

Adaptation in Online Social Learning

Virginia Bordignon
School of Engineering, EPFL
virginia.bordignon@epfl.ch

Vincenzo Matta
DIEM, University of Salerno
vmatta@unisa.it

Ali H. Sayed
School of Engineering, EPFL
ali.sayed@epfl.ch

Abstract—This work studies social learning under non-stationary conditions. Although designed for online inference, traditional social learning algorithms perform poorly under drifting conditions. To mitigate this drawback, we propose the Adaptive Social Learning (ASL) strategy. This strategy leverages an *adaptive Bayesian update*, where the adaptation degree can be modulated by tuning a suitable step-size parameter. The learning performance of the ASL algorithm is examined by means of a steady-state analysis. It is shown that, under the regime of small step-sizes: *i) consistent learning is possible; ii) and an accurate prediction of the performance can be furnished in terms of a Gaussian approximation.*

Index Terms—Social learning, Bayesian update, adaptive learning, diffusion strategy.

I. INTRODUCTION

In social learning strategies, a set of communicating agents seeks to update their beliefs about the state of nature in response to streaming information [1]–[4]. In most existing methods in the literature, as time evolves, the beliefs of the agents tend to concentrate on the true state [5]–[12], often at an exponentially fast rate of convergence. However, such remarkable convergence properties have the collateral effect of hindering adaptation.

Let us consider the following example. A network of 10 agents aims to solve a weather forecast problem using an online social learning algorithm. At each instant, these agents collect data originating from one among three possible hypotheses: “sunny”, “cloudy”, “rainy”. At first, data are consistent with the hypothesis “sunny”, but then from instant $i = 200$ onwards they indicate that the correct forecast is “rainy”. As we see in Fig. 1 for the behavior of agent 1, the social learning algorithm reacts with a considerable inertia to the drift in the hypothesis.

Fig. 1 shows that the agent learns well until instant $i = 200$, whereas from $i = 200$ onwards, the situation changes dramatically: the traditional social learning algorithm has a delayed reaction. First, agents perceive a change only at $i \approx 350$, but start opting for the wrong hypothesis “cloudy”. Then, after a prohibitive number of iterations, at $i \approx 550$, agents manage to overcome their stubbornness and opt for the correct hypothesis “rainy”. To tackle this problem, this work proposes an Adaptive Social Learning (ASL) strategy, whose performance is shown in the second column of the same Fig. 1 for the same example. We see that the ASL algorithm manages to track the target change at instant $i \approx$

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

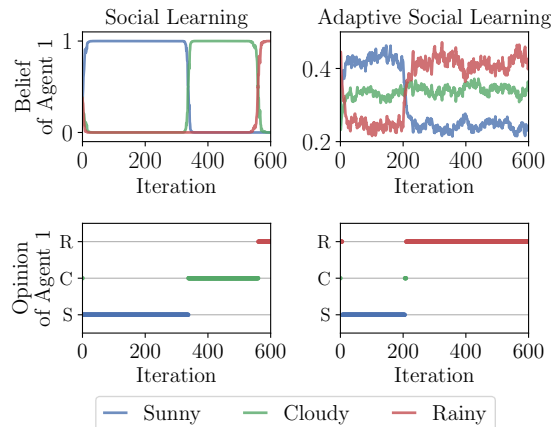


Fig. 1. *Traditional social learning vs. adaptive social learning. Top panels:* Belief evolution of agent 1, with θ_0 changing at time $i = 200$. *Bottom panels:* The instantaneous decision taken by agent 1 by choosing the hypothesis that maximizes the current belief.

200, exhibiting an adaptation capacity that is remarkably higher than that of the traditional social learning algorithm.

The main contributions of this work can be summarized as follows. First, we introduce a novel social learning strategy that enables adaptation. Then, by exploiting recent advances in the field of distributed detection over adaptive networks [13], we provide an accurate analytical characterization of this strategy in terms of *i) convergence of the system at steady state (Theorem 1); ii) achievability of consistent learning (Theorem 2); iii) and a Gaussian approximation for the learning performance (Theorem 3).* Due to space constraints, proofs are omitted.

