THE ROLE OF MEMORY IN SOCIAL LEARNING WHEN SHARING PARTIAL OPINIONS

Michele Cirillo[⋆] *Virginia Bordignon*[†] *Vincenzo Matta*[⋆] *Ali H. Sayed*[†]

[⋆] DIEM, University of Salerno, Fisciano (SA), Italy
[†] EPFL, School of Engineering, CH-1015 Lausanne, Switzerland

ABSTRACT

In social learning, a group of agents linked by a graph topology collect data and exchange opinions on some topic of interest, represented by a finite set of hypotheses. Traditional social learning algorithms allow all agents in the network to gain full confidence on the true underlying hypothesis as the number of observations increases. Under partial information sharing, agents can exchange opinions *only on a single hypothesis*. This introduces significant challenges as compared to the standard case of full opinion sharing. We propose a novel strategy where each agent forms a valid belief by completing the partial beliefs received from its neighbors. The completion process exploits the knowledge accumulated in the past beliefs, thanks to a principled *memory-aware* rule inspired by a Bayesian criterion. We provide a detailed characterization of the memory-aware strategy, which reveals novel learning dynamics and highlights its advantages over previously considered schemes.

***Index Terms*—** Social learning, Bayesian update, information diffusion, partial information.

1. INTRODUCTION AND RELATED WORK

Social learning refers to a family of distributed inferential strategies through which agents propagate their opinions over a network [1–11]. All agents collect *streaming data* related to a phenomenon of common interest, which is represented by a discrete-valued state, the *hypothesis* $\theta \in \Theta$. Each agent summarizes its opinions regarding these hypotheses into a probability vector (the *belief*). Under a fully-Bayesian approach, the beliefs should be iteratively updated over time by blending previous beliefs and the likelihood of the new incoming data through a Bayesian update. However, fully Bayesian approaches are usually not viable in social learning because a joint model encompassing all agents in the network is seldom available, and even when it is, the complexity in implementing the associated updates is a formidable task [12,13]. These challenges motivated the introduction of *non-Bayesian* social learning, which basically consists of a two-step algorithm: *i)* a Bayesian update performed *locally* by each agent using its *private* likelihood model and observations; and *ii)* a *cooperation* step where each agent combines the beliefs received from its immediate neighbors [3,4,8,14–18].

A common assumption in social learning is that each agent has access to the *entire* belief vector of its neighbors. However, this condition is not verified in several contexts, as it is often the case that agents share opinions on a *single* candidate state. For example, consider the problem of selecting one commercial product among brands $\{\theta_1, \theta_2, \theta_3\}$. A new product has been recently released by brand θ_1, which motivates the social group to exchange reviews regarding *only* the new release and ignore the remaining hypotheses. At the end of the learning process, the individuals end up with some updated convictions as regards the overall collection of brands $\{\theta_1, \theta_2, \theta_3\}$. From an engineering-design perspective, partial information sharing is also motivated by the requirement of devising multi-agent communication schemes that are parsimonious in terms of communication resources, or systems where sharing full information is forbidden or restricted due to privacy concerns [19,20].

We can distinguish two major approaches to deal with social learning under limited exchange of information. The first approach prescribes sharing the full beliefs, albeit by imposing some communication constraints. Under this category we find the work [21], where quantized belief ratios are exchanged. In comparison, an event-triggered algorithm is proposed in [22], where the communication burden is reduced by transmitting only at instants when the beliefs are deemed sufficiently innovative. Another strategy to reduce the communication load is to exchange beliefs with only one randomly-sampled neighbor at a time [23].

The second approach, known as *partial information sharing*, was introduced in [24,25], and prescribes that agents can only transmit their beliefs concerning a *single* hypothesis of interest, denoted by $\theta_{\rm TX}$. A critical part of this approach is that agents need to fill in the missing information regarding the non-transmitted components. The maximum-entropy choice was proposed in [24,25], where the non-transmitted belief components have equal mass. Since this filling strategy ignores past information, it is a *memoryless* strategy.

In the present work we propose a novel *filling strategy* where, instead of performing a blind allocation, each agent exploits its most updated *local* knowledge to fill in the missing entries in its neighbors' beliefs. Contrasted with the *memoryless* allocation, this alternative approach leads to a *memory-aware* strategy. The main contribution of this work is to provide a detailed characterization of the learning mechanism of the proposed strategy. The analysis reveals novel dynamics arising from partial sharing of opinions, highlighting a nontrivial interplay between memory and cooperation.

