OPTIMAL COMBINATION POLICIES FOR ADAPTIVE SOCIAL LEARNING

Ping Hu⋆*, Virginia Bordignon*⋆*, Stefan Vlaski*†*, Ali H. Sayed*⋆

⋆School of Engineering, École Polytechnique Fédérale de Lausanne (EPFL)
†Department of Electrical and Electronic Engineering, Imperial College London

ABSTRACT

This paper investigates the effect of combination policies on the performance of adaptive social learning in non-stationary environments. By analyzing the relation between the error probability and the underlying graph topology, we prove that in the slow adaptation regime, combination policies with a uniform Perron eigenvector will provide the smallest steady-state error probability. This result indicates that in terms of learning accuracy, doubly-stochastic combination policies yield optimal performance. Moreover, we estimate the adaptation time of adaptive social learning in the small signal-to-noise regime and show that in this regime, the influence of combination policies on the adaptation time is insignificant.

Index Terms— Social learning, combination policy, large deviations, adaptation time.

1. INTRODUCTION

Social learning is an inference process over multi-agent networks where agents work collaboratively to identify the true state from a set of admissible hypotheses. In each step, agents update their belief by combining the new local observations with the information from their neighboring agents using a given *combination policy*.

In recent years, variations of social learning strategies have been proposed, which provably enable truth learning for all agents. Representative algorithms include the linear social learning algorithms [1–3], log-linear social learning algorithms [4–7] and some newly proposed algorithms [8, 9]. All these variants provide almost-sure convergence guarantees for the learning process. As pointed out in [10], a remarkable learning performance, however, comes at the cost of reduced adaptation capabilities. Agents behave stubbornly when faced with state changes, which is detrimental for social learning in non-stationary environments.

To address this difficulty, the work [10] introduced an *adaptive social learning* (ASL) algorithm, which greatly improves the agents' adaptation ability for tracking drifts in the statistical properties of the data. This was achieved by introducing a new parameter δ to control the amount of weighting given to recent observations in relation to past observations. The parameter δ was shown to control a fundamental trade-off between steady-state learning performance and adaptation ability. In particular, it was shown that in the small-δ regime, the error probability decays exponentially with $1/\delta$. Moreover, the decaying rate is affected by the statistical properties of the data and the *centrality measure* of each agent.

The centrality of each agent is a function of the graph topology and the combination policy (left or doubly-stochastic) used by the social learning algorithm. In this work, we investigate the influence of combination policies on the learning performance. A similar problem has been considered in previous studies on learning methods, e.g., [7, 11]. However, since the error probability converges to zero almost surely in the non-adaptive scenario, the purpose of optimizing the combination policy in [7, 11] was only meant to accelerate the convergence rate of the learning process.

In the context of *adaptive* social learning, the steady-state error probability is non-zero, but decays exponentially with $1/\delta$. We thus formulate the optimization of the combination policy as finding the centrality vector (Perron eigenvector) that results in the largest *error exponent* of the steady-state error probability. Based on the large deviation analysis from [10], we prove that the uniform Perron eigenvector is an optimal solution. A direct implication of this result is that doubly-stochastic combination policies enable better steady-state accuracy for small step-sizes. In addition, we also discuss the effect of combination policies on the transient process of adaptive social learning. The *adaptation time* of the ASL strategy, which is an important index of learning behavior, is estimated and compared for different combination policies. Our results show that in the small signal-to-noise ratio (SNR) regime, combination policies play a minor role in influencing the adaptation time, and that it is sufficient to rely on doubly stochastic matrices. This conclusion is in contrast to the analogous result in non-adaptive social learning [7, 11], where a positive relation between the informativeness of agents and the centrality of agents is highlighted for improving the learning performance.