II. ASL STRATEGY

Consider a strongly-connected network of N agents trying to infer the true state of nature $\theta_0 \in \Theta$ given a set of H hypotheses, $\Theta = \{1, 2, \dots, H\}$. Each agent k , at time i , observes streaming data $\xi_{k,i}$, belonging to a certain space \mathcal{X}_k , drawn from a distribution that depends on the underlying hypothesis θ_0 . The data are assumed to be independent over time, i.e., across index i , whereas they can be dependent across agents. Moreover, it is assumed that the distribution of $\xi_{k,i}$ belongs to a set of H admissible models (likelihood functions) that are identified by the hypotheses $\theta \in \Theta = \{1, 2, \dots, H\}$. The likelihood of agent k evaluated at θ is denoted by $L_k(\xi|\theta)$ with $\xi \in \mathcal{X}_k$. Note that the likelihoods are allowed to vary across the agents.

We model the network using a strongly-connected graph, with a left-stochastic combination matrix $A \triangleq [a_{\ell k}]$. Element $a_{\ell k}$ weights information received by agent k from agent ℓ : $a_{\ell k}$ is a non-negative real number and it is equal to zero if $\ell \notin \mathcal{N}_k$, where \mathcal{N}_k is the neighborhood of agent k (k included). We define the Perron eigenvector π such that [14]:

$$A\pi = \pi, \quad \mathbb{1}^\top \pi = 1, \quad \pi \succ 0. \quad (1)$$

Agents will incorporate the information contained in their local observations and diffuse it across the network by iteratively updating and exchanging their belief vectors $\boldsymbol{\mu}_{k,i}$. The belief vector is a probability vector over the set of hypotheses Θ and each component $\mu_{k,i}(\theta)$ reflects the confidence of agent k at instant i that θ is the true hypothesis.

Assumption 1 (Positive initial beliefs). *All agents start with a strictly positive belief for all hypotheses, i.e., $\boldsymbol{\mu}_{k,0}(\theta) > 0$ for each agent k and all $\theta \in \Theta$.* \square

A. Adaptive Social Learning (ASL) Algorithm

In the adaptive scenario, system conditions can change over time, e.g., the true state of nature or the network topology might change. To address that setup we introduce the ASL strategy, which can be described in terms of the following iterative two-step algorithm. In the first step, each agent k constructs an *intermediate* belief vector $\boldsymbol{\psi}_{k,i}$ by incorporating the current observation $\boldsymbol{\xi}_{k,i}$ into the belief of the preceding time epoch, $\boldsymbol{\mu}_{k,i-1}$, through the following *adaptive Bayesian update*:

$$\boldsymbol{\psi}_{k,i}(\theta) = \frac{\boldsymbol{\mu}_{k,i-1}^{1-\delta}(\theta) L_k^\delta(\boldsymbol{\xi}_{k,i}|\theta)}{\sum_{\theta' \in \Theta} \boldsymbol{\mu}_{k,i-1}^{1-\delta}(\theta') L_k^\delta(\boldsymbol{\xi}_{k,i}|\theta')} \quad (2)$$

where $0 < \delta < 1$ is a parameter that will be referred to as the *step-size*. In the second step, each agent k aggregates all intermediate beliefs received from its neighbors into its updated belief vector $\boldsymbol{\mu}_{k,i}$ as

$$\boldsymbol{\mu}_{k,i}(\theta) = \frac{\exp\left\{\sum_{\ell \in \mathcal{N}_k} a_{\ell k} \log \boldsymbol{\psi}_{\ell,i}(\theta)\right\}}{\sum_{\theta' \in \Theta} \exp\left\{\sum_{\ell \in \mathcal{N}_k} a_{\ell k} \log \boldsymbol{\psi}_{\ell,i}(\theta')\right\}}. \quad (3)$$

Different than traditional social learning methods employed in [9]–[12], we see that each agent performs the first step by modulating, through the convex weights $1 - \delta$ and δ , the relative weight assigned to the past and new information. In particular, relatively large values of δ give more importance to the new data, whereas small values of δ give more importance to the past beliefs. A similar form of convex combination appeared in the statistical literature for defining the Chernoff information [15].