2. PARTIAL INFORMATION SHARING

We consider a network of K agents, cooperating to learn the true state of nature $\theta_0 \in \Theta \triangleq \{1, 2, \ldots, H\}$. Each agent k at instant i collects an observation $\boldsymbol{\xi}_{k,i} \in \mathcal{X}_k$ (we use bold font for random variables) and possesses a set of *private* models $L_k(\xi|\theta)$ for the probability (density or mass) function of $\xi \in \mathcal{X}_k$ given $\theta \in \Theta$. The observations $\boldsymbol{\xi}_{k,i}$ are distributed according to the true model $L_k(\xi|\theta_0)$. They are independent and identically distributed over time, but they can be dependent across agents.

The network is represented by a graph, whose nodes correspond to the agents, and whose edges correspond to communication links between the agents. The symbol $\mathcal{N}_k$ denotes the set of neighbors of agent k (including k itself). Each agent k associates to each agent

This work was supported in part by grant 205121-184999 from the Swiss National Science Foundation (SNSF).

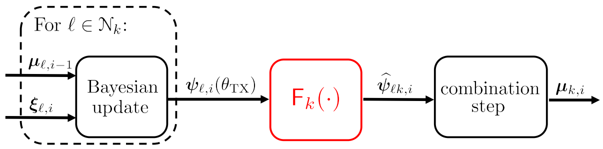

Fig. 1. Social learning under partial information sharing.

ℓ a nonnegative weight $a_{\ell k}$. These weights can be conveniently arranged into a left-stochastic matrix $A = [a_{\ell k}]$ satisfying:

$$\mathbb{1}^\top A = \mathbb{1}^\top, \quad a_{\ell k} \geq 0, \quad a_{\ell k} = 0 \text{ if } \ell \notin \mathcal{N}_k. \tag{1}$$

The main output of social learning is the *belief vector* $\boldsymbol{\mu}_{k,i} \in \Delta_H$, where Δ_H denotes the probability simplex with dimension H. The θ-th component $\boldsymbol{\mu}_{k,i}(\theta)$ quantifies how confident agent k is at instant i that θ is the true state of nature. We assume throughout the work that the initial belief vectors are uniform, i.e., $\mu_{k,0}(\theta) = 1/H$ for all $\theta \in \Theta$ and all $k = 1, 2, \dots, K$. In *traditional* social learning, each agent k iteratively updates its belief vector $\boldsymbol{\mu}_{k,i}$ using the following two-step procedure:

$$\boldsymbol{\psi}_{k,i}(\theta) \propto L_k(\boldsymbol{\xi}_{k,i}|\theta)\boldsymbol{\mu}_{k,i-1}(\theta), \tag{2a}$$

$$\boldsymbol{\mu}_{k,i}(\theta) \propto \prod_{\ell \in \mathcal{N}_k} [\boldsymbol{\psi}_{\ell,i}(\theta)]^{a_{\ell k}}. \tag{2b}$$

As is standard in Bayesian theory, the proportionality sign hides the constant (independent of θ) that is necessary to obtain a probability vector. Step (2a) outputs an *intermediate belief vector* $\boldsymbol{\psi}_{k,i}$ through a local Bayesian update with prior $\boldsymbol{\mu}_{k,i-1}$ and likelihood $L_k(\boldsymbol{\xi}_{k,i}|\theta)$. The intermediate beliefs are then shared over the network and, in step (2b), each agent k combines the intermediate beliefs in its neighborhood $\mathcal{N}_k$ using a weighted geometric average with combination weights $a_{\ell k}$. It is known that with strategy (2a)-(2b) all agents learn the truth almost-surely as $i \to \infty$, that is, $\boldsymbol{\mu}_{k,i}(\theta_0) \xrightarrow{\text{a.s.}} 1$ [14, 17, 18].

