The proof of Corollary 2 can be referred to [18]. Corollary 2 provides the theoretical upper bound of beliefs on wrong states under given confidence level and as t is large enough.

IV. Numerical Examples

A. Tuning convergence rate

We first demonstrate the capability of the PSL strategies in tuning the convergence rate.

Example 2: Consider a strongly-connected digraph consisting of 5 agents, whose structure is shown in Fig. 1(a). The weight of each directed edge is randomly generated from $(0,1)$ and then normalized to satisfy the weight matrix A is row-stochastic. We assume there are four possible states $\Theta = \{\theta_1, \theta_2, \theta_3, \theta_4\}$ and θ_1 is set as the true state. The initial beliefs are also uniformly generated from interval $(0,1)$ and subject to $\sum_{j=1}^{4} \mu_{i,0}(\theta_j) = 1, \forall i = 1, \cdots, 5$. At each time step t, a signal s_t is randomly generated, following the normal distribution $\mathcal{N}(\mu_1, 1)$ and observed by all agents. As θ_1 is the underlying true state, from assumptions we have $\ell_i(\cdot|\theta_1) = \mathcal{N}(\mu_1, 1), \forall i = 1, \cdots, 5$. The likelihood functions of other states are randomly assigned as $\mathcal{N}(\mu_1, 1)$ or $\mathcal{N}(\mu_2, 1)$ with $\mu_1 \neq \mu_2$, while $\theta^* = \theta_1$ is globally identifiable.

Here all agents share the same value of stubbornness, i.e., $\delta_i = \delta$, hence $\gamma_i = \gamma$, for $i = 1, \cdots, n$. We focus on the number of iterations for update rules (3) and (4) to collectively learn the true state, where successful learning is considered to be reached if $\sum_{i=1}^{5} |\mu_{i,t}(\theta^*) - 1| \leqslant 10^{-3}$. According to (12), we define the theoretical asymptotic curve for the convergence as $L(t) = 1/(1 + \sum_{\theta \neq \theta^*} \exp(-\sum_{j=1}^{n} \pi_j \gamma_j K_j(\theta^*, \theta) t))$.

Results in Fig. 2 validate the results in Theorem 1, which further show that larger value of γ (smaller δ) leads to faster convergence. Thus PSL strategies indeed enable fine tuning of the convergence rate through the introduced parameter δ.

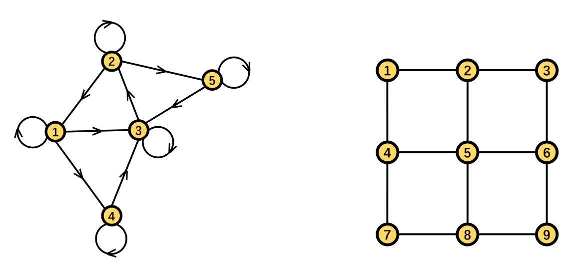

(a) (b)

Fig. 1. Two network structures used in this work.

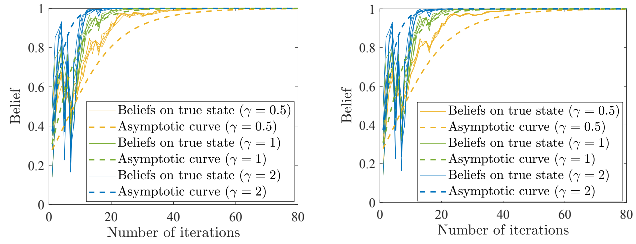

(a) (b)

Fig. 2. The evolution of beliefs on true state θ^* for 5 agents interacted in structure shown in Fig. 1(a) with different stubbornness. (a) and (b) are results for update rules (3) and (4) respectively, where successful learning are both reached at iteration number of 22 ($\gamma = 2$), 42 ($\gamma = 1$) and 72 ($\gamma = 0.5$).

B. Learning with completely stubborn agent

The PSL strategies prevail over existing methods on aspect of its capability in learning with agents of diverse stubbornness. Here we provide examples showing interesting sociological phenomena how social influence and information content of agents affect the group's learning rate.

Example 3: Let $\Theta = \{\theta_1, \theta_2, \theta_3, \theta_4\}$ and $\theta^* = \theta_1$. Here we consider 9 agents interacted in a 3×3 grid network [Fig. 1(b)] with the weights determined by

$$a_{ij} = \begin{cases} \frac{1}{|N_i|}, & \text{if } j \in N_i, \\ 0, & \text{otherwise,} \end{cases}$$

where N_i denotes agent i's neighboring set. For the defined weight matrix, the eigenvector centrality of agent i, which reflects its social influence, is in proportion to $|N_i|$ (the number of neighbors). The initial beliefs are uniformly generated from interval $(0, 1)$ and subject to $\sum_{j=1}^{4} \mu_{i,0}(\theta_j) = 1, \forall i = 1, \cdots, 9$.