As usual in the theory of adaptation and learning, the learning performance is characterized in the steady-state regime [14]. In steady state, the true hypothesis θ_0 is kept constant over time, yielding:

$$\boldsymbol{\xi}_{k,i} \sim L_k(\boldsymbol{\xi}|\theta_0), \quad k = 1, 2, \dots, N, \quad i = 1, 2, \dots \quad (4)$$

and the data $\{\boldsymbol{\xi}_{k,i}\}$ are independent and identically distributed (i.i.d.) over time.

Let us define the log-likelihood ratio as

$$\boldsymbol{x}_{k,i}(\theta) \triangleq \log \left(\frac{L_k(\boldsymbol{\xi}_{k,i}|\theta_0)}{L_k(\boldsymbol{\xi}_{k,i}|\theta)} \right). \quad (5)$$

Assumption 2 (Finiteness of KL divergences). *For each $k = 1, 2, \dots, N$ and $\theta \neq \theta_0$:*

$$d_k(\theta) \triangleq \mathbb{E}[\boldsymbol{x}_{k,i}(\theta)] < \infty. \quad (6)$$

\square

To motivate cooperation among agents, we introduce the following identifiability assumption, which implies that the inference problem need not be locally identifiable.

Assumption 3 (Global identifiability). *For each wrong hypothesis $\theta \neq \theta_0$, there exists at least one agent k_θ that has strictly positive KL divergence, $d_{k_\theta}(\theta) > 0$.* \square

In order to characterize the learning performance, it is useful to introduce the logarithm of the ratio between the belief evaluated at θ_0 and the belief evaluated at $\theta \neq \theta_0$:

$$\boldsymbol{\lambda}_{k,i}^{(\delta)}(\theta) \triangleq \log \left(\frac{\boldsymbol{\mu}_{k,i}(\theta_0)}{\boldsymbol{\mu}_{k,i}(\theta)} \right). \quad (7)$$

When we omit the argument θ and write $\boldsymbol{\lambda}_{k,i}^{(\delta)}$, we are referring to the $(H - 1) \times 1$ vector concatenating the log-belief ratios $\lambda_{k,i}^{(\delta)}(\theta)$ for $\theta \neq \theta_0 \in \Theta$. When we omit the subscript i we are referring to a random variable characterized at the *steady state*, i.e., as $i \rightarrow \infty$. Thus, $\boldsymbol{\lambda}_k^{(\delta)}(\theta)$ and $\boldsymbol{\lambda}_k^{(\delta)}$ are, respectively, the *steady-state* log-belief ratio evaluated at θ , and the *steady-state* vector of log-belief ratios.

Remark 1 (Positive beliefs). *In view of Assumption 1 and the ASL algorithm (2)–(3), we see that the belief $\boldsymbol{\mu}_{k,i}(\theta)$ remains nonzero for any θ across time. So, under stationary conditions, at any instant i_0 , the belief vector $\boldsymbol{\mu}_{k,i_0}$ fulfills Assumption 1. This property allows performing the steady-state analysis from $i = 0$ without losing generality. It is also relevant to avoid ill-defined log-belief ratios.* \square

For each i , we can define the instantaneous decision of agent k as corresponding to the hypothesis that maximizes the belief, which leads to the following error probability:

$$p_{k,i}^{(\delta)} = \mathbb{P} \left(\arg \max_{\theta \in \Theta} \boldsymbol{\mu}_{k,i}(\theta) \neq \theta_0 \right) \xrightarrow{i \rightarrow \infty} p_k^{(\delta)}, \quad (8)$$

where $p_k^{(\delta)}$ is the *steady-state* error probability.¹

B. Network Average of Log-Likelihood Ratios

First, a useful concept to introduce is the *network average* of log-likelihood ratios and its expectation, for all $\theta \neq \theta_0$:

$$\boldsymbol{x}_{\text{ave},i}(\theta) = \sum_{\ell=1}^N \pi_\ell \boldsymbol{x}_{\ell,i}(\theta), \quad (9)$$

$$\boldsymbol{m}_{\text{ave}}(\theta) \triangleq \mathbb{E}[\boldsymbol{x}_{\text{ave},i}(\theta)] = \sum_{\ell=1}^N \pi_\ell d_\ell(\theta). \quad (10)$$

¹The existence of the limit in (8) relies on the convergence proved in Theorem 1 (details omitted for space constraints).