In traditional social learning, each agent k receives from each neighbor $\ell \in \mathcal{N}_k$ the *complete* belief vector $\boldsymbol{\psi}_{\ell,i}$. We here consider instead the challenging scenario of *partial information sharing*, where agents share a *single* component of their belief vector corresponding to a *hypothesis of interest* $\theta_{\rm TX} \in \Theta$. The new learning algorithm consists of the following three steps — see Fig. 1:

$$\boldsymbol{\psi}_{k,i}(\theta) \propto L_k(\boldsymbol{\xi}_{k,i}|\theta)\boldsymbol{\mu}_{k,i-1}(\theta), \tag{3a}$$

$$\widehat{\boldsymbol{\psi}}_{\ell k,i} = \mathsf{F}_k\big(\boldsymbol{\psi}_{\ell,i}(\theta_{\rm TX})\big), \text{ for } \ell \in \mathcal{N}_k \tag{3b}$$

$$\boldsymbol{\mu}_{k,i}(\theta) \propto \prod_{\ell \in \mathcal{N}_k} [\widehat{\boldsymbol{\psi}}_{\ell k,i}(\theta)]^{a_{\ell k}}. \tag{3c}$$

Steps (3a) and (3c) are still a Bayesian update and a geometric averaging rule, respectively. The novelty is in step (3b), where agent k, upon receiving the *single* belief component $\boldsymbol{\psi}_{\ell,i}(\theta_{\rm TX})$ from a neighboring agent ℓ, builds a complete *estimated* belief vector $\widehat{\boldsymbol{\psi}}_{\ell k,i}$. This is the estimated belief vector constructed by agent k *relative* to its neighbor ℓ, in the sense that, starting from the received component $\boldsymbol{\psi}_{\ell,i}(\theta_{\rm TX})$, agent k fills in the missing entries according to a *filling strategy* $\mathsf{F}_k : \mathbb{R} \mapsto \Delta_H$. The mapping F_k will be allowed to be *random* since it can depend on the beliefs possessed by agent k.

The choice of the filling strategy F_k is critical to achieve correct learning. We now show how it can be derived from a Bayesian approach. Assume that in the construction of the estimated belief $\widehat{\boldsymbol{\psi}}_{\ell k,i}$, agent k trusts agent ℓ and, hence, it sets:

$$\widehat{\boldsymbol{\psi}}_{\ell k,i}(\theta_{\rm TX}) = \boldsymbol{\psi}_{\ell,i}(\theta_{\rm TX}). \tag{4}$$

Once (4) is enforced, the residual mass assigned to the set $\mathcal{T} = \{\theta \neq \theta_{\rm TX}\}$ is $1 - \boldsymbol{\psi}_{\ell,i}(\theta_{\rm TX})$. From Bayes' rule, this implies that the belief *conditioned* on set $\mathcal{T}$, say $\boldsymbol{b}_k(\theta|\mathcal{T})$, must fulfill the condition:

$$\widehat{\boldsymbol{\psi}}_{\ell k,i}(\theta) = \boldsymbol{b}_k(\theta|\mathcal{T})(1 - \boldsymbol{\psi}_{\ell,i}(\theta_{\rm TX})), \qquad \text{for all } \theta \neq \theta_{\rm TX}. \tag{5}$$

To complete the filling strategy, it is necessary to choose the form of $\boldsymbol{b}_k(\theta|\mathcal{T})$. In [25], the maximum-entropy choice was proposed:

$$\boldsymbol{b}_k(\theta|\mathcal{T}) = \frac{1}{H-1}. \tag{6}$$

In this work, we propose a novel strategy that exploits more fully the knowledge that agent k has accumulated up to time i. To this end, we consider the most updated belief available at agent k at time i, namely, $\boldsymbol{\psi}_{k,i}$, which leads then to the following conditional belief:

$$\boldsymbol{b}_k(\theta|\mathcal{T}) = \frac{\boldsymbol{\psi}_{k,i}(\theta)}{1 - \boldsymbol{\psi}_{k,i}(\theta_{\rm TX})}. \tag{7}$$

While strategy (6) forgets any evidence from the past and performs a uniform allocation, strategy (7) diversifies the allocation based on the available knowledge stored in the intermediate belief vector $\boldsymbol{\psi}_{k,i}$. We shall accordingly refer to (6) as the *memoryless* strategy, and to (7) as the *memory-aware* strategy.

Note also that strategy (7) automatically ensures *self-awareness*, meaning that for $\ell = k$ we get $\widehat{\boldsymbol{\psi}}_{kk,i} = \boldsymbol{\psi}_{k,i}$. It is interesting to remark that this compelling property arose naturally from our Bayesian interpretation of the filling strategy, once we allowed it to incorporate the information contained in $\boldsymbol{\psi}_{k,i}$. In comparison, strategy (6) does not provide self-awareness.