2. PROBLEM SETTING

We consider a collection of N agents working collectively to agree on a hypothesis that best explains the streaming, dis-

This work was supported in part by SNSF grant 205121-184999. E-mails: ping.hu@epfl.ch, virginia.bordignon@epfl.ch, s.vlaski@imperial.ac.uk, ali.sayed@epfl.ch.

tributed observations in the network. At each time instant i, each agent k observes a private signal $\boldsymbol{\xi}_{k,i}$ (boldface notation for random variables) belonging to a certain space $\mathcal{X}_k$. The private signals at every agent, which are assumed to be statistically independent over time and space after conditioning on the true state, are realizations of a random variable following an unknown distribution L_k. The family of the distribution L_k is parameterized by H hypothesis $\theta \in \Theta = \{\theta_0, \theta_1, \dots, \theta_{H-1}\}$. The likelihood of the signal $\boldsymbol{\xi}_{k,i}$ conditioned on hypothesis θ is denoted by

$$L_k(\boldsymbol{\xi}_{k,i}|\theta), \quad \boldsymbol{\xi}_{k,i} \in \mathcal{X}_k. \tag{1}$$

To infer the true model using adaptive social learning [10], each agent k holds a local belief vector $\boldsymbol{\mu}_{k,i}$, which represents a probability mass function over the set of hypotheses Θ. The ASL algorithm [10] is described by:

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

$$\boldsymbol{\mu}_{k,i}(\theta) = \frac{\exp\big\{\sum_{\ell \in \mathcal{N}_k} a_{\ell k} \log \boldsymbol{\psi}_{\ell,i}(\theta)\big\}}{\sum_{\theta' \in \Theta} \exp\big\{\sum_{\ell \in \mathcal{N}_k} a_{\ell k} \log \boldsymbol{\psi}_{\ell,i}(\theta')\big\}} \tag{2b}$$

where $\mathcal{N}_k$ is the set of neighbors of agent k, and $a_{\ell k}$ is the combination weight that agent k places on the information received from the neighboring agent ℓ. The combination matrix $A = [a_{\ell k}]$ satisfies

$$A^\top \mathbb{1} = \mathbb{1}, \quad a_{\ell k} > 0, \quad \forall \ell \in \mathcal{N}_k \tag{3}$$

and $a_{\ell k} = 0$ for $\ell \notin \mathcal{N}_k$, where $\mathbb{1}$ denotes the N-dimensional vector of all ones. In addition, we assume that the communication network is strongly connected, which ensures that the Perron eigenvector π of matrix $A = [a_{\ell k}]$ will have strictly positive entries [12, 13]. That is,

$$A\pi = \pi, \quad \mathbb{1}^\top \pi = 1, \quad \pi_\ell > 0, \ \forall \ell = 1, 2, \dots, N. \tag{4}$$

We further impose the following three assumptions [1, 5, 8, 10].

Assumption 1 (Finiteness of Kullback-Leibler (KL) divergence). *The KL divergence between $L_k(\cdot|\theta)$ and $L_k(\cdot|\theta')$ is finite for each pair of distinct hypotheses (θ, θ') and for each agent.* □

Assumption 2 (Global identifiability). *For each wrong hypothesis $\theta \neq \theta_0$, there is at least one agent k for which the KL divergence between $L_k(\cdot|\theta_0)$ and $L_k(\cdot|\theta)$ is positive.* □

Assumption 3 (Positive initial belief). *For each hypothesis $\theta \in \Theta$ and each agent $k = 1, 2, \dots, N$, the initial belief $\boldsymbol{\mu}_{k,0}(\theta)$ is positive.* □

For the analysis in the sequel, we introduce the following normalized variables representing the log-likelihood ratio $\boldsymbol{x}_{k,i}(\theta)$ and the log-belief ratio $\boldsymbol{\lambda}_{k,i}(\theta)$ for all θ:

$$\boldsymbol{x}_{k,i}(\theta) \triangleq \log \frac{L_k(\boldsymbol{\xi}_{k,i}|\theta_0)}{L_k(\boldsymbol{\xi}_{k,i}|\theta)}, \ \boldsymbol{\lambda}_{k,i}(\theta) \triangleq \log \frac{\boldsymbol{\mu}_{k,i}(\theta_0)}{\boldsymbol{\mu}_{k,i}(\theta)}. \tag{5}$$