Second, if the log-likelihoods have finite variances², we can compute the covariance between $\mathbf{x}_{k,i}(\theta)$ and $\mathbf{x}_{k,i}(\theta')$ as

$$\rho_\ell(\theta, \theta') = \mathbb{E} \left[\left(\mathbf{x}_{\ell,i}(\theta) - d_\ell(\theta) \right) \left(\mathbf{x}_{\ell,i}(\theta') - d_\ell(\theta') \right) \right]. \quad (11)$$

Finally, if data are *independent across the agents*, the covariance between variables $\mathbf{x}_{\text{ave}}(\theta)$ and $\mathbf{x}_{\text{ave}}(\theta')$ is given as

$$\mathbf{c}_{\text{ave}}(\theta, \theta') \triangleq \sum_{\ell=1}^N \pi_\ell^2 \rho_\ell(\theta, \theta'). \quad (12)$$

III. STEADY-STATE ANALYSIS

As seen in the introductory example, in the *adaptive* setting the belief will not converge (in the almost-sure sense) as $i \rightarrow \infty$. On the contrary, the belief of each agent will exhibit an *oscillatory* behavior: a feature that results from adaptation. As we will see, because of this asymptotic *random* character, steady-state analysis is not trivial. The analysis of this random behavior is established in Theorem 1.

Before stating the theorem, let us examine the evolution of the log-belief ratios. Manipulating (2) and (3) in the log domain, for every $\theta \neq \theta_0$ we have:

$$\lambda_{k,i}^{(\delta)}(\theta) = (1-\delta) \sum_{\ell \in \mathcal{N}_k} a_{\ell k} \lambda_{\ell,i-1}^{(\delta)}(\theta) + \delta \sum_{\ell \in \mathcal{N}_k} a_{\ell k} \mathbf{x}_{\ell,i-1}(\theta). \quad (13)$$

The recursion in (13) is in the form of a *diffusion* algorithm with *step-size* δ — see, e.g., [14]. Developing the recursion in (13) we can write, for all $\theta \neq \theta_0$:

$$\begin{aligned} \lambda_{k,i}^{(\delta)}(\theta) &= (1-\delta)^i \sum_{\ell=1}^N [A^\ell]_{\ell k} \lambda_{k,0}(\theta) \\ &+ \delta \sum_{m=0}^{i-1} \sum_{\ell=1}^N (1-\delta)^m [A^{m+1}]_{\ell k} \mathbf{x}_{\ell,i-m}(\theta). \end{aligned} \quad (14)$$

Since the first term on the RHS of (14) vanishes as $i \rightarrow \infty$, for the steady-state analysis we can rewrite with slight abuse of notation:

$$\lambda_{k,i}^{(\delta)}(\theta) = \delta \sum_{m=0}^{i-1} \sum_{\ell=1}^N (1-\delta)^m [A^{m+1}]_{\ell k} \mathbf{x}_{\ell,i-m}(\theta). \quad (15)$$

Theorem 1 (Stability of log-belief ratios). *Let*

$$\lambda_k^{(\delta)}(\theta) = \sum_{\ell=1}^N \delta \sum_{m=0}^{\infty} (1-\delta)^m [A^{m+1}]_{\ell k} \mathbf{x}_{\ell,m+1}(\theta), \quad (16)$$

where the ordering of the summations in (16) means that the N inner series are all almost-surely convergent. Then, under Assumptions 1 and 2 we have that:

$$\lambda_{k,i}^{(\delta)} \overset{i \rightarrow \infty}{\rightsquigarrow} \lambda_k^{(\delta)} \quad (17)$$

where \rightsquigarrow indicates convergence in distribution. ■

Theorem 1 shows that, as long as the first moment of $\mathbf{x}_{k,i}$ exists, the statistical distribution of $\lambda_{k,i}^{(\delta)}$ converges to

²Remarkably, the existence of second moments is not required in Theorems 1 and 2, and is used only in Theorem 3.

the distribution of a stable (i.e., well-defined) random vector $\lambda_k^{(\delta)}$ as $i \rightarrow \infty$. We remark that this does not imply that the partial sum in (15) will converge almost surely to (16) as $i \rightarrow \infty$. The subtlety here is that while

$$\tilde{\lambda}_{k,i}^{(\delta)}(\theta) \triangleq \delta \sum_{m=0}^{i-1} \sum_{\ell=1}^N (1-\delta)^m [A^{m+1}]_{\ell k} \mathbf{x}_{\ell,m+1}(\theta) \quad (18)$$

is almost-surely convergent (which can be deduced from part 1) of Lemma 1 in the appendix, the summation in (15) is not, due to the reversed ordering of the summands.