3. ASSUMPTIONS

For ease of reference, we collect in this section the assumptions involved in our analysis, which are common in the study of social learning methods. The first assumption is a standard condition of network connectivity.

Assumption 1 (**Strongly-Connected Network**). *There exists at least one path linking every two nodes in both directions and at least one self-loop, i.e., $a_{kk} > 0$ for some agent k.* □

Since A is left-stochastic, strong connectivity implies that A is a primitive matrix [26, 27]. Then, the Perron-Frobenius theorem implies the existence of a vector $v = [v_k]$, a.k.a. the Perron eigenvector, which satisfies the following conditions [26, 27]:

$$Av = v, \quad \mathbb{1}^\top v = 1, \quad v_k > 0 \ \forall k. \tag{8}$$

Next we list two regularity conditions on the likelihoods.

Assumption 2 (**Finiteness of KL Divergences**). *For each $k = 1, 2, \dots, K$ and each pair θ, θ': $D_{\rm KL}(L_{k,\theta_0}||L_{k,\theta}) < \infty$, where $D_{\rm KL}(L_{k,\theta_0}||L_{k,\theta})$ denotes the Kullback-Leibler divergence between $L_k(\xi|\theta_0)$ and $L_k(\xi|\theta)$ [28].* □

The next assumption excludes some redundancy from the likelihood models. It rules out the case where at some agent the likelihood of the true hypothesis is equal to a convex combination of the likelihoods of the distinguishable hypotheses. Preliminarily, we need to introduce the set of hypotheses *distinguishable for agent k*.

$$\mathcal{D}_k \triangleq \{\theta \in \Theta \setminus \{\theta_0\} : D_{\rm KL}(L_{k,\theta_0}||L_{k,\theta}) > 0\}. \tag{9}$$

Assumption 3 (**Regular Likelihoods**). *Let $\mathcal{D}_k \neq \emptyset$. Consider a convex combination vector $\alpha = [\alpha(\theta)]_{\theta\in\mathcal{D}_k}$ and introduce the convex combination:*

$$f_{k,\alpha}(\xi) \triangleq \sum_{\theta\in\mathcal{D}_k} \alpha(\theta) L_k(\xi|\theta). \tag{10}$$

We assume that $\min_{\alpha\in\Delta_{|\mathcal{D}_k|}} D_{\mathrm{KL}}(L_{k,\theta_0}||f_{k,\alpha}) > 0.$ ☐

We continue by introducing the set of hypotheses *indistinguishable for agent* k:

$$\mathcal{I}_k \triangleq \{\theta \in \Theta \setminus \{\theta_0\} : D_{\mathrm{KL}}(L_{k,\theta_0}\|L_{k,\theta}) = 0\}. \tag{11}$$

Note that if agent k learns in isolation, hypothesis θ_0 is identifiable only if $\mathcal{I}_k = \emptyset$, otherwise θ_0 would be indistinguishable from any hypothesis belonging to $\mathcal{I}_k$. However, in social learning it is often sufficient to relax this *local* identifiability notion and to resort instead to the following condition of identifiability at the *network* level.

Assumption 4 (**Global Identifiability**). *For each $\theta \neq \theta_0$, there exists at least one agent k with $\theta \notin \mathcal{I}_k$, namely, $\bigcap_{k=1}^{K} \mathcal{I}_k = \emptyset$.* ☐

Under global identifiability, if agent k is able to distinguish one hypothesis θ from the true one (i.e., if $\theta \in \mathcal{D}_k$), then the entire network can benefit from this ability, including some other agent ℓ for which $\theta \in \mathcal{I}_\ell$. A special case of global identifiability is when there exists one powerful agent which is able to distinguish all the false hypotheses. We refer to this case as *strong-agent identifiability*.

Assumption 5 (**Strong-Agent Identifiability**). *There exists one agent k with $\mathcal{I}_k = \emptyset$.* ☐

Note that "strong-agent identifiability" is significantly milder than local identifiability, since it does not require identifiability at *all* agents, but just at *one* powerful agent.