Using these variables, the ASL algorithm (2) can be rewritten as the two-step recursion

$$\begin{cases} \boldsymbol{\nu}_{k,i}(\theta) = (1-\delta)\boldsymbol{\lambda}_{k,i-1}(\theta) + \delta \boldsymbol{x}_{k,i}(\theta) \\ \boldsymbol{\lambda}_{k,i}(\theta) = \sum_{\ell \in \mathcal{N}_k} a_{\ell k} \boldsymbol{\nu}_{\ell,i}(\theta) \end{cases} \tag{6}$$

which has the form of a standard diffusion learning rule [12, 13]. We also introduce the weighted network average of log-likelihood ratios for all θ,

$$\boldsymbol{x}_{\mathsf{ave},i}(\theta) = \textstyle\sum_{\ell=1}^N \pi_\ell \boldsymbol{x}_{\ell,i}(\theta). \tag{7}$$

The expectations and variances of $\boldsymbol{x}_{k,i}(\theta)$ and $\boldsymbol{x}_{\mathsf{ave},i}(\theta)$ relative to the distribution $\prod_{k=1}^N L_k(\cdot|\theta_0)$ are denoted by $d_k(\theta)$, $\mathsf{m}_{\mathsf{ave}}(\theta)$, $\rho_k(\theta)$ and $\mathsf{c}_{\mathsf{ave}}(\theta)$. For agent k, the instantaneous error probability of social learning at time instant i is expressed as

$$\begin{aligned} p_{k,i} &= \mathbb{P}\Big[\arg\max_{\theta \in \Theta} \boldsymbol{\mu}_{k,i}(\theta) \neq \theta_0\Big] \\ &= \mathbb{P}\Big[\exists \theta \neq \theta_0 : \boldsymbol{\lambda}_{k,i}(\theta) \leq 0\Big]. \end{aligned} \tag{8}$$

As i approaches infinity, we obtain the steady-state error probability p_k for agent k, i.e.,

$$p_k = \lim_{i \to \infty} p_{k,i}, \tag{9}$$

which was proven to exist in [10].

3. MINIMIZING PROBABILITY OF ERROR

Let $\Lambda_k(t;\theta)$ and $\Lambda_{\mathsf{ave}}(t;\theta)$ denote the Logarithmic Moment Generating Function (LMGF) of the log-likelihood ratio $\boldsymbol{x}_{k,i}(\theta)$ and the average log-likelihood ratio $\boldsymbol{x}_{\mathsf{ave},i}(\theta)$:

$$\Lambda_k(t;\theta) = \log \mathbb{E}\Big[e^{t\boldsymbol{x}_{k,i}(\theta)}\Big], \tag{10}$$

$$\Lambda_{\mathsf{ave}}(t;\theta) = \log \mathbb{E}\Big[e^{t\boldsymbol{x}_{\mathsf{ave},i}(\theta)}\Big] \overset{(7)}{=} \sum_{\ell=1}^N \Lambda_\ell(\pi_\ell t;\theta). \tag{11}$$

Using the Gartner-Ellis Theorem [14], it was proven in [10] that the steady-state error probability (9) obeys a Large Deviation Principle (LDP) with a rate function determined by the LMGFs. We report here Theorem 4 from [10] for ease of reference.

Lemma 1 (Theorem 4 in [10]). *Assume that the LMGF of $\boldsymbol{x}_{k,i}(\theta)$ exists everywhere, namely, $\Lambda_k(t;\theta) < +\infty, \forall t \in \mathbb{R}$ for all $k = 1, 2, \dots, N$ and $\theta \neq \theta_0$. Let*

$$\phi(t;\theta) = \int_0^t \frac{\Lambda_{\mathsf{ave}}(\tau;\theta)}{\tau} d\tau, \tag{12}$$

$$\Phi(\theta) = -\inf_{t \in \mathbb{R}} \phi(t;\theta). \tag{13}$$

Then, under Assumptions 1–3, the error probability p_k obeys the LDP:

$$p_k \simeq e^{-\Phi/\delta} \tag{14}$$

with the error exponent $\Phi = \min_{\theta \neq \theta_0} \Phi(\theta)$, where the notation $\simeq$ denotes equality to the leading order in the exponent as δ goes to zero. ■