IV. SMALL- δ ANALYSIS

We will proceed with the asymptotic analysis of $\lambda_k^{(\delta)}$ now in the regime of small δ . As seen in [13], to deal with this asymptotic behavior in the adaptation context, we first introduce a steady-state vector $\lambda_k^{(\delta)}$ that already embodies the effect of summing an infinite number of terms. Only then, we proceed to characterize the asymptotic behavior of the steady-state random vector $\lambda_k^{(\delta)}$ as δ goes to zero. To support the results that follow, we will rely on Lemma 1, which can be found enunciated in Appendix A.

A. Consistent Social Learning

Theorem 2 (Consistency of ASL). *Under Assumptions 1 and 2, we have the following convergence:*

$$\lambda_k^{(\delta)} \xrightarrow{\delta \rightarrow 0} \mathbf{m}_{\text{ave}} \quad \text{in probability.} \quad (19)$$

Since under Assumption 3 all entries of \mathbf{m}_{ave} are strictly positive, Eq. (19) implies that for all $\theta \neq \theta_0$:

$$\lim_{\delta \rightarrow 0} p_k^{(\delta)} = 0 \quad (20)$$

i.e., each agent learns the truth as $\delta \rightarrow 0$. ■

The result of Theorem 2 relies on the weak law of small step-sizes proved in Lemma 1, part 3). This result requires the existence of the first moments $d_\ell(\theta)$, which is guaranteed by Assumption 2. Moreover, it requires that $\mathbf{m}_{\text{ave}}(\theta) > 0$ for all $\theta \neq \theta_0$, which is ensured by Assumption 3 and the strict-positivity of the Perron eigenvector.

B. Normal Approximation for Small δ

Let us examine the behavior of the first two moments of the log-belief ratios. In view of Lemma 1, part 2), we conclude that the expectation of the steady-state random vector $\lambda_k^{(\delta)}$ can be expressed as:

$$\mathbf{m}_k^{(\delta)}(\theta) \triangleq \mathbb{E} \left[\lambda_k^{(\delta)}(\theta) \right] = \mathbf{m}_{\text{ave}}(\theta) + O(\delta), \quad (21)$$

where $O(\delta)$ is a quantity such that the ratio $O(\delta)/\delta$ remains bounded as $\delta \rightarrow 0$. Likewise, using part 4) of Lemma 1, we conclude that the covariance of the steady-state random vector $\lambda_k^{(\delta)}$ results in:

$$\begin{aligned} \mathbf{c}_k^{(\delta)}(\theta, \theta') &\triangleq \mathbb{E} \left[\left(\lambda_k^{(\delta)}(\theta) - \mathbf{m}_k^{(\delta)}(\theta) \right) \left(\lambda_k^{(\delta)}(\theta') - \mathbf{m}_k^{(\delta)}(\theta') \right) \right] \\ &= \frac{\mathbf{c}_{\text{ave}}(\theta, \theta')}{2} \delta + O(\delta^2). \end{aligned} \quad (22)$$

Note that (21) and (22) can be rewritten in vector and matrix form, respectively as:

$$\mathbf{m}_k^{(\delta)} = \mathbf{m}_{\text{ave}} + O(\delta), \quad \mathbf{C}_k^{(\delta)} = \frac{\mathbf{C}_{\text{ave}}}{2} \delta + O(\delta^2) \quad (23)$$

where $\mathbf{C}_k^{(\delta)} = [\mathbf{c}_k^{(\delta)}(\theta, \theta')]$ and $\mathbf{C}_{\text{ave}} = [\mathbf{c}_{\text{ave}}(\theta, \theta')]$. The first equation in (23) shows that the expectation vector of the steady-state log-belief ratios, $\mathbf{m}_k^{(\delta)}$, approximates, for small δ , the expectation vector of the average log-likelihood ratios, \mathbf{m}_{ave} . Moreover, the second equation in (23) reveals that the covariance matrix of the steady-state log-belief ratios, $\mathbf{C}_k^{(\delta)}$, goes to zero as $\mathbf{C}_{\text{ave}} \delta/2$, where \mathbf{C}_{ave} is the covariance matrix of the average log-likelihood ratios.