4. MAIN THEOREMS

In the following we use the notation, for any $\mathcal{S} \subseteq \Theta$:

$$\boldsymbol{\mu}_{k,i}(\mathcal{S}) \triangleq \sum_{\theta\in\mathcal{S}} \boldsymbol{\mu}_{k,i}(\theta), \tag{12}$$

with the convention that $\boldsymbol{\mu}_{k,i}(\emptyset) = 0$. Moreover, the cardinalities of the indistinguishable sets $\mathcal{I}_k$ will be denoted by J_k, and we introduce the *network average* of the individual cardinalities:

$$J \triangleq \prod_{k=1}^{K} J_k^{v_k}, \tag{13}$$

that is, a weighted geometric average of the individual $\{J_k\}$, with weights given by the entries $\{v_k\}$ of the Perron eigenvector.

Our main result is collected in the next theorems, whose proofs, omitted for space constraints, are available in [29].

Theorem 1 (**Convergence in the case $\theta_{\mathrm{TX}} \neq \theta_0$**). *Under Assumptions 1–4, we have the following properties for each $k = 1, 2, \dots, K$:*

- *Belief of the transmitted hypothesis:*

$$\boldsymbol{\mu}_{k,i}(\theta_{\mathrm{TX}}) \xrightarrow{\text{a.s.}} 0. \tag{14}$$

- *Beliefs of the non-transmitted distinguishable hypotheses $\theta \in \mathcal{D}_k \setminus \{\theta_{\mathrm{TX}}\}$:*

$$\boldsymbol{\mu}_{k,i}(\theta) \xrightarrow{\text{a.s.}} 0. \tag{15}$$

- *Beliefs of the true hypothesis and of the non-transmitted indistinguishable hypotheses $\theta \in \{\theta_0\} \cup \big(\mathcal{I}_k \setminus \{\theta_{\mathrm{TX}}\}\big)$:*

$$\boldsymbol{\mu}_{k,i}(\theta) \xrightarrow{\text{a.s.}} \frac{1}{1 + |\mathcal{I}_k \setminus \{\theta_{\mathrm{TX}}\}|}, \tag{16}$$

i.e., the mass is equipartitioned over the set comprising the true and the non-transmitted indistinguishable hypotheses.

Theorem 2 (**Convergence in the case $\theta_{\mathrm{TX}} = \theta_0$**). *Under Assumptions 1–4, we have the following properties for each $k = 1, 2, \dots, K$:*

- *Belief of the true hypothesis:*

$$\boldsymbol{\mu}_{k,i}(\theta_0) \xrightarrow{\text{a.s.}} \frac{1}{1 + J}. \tag{17}$$

- *Beliefs of the distinguishable hypotheses $\theta \in \mathcal{D}_k$:*

$$\boldsymbol{\mu}_{k,i}(\theta) \xrightarrow{\text{a.s.}} 0. \tag{18}$$

- *Beliefs of the indistinguishable hypotheses $\theta \in \mathcal{I}_k$:*

$$\boldsymbol{\mu}_{k,i}(\theta) \xrightarrow{\text{a.s.}} \frac{J}{J_k}\,\frac{1}{1 + J}. \tag{19}$$

Theorems 1 and 2 capture the fundamental learning mechanism of the memory-aware strategy. Ideally, we want a learning strategy that places negligible confidence on θ_{TX} when $\theta_{\mathrm{TX}} \neq \theta_0$, and full confidence otherwise [25]. As a matter of fact, Theorem 1 guarantees that all agents are able to learn well when $\theta_{\mathrm{TX}} \neq \theta_0$, since they end up placing zero mass on the (false) shared hypothesis. Let us switch to the case $\theta_{\mathrm{TX}} = \theta_0$, which is covered by Theorem 2. Equation (17) reveals that the belief of the true hypothesis converges to $1/(1 + J)$. From (13) we know that J is a weighted geometric average of the single-agent cardinalities $\{J_k\}$. As such, J is zero even if *a single agent* k has $\mathcal{I}_k = \emptyset$, i.e., if Assumption 5 holds. When $J = 0$, from (17) we have $\boldsymbol{\mu}_{k,i}(\theta_{\mathrm{TX}}) \xrightarrow{\text{a.s.}} 1$, and we conclude that *truth learning is achievable under strong-agent identifiability*.