We denote agents that have access to external information ($\gamma > 0$) as *open agents*, and assume there is only one open agent in the network and the others are completely stubborn. Furthermore, all agents in the network are assumed to be equivalently informative to the true state, with $\ell_i(\cdot|\theta_k) = \mathcal{N}(k, 1), i = 1, \cdots, 9$. At each time t, a signal vector $\boldsymbol{s}_t$ is randomly generated, following the Gaussian distribution $\boldsymbol{\ell}(\cdot|\theta^*) = \boldsymbol{\mathcal{N}}(\mathbf{1}_9, I_{9\times 9})$ and observed by all agents.

By assigning the open agent in the corner (1, 3, 7, or 9), margin (2, 4, 6, or 8), and centre (5), it has relatively small, middle, and large eigenvector centrality respectively. Fig. 3 shows the convergence of agent 9, following update rules (3) and (4) and demonstrating that the more social influence the open agent is of, the faster the group can learn the true state.

Example 4: Consider the same network structure as in Example 3 and let $\Theta = \{\theta_1, \theta_2\}$, $\theta^* = \theta_1$. In this example we fix the only open agent at corner, hence with same eigenvector centrality. At each time step t, a signal vector $\boldsymbol{s}_t$ is randomly generated, following the Gaussian distribution $\boldsymbol{\ell}(\cdot|\theta^*) = \boldsymbol{\mathcal{N}}(\mathbf{1}_9, I_{9\times 9})$ and observed by the agents. For the signal structures, we set $\ell_i(\cdot|\theta_2) = \mathcal{N}(i+1, 1), i \in \{1, 3, 7, 9\}$, and $\ell_i(\cdot|\theta_2) = \mathcal{N}(1, 1), i \in \{2, 4, 5, 6, 8\}$, leading to that $K_9(\theta^*, \theta) > K_7(\theta^*, \theta) > K_3(\theta^*, \theta) > K_1(\theta^*, \theta) > 0$ and $K_j(\theta^*, \theta) = 0, j \in \{2, 4, 5, 6, 8\}$. We denote agent 9 as the most informative agent, which is the most helpful to distinguish between true and wrong states. Fig. 4 illustrates that open agent being more informative contributes to higher convergence rate, validating our analysis in Section III.

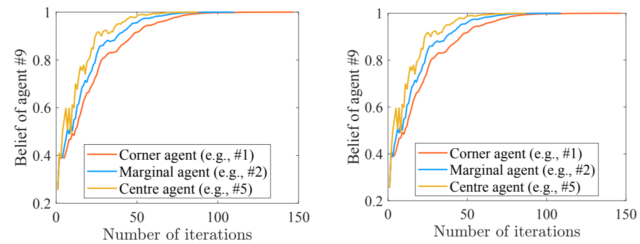

(a) (b)

Fig. 3. The evolution of beliefs on true state θ^* for agent 9 with the only open agent locating at different positions. (a) and (b) are results for update rules (3) and (4) respectively. Notice that when the open agent is located at centre or corner of the grid, the convergence rate is fastest or slowest respectively.

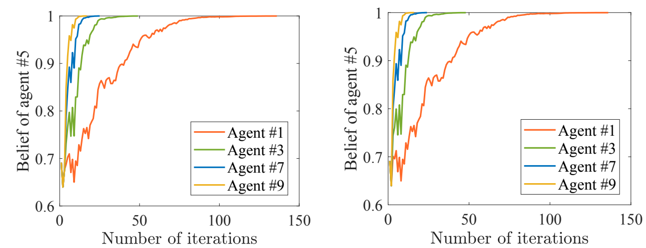

(a) (b)

Fig. 4. The evolution of beliefs on true state θ^* for agent 5 with the only open agent being different extent of informative. (a) and (b) are results for update rules (3) and (4) respectively. Notice that when the open agent is more informative, the convergence rate is faster.

V. Conclusion and Future Work

To conclude, we have proposed a non-Bayesian social learning strategy by introducing a stubbornness parameter balancing the agents' internal belief and external information. The update rules have been directly derived from an optimization problem and their convergence have been analytically estimated. Essentiality of tuning convergence rate by adjusting the stubbornness parameters has been demonstrated. Several sociological insights have been revealed both analytically and numerically during the analyses of learning with completely stubborn agents.