Theorem 3 (Asymptotic normality). *Assume that the data $\{\xi_{k,i}\}$ are independent across the agents (recall that they are always assumed i.i.d. over time), and that the log-likelihood ratios have finite variance. Then, under Assumptions 1, 2 and 3, the following convergence in distribution holds:*

$$\frac{\lambda_k^{(\delta)} - \mathbf{m}_{\text{ave}}}{\sqrt{\delta}} \stackrel{\delta \rightarrow 0}{\rightsquigarrow} \mathcal{G}\left(0, \frac{\mathbf{C}_{\text{ave}}}{2}\right) \quad (24)$$

where $\mathcal{G}(0, C)$ is a zero-mean multivariate Gaussian with covariance matrix equal to C . ■

The result in Theorem 3 comes from Lemma 1, part 5). As $\delta \rightarrow 0$, Theorem 3 suggests the approximation:

$$\lambda_k^{(\delta)} \approx \mathcal{G}\left(\mathbf{m}_{\text{ave}}, \frac{\mathbf{C}_{\text{ave}}}{2} \delta\right). \quad (25)$$

Using the expressions in (21) and (22) instead of the limiting \mathbf{m}_{ave} and $\mathbf{C}_{\text{ave}} \delta/2$, we get the alternative approximation:

$$\lambda_k^{(\delta)} \approx \mathcal{G}\left(\mathbf{m}_k^{(\delta)}, \mathbf{C}_k^{(\delta)}\right), \quad (26)$$

which can capture different performance across agents, since $\mathbf{m}_k^{(\delta)}$ and $\mathbf{C}_k^{(\delta)}$ depend on k .

V. SIMULATION RESULTS

We consider the network topology displayed in Fig. 2 (additionally, we allow a self-loop for each agent). The combination matrix is designed using an averaging rule, resulting in a left-stochastic matrix [14].




Fig. 2. Network topology with 10 agents: agent 1 is highlighted.

The network is faced with the following statistical learning problem. We consider a family of Laplace likelihood functions with scale parameter equal to 1, and with different expectations parametrized with $n \in \{1, 2, 3\}$ as follows:

$$f_n(\xi) = \frac{1}{2} \exp\{-|\xi - 0.5n|\}. \quad (27)$$

We assume that the inference problem is *not locally identifiable* since we consider the setup in Table I for each agent's family of likelihood functions.

TABLE I
IDENTIFIABILITY SETUP FOR THE NETWORK IN FIG. 2.

| Agent k | Likelihood Function: $L_k(\theta)$ | | |
|-----------|------------------------------------|--------------|--------------|
| | $\theta = 1$ | $\theta = 2$ | $\theta = 3$ |
| 1 – 3 | $f_1(\xi)$ | $f_1(\xi)$ | $f_3(\xi)$ |
| 4 – 6 | $f_1(\xi)$ | $f_3(\xi)$ | $f_3(\xi)$ |
| 7 – 10 | $f_1(\xi)$ | $f_2(\xi)$ | $f_1(\xi)$ |

A. Consistency

We consider that all agents are running the ASL algorithm for a fixed $\theta_0 = 1$ over 10000 time samples (after which we consider that they achieved the steady state). From Theorem 2, we saw that as δ approaches zero, all agents k are able to consistently learn — see (19). To show this effect, for each value of δ (50 sample points in the interval $\delta \in [0.001, 1]$ are taken), we consider a different realization of the observations. In Fig. 3, for agent 1 and $\theta = 2, 3$, we show how the log-belief ratios $\lambda_1^{(\delta)}(\theta)$ behave for decreasing values of δ . We see the effect of the weak law of small stepsizes, since the limiting log-belief ratios tend to concentrate around \mathbf{m}_{ave} .




Fig. 3. Evolution of steady-state log-belief ratios for agent 1 as $\delta \rightarrow 0$.