It remains to examine the memory-aware strategy when strong-agent identifiability does not hold. To this end, observe that the geometric average of a group of numbers is bounded by the minimum and maximum values in the group. Thus, in the network we will have two subsets of agents: one subset of agents with cardinality J_k greater than the average cardinality J, and another subset of agents with $J_k \leq J$. Consider first the case $J_k > J$. From (17) and (19) we see that, almost surely for large i:

$$\boldsymbol{\mu}_{k,i}(\theta_0) > \boldsymbol{\mu}_{k,i}(\theta), \qquad \text{for all } \theta \in \Theta \setminus \{\theta_0\}. \tag{20}$$

Conversely, agents with $J_k \leq J$ end up with:

$$\boldsymbol{\mu}_{k,i}(\theta_0) \leq \boldsymbol{\mu}_{k,i}(\theta), \qquad \text{for all } \theta \in \Theta \setminus \{\theta_0\}. \tag{21}$$

Examining jointly (20) and (21), we discover an interesting and perhaps unexpected implication. After cooperation, the agents that were *individually more confused* (larger J_k) truly benefit from cooperation, ending up with a belief that is maximized at the true hypothesis. The situation is reversed for the agents that were *individually less confused* (smaller J_k) but then end up with a belief that is no longer maximized at the true hypothesis (or, if it is, there are multiple indistinguishable maxima). In the next section we show how this issue can be resolved since we prove that, when strong-agent identifiability is violated, it is still possible to implement a decision rule that leads to correct learning.

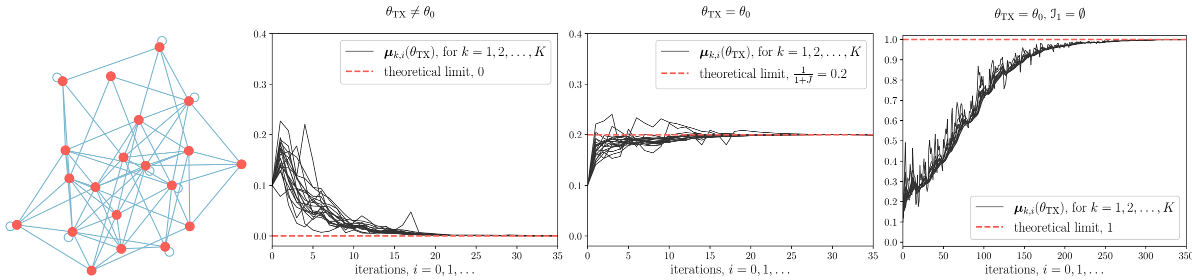

Fig. 2. *First plot.* Network graph. *Second and third plots.* The algorithm is run with indistinguishable sets $\{\mathcal{I}_k\}$ provided in (28). Therefore, only global (i.e., not strong-agent) identifiability holds. *Fourth plot.* Strong-agent identifiability is enforced by letting $\mathcal{I}_1 = \emptyset$.

4.1. A Consistent Decision Rule

Consider the following threshold test, for some small ε:

$$\begin{cases} \boldsymbol{\mu}_{k,i}(\theta_{\rm TX}) > (1/H) - \varepsilon \Rightarrow \text{accept } \theta_{\rm TX} \\ \boldsymbol{\mu}_{k,i}(\theta_{\rm TX}) \leq (1/H) - \varepsilon \Rightarrow \text{reject } \theta_{\rm TX} \end{cases} \tag{22}$$

Since $J_k \leq H - 1$, which implies that $J \leq H - 1$, we have:

$$0 < \underbrace{\frac{1}{H} - \varepsilon}_{\text{threshold}} < \frac{1}{1+J}. \tag{23}$$

From Theorems 1 and 2 we know that $\boldsymbol{\mu}_{k,i}(\theta_{\rm TX})$ converges to zero if $\theta_{\rm TX} \neq \theta_0$, and to $1/(1+J)$ otherwise. Using (23), we conclude that the decision rule (22) guarantees always correct learning as $i \to \infty$.

The rationale behind (22) is that the transmitted hypothesis is accepted if the observed belief exceeds (but for a small ϵ) the uniform value $1/H$ corresponding to the initial belief assignment. Note that this rule is *not* based on the maximization of the belief. In contrast, the rule accepts $\theta_{\rm TX}$ provided the minimal requirement of exceeding the average probability, even if there could be other indistinguishable hypotheses featuring a higher belief. In other words, the decision rule (22) is biased in favor of $\theta_{\rm TX}$, because the agents are conscious of the limitations of partial information sharing, and they try to overcome them by adjusting the decision threshold.