Our work can assist network designers to realize adjustable distributed estimation tasks and minimize losses when some agents encounter sudden incapacitation of perception. Because of its sociological nature, the PSL strategy has potential theoretical and applied values. Future work includes designing learning models with time delay, and considering partial stubbornness on certain states or time-varying stubbornness.

REFERENCES

[1] F. Fatas-Villafranca, D. Saura, and F. J. Vázquez, "A dynamic model of public opinion formation," *Journal of Public Economic Theory*, vol. 13, no. 3, pp. 417–441, 2011.

[2] A. Patania, G. Petri, and F. Vaccarino, "The shape of collaborations," *EPJ Data Science*, vol. 6, pp. 1–16, 2017.

[3] X. Qiu, L. Yang, C. Guan, and S. Leng, "Closed-loop control of higher-order complex networks: Finite-time and pinning strategies," *Chaos, Solitons & Fractals*, vol. 173, p. 113677, 2023.

[4] A. V. Banerjee, "A simple model of herd behavior," *The Quarterly Journal of Economics*, vol. 107, no. 3, pp. 797–817, 1992.

[5] S. Bikhchandani, D. Hirshleifer, and I. Welch, "A theory of fads, fashion, custom, and cultural change as informational cascades," *Journal of Political Economy*, vol. 100, no. 5, pp. 992–1026, 1992.

[6] L. Smith and P. Sørensen, "Pathological outcomes of observational learning," *Econometrica*, vol. 68, no. 2, pp. 371–398, 2000.

[7] D. Gale and S. Kariv, "Bayesian learning in social networks," *Games and Economic Behavior*, vol. 45, no. 2, pp. 329–346, 2003.

[8] V. Bordignon, S. Vlaski, V. Matta, and A. H. Sayed, "Learning from heterogeneous data based on social interactions over graphs," *IEEE Transactions on Information Theory*, vol. 69, no. 5, pp. 3347–3371, 2022.

[9] V. Bordignon, V. Matta, and A. H. Sayed, "Partial information sharing over social learning networks," *IEEE Transactions on Information Theory*, vol. 69, no. 3, pp. 2033–2058, 2022.

[10] 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.

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

[12] 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.

[13] P. Molavi, A. Tahbaz-Salehi, and A. Jadbabaie, "Foundations of non-bayesian social learning," *Columbia Business School Research Paper*, no. 15-95, 2017.

[14] P. Molavi, K. R. Rad, A. Tahbaz-Salehi, and A. Jadbabaie, "On consensus and exponentially fast social learning," in *2012 American Control Conference (ACC)*. IEEE, 2012, pp. 2165–2170.

[15] K. R. Rad and A. Tahbaz-Salehi, "Distributed parameter estimation in networks," in *49th IEEE Conference on Decision and Control (CDC)*. IEEE, 2010, pp. 5050–5055.

[16] S. Kar, J. M. 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.

[17] S. Shahrampour, A. Rakhlin, and A. Jadbabaie, "Distributed detection: Finite-time analysis and impact of network topology," *IEEE Transactions on Automatic Control*, vol. 61, no. 11, pp. 3256–3268, 2015.

[18] A. Nedić, A. Olshevsky, and C. A. Uribe, "Nonasymptotic convergence rates for cooperative learning over time-varying directed graphs," in *2015 American Control Conference (ACC)*. IEEE, 2015, pp. 5884–5889.

[19] ——, "Fast convergence rates for distributed non-bayesian learning," *IEEE Transactions on Automatic Control*, vol. 62, no. 11, pp. 5538–5553, 2017.

[20] 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.

[21] 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.

[22] 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.

[23] V. Bordignon, V. Matta, and A. H. Sayed, "Adaptive social learning," *IEEE Transactions on Information Theory*, vol. 67, no. 9, pp. 6053–6081, 2021.

[24] 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.

[25] Y. Ha, Y. Guo, and W. Lin, "Non-bayesian social learning model with periodically switching structures," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 31, no. 4, p. 043137, 2021.

[26] S. Boccaletti, V. Latora, Y. Moreno, M. Chavez, and D.-U. Hwang, "Complex networks: Structure and dynamics," *Physics Reports*, vol. 424, no. 4-5, pp. 175–308, 2006.

[27] S. Leng, Z. Xu, and H. Ma, "Reconstructing directional causal networks with random forest: Causality meeting machine learning," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 29, no. 9, 2019.

[28] L. Yang, W. Lin, and S. Leng, "Conditional cross-map-based technique: From pairwise dynamical causality to causal network reconstruction," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 33, no. 6, 2023.

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