B. Asymptotic Normality

From Theorem 3, we saw that we can approximate the steady-state log-belief ratios by a multivariate Gaussian, see Eqs. (25) and (26). In Fig. 4, we display the log-belief ratios for instant $i = 10000$. The experiment is repeated over 100 Monte Carlo runs, such that we obtain 100 realizations of the steady-state variable $\lambda_k^{(\delta)}$. Moreover, we consider four decreasing values of δ .

In dashed blue lines we see the ellipses representing the confidence intervals relative to one and two standard deviations computed for the empirical Gaussian approximation seen in (26): the smaller ellipse encompasses approximately 68% of the samples whereas the larger ellipse encompasses 95%. In red dotted lines, we see the corresponding ellipses for the limiting theoretical Gaussian approximation seen in (25), with the red cross indicating the limiting theoretical expectation \mathbf{m}_{ave} . Note how as δ decreases, the ellipses tend to be smaller, which is in accordance with the scaling of the covariance matrices by δ in (25) and (26), and the distributions tend to overlap, which is in accordance with the behavior predicted by Theorem 3.

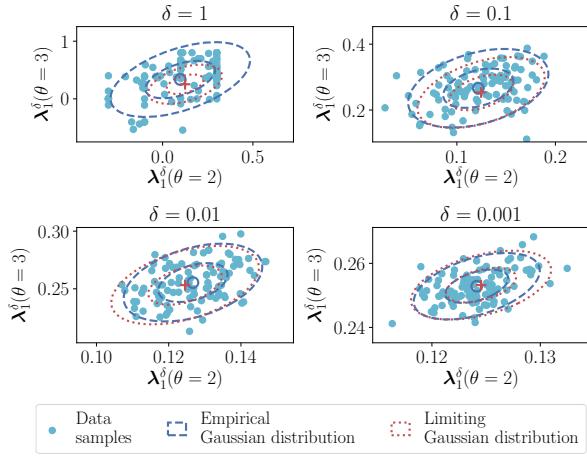


Fig. 4. Distribution of data samples at steady state compared with the limiting and empirical Gaussian distributions for decreasing δ .

VI. CONCLUSION

In this paper, we proposed the Adaptive Social Learning strategy as a way to address the significant inertia of the traditional social learning to adapting. We have first characterized the behavior of the algorithm in steady state, by showing that the log-belief ratios converge to a stable random variable. Then, exploring the regime of small δ , we could verify the algorithm's learning consistency and the limiting Gaussian behavior of the steady-state log-belief ratios.

APPENDIX A

Lemma 1 (Asymptotic properties of a useful random series). *For $m = 0, 1, \dots$, let $\{z_m\}$ be a sequence of i.i.d. integrable random variables with $m_z \triangleq \mathbb{E}[z_m]$ and $m_z^{\text{abs}} \triangleq \mathbb{E}[|z_m|] < \infty$. Let also $0 < \delta < 1$, and consider the following partial sums:*

$$s_i(\delta) = \delta \sum_{m=0}^i (1-\delta)^m \alpha_m z_m, \quad (28)$$

where $0 < \alpha_m < 1$, with α_m converging to some value α and obeying the following upper bound for all m :

$$|\alpha_m - \alpha| \leq \kappa \beta^m, \quad (29)$$

for some constant $\kappa > 0$ and for some $0 < \beta < 1$. Then we have the following asymptotic properties.

- 1) **Steady-state stability.** *The partial sums in (28) are almost-surely absolutely convergent, namely, we can define the (almost-surely) convergent series:*

$$s^{\text{abs}}(\delta) \triangleq \delta \sum_{m=0}^{\infty} (1-\delta)^m \alpha_m |z_m|, \quad (30)$$

$$s(\delta) \triangleq \delta \sum_{m=0}^{\infty} (1-\delta)^m \alpha_m z_m. \quad (31)$$

- 2) **First moment.** *The expectation of $s(\delta)$ is:*

$$\mathbb{E}[s(\delta)] = m_z \delta \sum_{m=0}^{\infty} (1-\delta)^m \alpha_m = \alpha m_z + O(\delta), \quad (32)$$

where $O(\delta)$ is a quantity such that the ratio $O(\delta)/\delta$ remains bounded as $\delta \rightarrow 0$.