4.2. Comparison Against the Memoryless Strategy

It was shown in [25] that with the memoryless filling strategy (6), the social learning problem can be reinterpreted as a binary detection problem involving the comparison between the likelihood of the true hypothesis, $L_k(\xi|\theta_0)$, and a fictitious distribution:

$$f_k(\xi) = \frac{1}{H-1} \sum_{\tau \neq \theta_{\rm TX}} L_k(\xi|\tau). \tag{24}$$

Then, the following KL divergences were introduced:

$$d_{\rm TX} \triangleq \sum_{k=1}^{K} v_k \, D_{\rm KL}(L_{k,\theta_0} \| L_{k,\theta_{\rm TX}}), \tag{25}$$

$$d_{\rm fict} \triangleq \sum_{k=1}^{K} v_k \, D_{\rm KL}(L_{k,\theta_0} \| f_k). \tag{26}$$

Consider first the case $\theta_{\rm TX} \neq \theta_0$. We know that in this case the memory-aware strategy *always* learns well. According to [25, Theorem 3], the memoryless strategy learns correctly if $d_{\rm fict} < d_{\rm TX}$, which means that, to declare that $\theta_{\rm TX}$ is false, the true distribution must be more similar to the fictitious distribution than to the distribution of the shared hypothesis. Notably, when $d_{\rm fict} > d_{\rm TX}$, the memoryless strategy incurs *mislearning*, i.e., it is completely fooled and places full mass on the wrong hypothesis [25].

Conversely, in the case $\theta_{\rm TX} = \theta_0$, the memoryless strategy always learns [25], while the memory-aware strategy requires strong-agent identifiability. However, the memory-aware strategy does not display the mislearning behavior since it places a nonzero mass on the true hypothesis, and thanks to this property, we have shown that it can be driven to the right decision by means of (22).

5. SIMULATION RESULTS

We consider the network graph with $K = 20$ agents shown in Fig. 2, equipped with a Metropolis combination matrix [27]:

$$a_{\ell k} = \begin{cases} 1/\max\{|\mathcal{N}_\ell|, |\mathcal{N}_k|\}, & \ell \in \mathcal{N}_k \setminus \{k\} \\ a_{\ell k} = 0, & \ell \notin \mathcal{N}_k, \end{cases} \tag{27}$$

with $a_{kk} = 1 - \sum_{\ell \in \mathcal{N}_k} a_{\ell k}$. Note that the Perron eigenvector is uniform since A is doubly stochastic. We use this combination matrix in the social learning algorithm (3a)–(3c) with filling strategy (7).

The learning problem has $H = 10$ hypotheses, and $L_k(\xi|\theta)$ is Gaussian with unit variance and mean equal to θ. The true hypothesis is $\theta_0 = 1$. The sets of indistinguishable hypotheses have the following cardinalities:

$$J_k = \begin{cases} 4, & \text{for } k \in [1, 10], \\ 8, & \text{for } k \in [11, 15], \\ 2, & \text{for } k \in [16, 20], \end{cases} \quad \text{and} \quad \bigcap_{k=1}^{K} \mathcal{I}_k = \emptyset. \tag{28}$$

Note that (28) implies that strong-agent identifiability is not satisfied, since $\mathcal{I}_k \neq \emptyset$ for any k, whereas global identifiability is satisfied. In view of (13), we have $J = (4^{10} \cdot 8^5 \cdot 2^5)^{1/20} = 4$.

In the second plot of Fig. 2, we illustrate the dynamics of the agents' belief relative to the transmitted hypothesis, for the case $\theta_{\rm TX} \neq \theta_0$. We see that, according to (14), the beliefs all converge to 0. In the third plot we consider instead $\theta_{\rm TX} = \theta_0$. We see that, as predicted by (17), the beliefs now converge to $1/(1 + J) = 1/5$. Finally, we consider a variation of the previous experiment where we enforce strong-agent identifiability at agent 1. We run the social learning algorithm for $\theta_{\rm TX} = \theta_0$ and, as shown in the fourth plot of Fig. 2, the beliefs of all agents converge to 1, as it must be.

6. REFERENCES

[1] D. Acemoglu and A. Ozdaglar, "Opinion dynamics and learning in social networks," *Dynamic Games and Applications*, vol. 1, no. 1, pp. 3–49, 2011.

[2] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar, "Bayesian learning in social networks," *The Review of Economic Studies*, vol. 78, no. 4, pp. 1201–1236, 2011.

[3] 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.

[4] X. Zhao and A. H. Sayed, "Learning over social networks via diffusion adaptation," in *Proc. Asilomar Conference on Signals, Systems, and Computers (ACSSC)*, Pacific Grove, CA, USA, November 2012, pp. 709–713.

[5] C. Chamley, A. Scaglione, and L. Li, "Models for the diffusion of beliefs in social networks: An overview," *IEEE Signal Processing Magazine*, vol. 30, no. 3, pp. 16–29, 2013.

[6] V. Krishnamurthy and H. V. Poor, "Social learning and Bayesian games in multiagent signal processing: How do local and global decision makers interact?," *IEEE Signal Processing Magazine*, vol. 30, no. 3, pp. 43–57, 2013.

[7] E. Mossel and O. Tamuz, "Opinion exchange dynamics," *Probability Surveys*, vol. 14, pp. 155–204, 2017.

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

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

[10] S. Bikhchandani, D. Hirshleifer, and I. Welch, "Learning from the behavior of others: Conformity, fads, and informational cascades," *Journal of Economic Perspectives*, vol. 12, no. 3, pp. 151–170, 1998.

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

[12] E. Mossel and O. Tamuz, "Making consensus tractable," *ACM Transactions on Economics and Computation*, vol. 1, no. 4, pp. 1–19, 2013.

[13] J. Hązła, A. Jadbabaie, E. Mossel, and M. A. Rahimian, "Bayesian decision making in groups is hard," *Operations Research*, vol. 69, no. 2, pp. 632–654, 2021.

[14] 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.

[15] P. Molavi, A. Tahbaz-Salehi, and A. Jadbabaie, "A theory of non-Bayesian social learning," *Econometrica*, vol. 86, no. 2, pp. 445–490, 2018.

[16] 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.

[17] 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.

[18] V. Matta, V. Bordignon, A. Santos, and A. H. Sayed, "Interplay between topology and social learning over weak graphs," *IEEE Open Journal of Signal Processing*, vol. 1, pp. 99–119, 2020.

[19] A. Koloskova, S. U. Stich, and M. Jaggi, "Decentralized stochastic optimization and gossip algorithms with compressed communication," in *Proc. International Conference on Machine Learning (ICML)*, Long Beach, CA, USA, June 2019, pp. 3478–3487.

[20] J. Mills, J. Hu, and G. Min, "Communication-efficient federated learning for wireless edge intelligence in IoT," *IEEE Internet of Things Journal*, vol. 7, no. 7, pp. 5986–5994, 2020.

[21] M. T. Toghani and C. A. Uribe, "Communication-efficient distributed cooperative learning with compressed beliefs," *IEEE Transactions on Control of Network Systems*, vol. 9, no. 3, pp. 1215–1226, 2022.

[22] A. Mitra, S. Bagchi, and S. Sundaram, "Event-triggered distributed inference," in *Proc. Conference on Decision and Control (CDC)*, Jeju Island, South Korea, December 2020, pp. 6228–6233.

[23] Y. Inan, M. Kayaalp, E. Telatar, and A. H. Sayed, "Social learning under randomized collaborations," in *Proc. IEEE International Symposium on Information Theory (ISIT)*, Helsinki, Finland, June 2022, pp. 115–120.

[24] V. Bordignon, 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.

[25] V. Bordignon, V. Matta, and Ali H. Sayed, "Partial information sharing over social learning networks," *IEEE Transactions on Information Theory*, vol. 69, no. 3, pp. 2033–2058, 2023.

[26] R. A. Horn and C. R. Johnson, *Matrix Analysis*, Cambridge University Press, 2012.

[27] A. H. Sayed, "Adaptation, learning, and optimization over networks," *Foundations and Trends in Machine Learning*, vol. 7, no. 4-5, pp. 311–801, 2014.

[28] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, John Wiley & Sons, 1991.

[29] M. Cirillo, V. Bordignon, V. Matta, and Ali H. Sayed, "Memory-aware social learning under partial information sharing," *under review, available online at arXiv:2301.10688v1*, 2023.