- 3) **Weak law of small step-sizes.** *The series $s(\delta)$ converges to αm_z in probability as $\delta \rightarrow 0$, namely, for all $\epsilon > 0$ we have that:*

$$\lim_{\delta \rightarrow 0} \mathbb{P}[|s(\delta) - \alpha m_z| > \epsilon] = 0. \quad (33)$$

- 4) **Second moment.** *If $\sigma_z^2 \triangleq \text{VAR}[z_m] < \infty$, then:*

$$\begin{aligned} \text{VAR}[s(\delta)] &= \sigma_z^2 \delta^2 \sum_{m=0}^{\infty} (1-\delta)^{2m} \alpha_m^2 \\ &= \frac{\alpha^2 \sigma_z^2}{2} \delta + O(\delta^2). \end{aligned} \quad (34)$$

- 5) **Asymptotic normality.** *If z_m has finite variance σ_z^2 , then the following convergence in distribution holds:*

$$\frac{s(\delta) - m_z}{\sqrt{\delta}} \xrightarrow{\delta \rightarrow 0} \mathcal{G}\left(0, \alpha^2 \sigma_z^2 / 2\right), \quad (35)$$

and, hence, $s(\delta)$ is asymptotically normal as $\delta \rightarrow 0$.

REFERENCES

- [1] C. Chamley, *Rational Herds: Economic Models of Social Learning*. Cambridge University Press, UK, 2004.
- [2] A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi, "Non-Bayesian social learning," *Games and Economic Behavior*, vol. 76, no. 1, pp. 210–225, Sep. 2012.
- [3] 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 Process. Mag.*, vol. 30, no. 3, pp. 43–57, May 2013.
- [4] C. Chamley, A. Scaglione, and L. Li, "Models for the diffusion of beliefs in social networks: An overview," *IEEE Signal Process. Mag.*, vol. 30, no. 3, pp. 16–29, May 2013.
- [5] D. Acemoglu and A. Ozdaglar, "Opinion dynamics and learning in social networks," *Dyn. Games Appl.*, vol. 1, no. 1, pp. 3–49, 2011.
- [6] A. Jadbabaie, P. Molavi, and A. Tahbaz-Salehi, "Information heterogeneity and the speed of learning in social networks," *Columbia Business School Research Paper*, pp. 13–28, May 2013.
- [7] X. Zhao and A. H. Sayed, "Learning over social networks via diffusion adaptation," in *Proc. Asilomar Conference on Signals, Systems and Computers*, Nov. 2012, pp. 709–713.
- [8] H. Salami, B. Ying, and A. H. Sayed, "Social learning over weakly connected graphs," *IEEE Trans. Signal Inf. Process. Netw.*, vol. 3, no. 2, pp. 222–238, Jun. 2017.
- [9] A. Nedić, A. Olshevsky, and C. A. Uribe, "Fast convergence rates for distributed non-Bayesian learning," *IEEE Trans. Autom. Control*, vol. 62, no. 11, pp. 5538–5553, Nov. 2017.
- [10] A. Lalitha, T. Javidi, and A. D. Sarwate, "Social learning and distributed hypothesis testing," *IEEE Trans. Inf. Theory*, vol. 64, pp. 6161–6179, Sep. 2018.
- [11] V. Matta, A. Santos, and A. H. Sayed, "Exponential collapse of social beliefs over weakly-connected heterogeneous networks," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Brighton, UK, May 2019, pp. 5267–5271.
- [12] V. Matta, V. Bordinon, A. Santos, and A. H. Sayed, "Interplay between topology and social learning over weak graphs," *submitted for publication*, available online as arXiv:1910.13905v1 [cs.MA], Oct. 2019.
- [13] V. Matta and A. H. Sayed, "Estimation and detection over adaptive networks," in *Cooperative and Graph Signal Processing*, P. Djuric and C. Richard, Eds. Elsevier, 2018, pp. 69–106.
- [14] A. H. Sayed, "Adaptation, Learning, and Optimization over Networks," *Found. Trends Mach. Learn.*, vol. 7, no. 4-5, pp. 311–801, 2014.
- [15] H. Chernoff, "A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations," *The Annals of Mathematical Statistics*, vol. 23, no. 4, pp. 493–507, Dec. 1952.