According to (12) and (13), the error exponent Φ is determined by the LMGF of $\boldsymbol{x}_{\mathsf{ave},i}(\theta)$, which as seen in (7), depends on the Perron eigenvector of the combination matrix. To find the best Perron eigenvector that provides the largest Φ, we formulate the following optimization problem:

$$\max_{\pi} \Phi \tag{15}$$

$$s.t. \quad \mathbb{1}^\top \pi = 1, \tag{16}$$

$$\pi_\ell > 0, \quad \forall \ell = 1, 2, \dots, N. \tag{17}$$

We show next that the uniform Perron eigenvector is an optimal solution.

Theorem 1 (**Optimal Perron eigenvector**). *The maximum error exponent of the steady-state error probability is achieved when the Perron eigenvector is uniform, i.e.,*

$$\frac{1}{N}\mathbb{1} \in \arg\max_{\pi} \Phi \quad s.t. \text{ (16) } \textit{and} \text{ (17)}. \tag{18}$$

Proof. The proof relies on the strict convexity of the LMGF $\Lambda_\ell(t;\theta)$ and the rate functions $\phi(t;\theta)$, as well as the property that $\Lambda_\ell(t;\theta) = 0$ for $t = 0$ and -1. From (12), we have

$$\begin{aligned}
\phi(t;\theta) &= \int_0^t \frac{\Lambda_{\mathsf{ave}}(\tau;\theta)}{\tau} d\tau = \sum_\ell \int_0^{\pi_\ell t} \frac{\Lambda_\ell(\tau;\theta)}{\tau} d\tau \\
&= \sum_\ell \int_0^{-1} \frac{\Lambda_\ell(\tau;\theta)}{\tau} d\tau + \sum_{\ell:\pi_\ell t < -1} \int_{-1}^{\pi_\ell t} \underbrace{\frac{\Lambda_\ell(\tau;\theta)}{\tau}}_{\leq 0} d\tau \\
&\quad + \sum_{\ell:\pi_\ell t > -1} \int_{-1}^{\pi_\ell t} \underbrace{\frac{\Lambda_\ell(\tau;\theta)}{\tau}}_{\geq 0} d\tau \\
&\geq \sum_\ell \int_0^{-1} \frac{\Lambda_\ell(\tau;\theta)}{\tau} d\tau,
\end{aligned} \tag{19}$$

where the equality holds if $\pi_\ell t = -1$ for all ℓ. Due to the constraints (16) and (17), we deduce that $\pi^\star = \frac{1}{N}\mathbb{1}$ and $t^\star = -N$ is an optimal solution to (15). ∎

Since the error exponent determines the leading-order decay rate in (14), we conclude from Theorem 1 that in the small-δ regime, any doubly-stochastic combination policy will be preferable for reducing the steady-state error probability of adaptive social learning.

4. MINIMIZING ADAPTATION TIME IN THE SMALL SNR REGIME

The *adaptation time* is defined as the critical time instant i after which the instantaneous error probability is decaying with an error exponent $(1-\epsilon)\Phi$ for some small $\epsilon > 0$:

$$p_{k,i} \leq e^{-\frac{1}{\delta}[(1-\epsilon)\Phi + \mathcal{O}(\delta)]} \tag{20}$$

where the notation $\mathcal{O}(\delta)$ signifies that the ratio $\mathcal{O}(\delta)/\delta$ stays bounded as $\delta \to 0$.

We examine the learning task in the small SNR regime where the error probabilities need not be too small [15] and $\Lambda_{\mathsf{ave}}(t;\theta)$ can be approximated by a second-order polynomial for $t \in [t_\theta^\star, 0]$. Here, $t_\theta^\star < 0$ is the unique solution to $\Lambda_{\mathsf{ave}}(t_\theta^\star;\theta) = 0$. Then,

$$\Lambda_{\mathsf{ave}}(t;\theta) = \sum_{n=1}^{\infty} \frac{\kappa_n t^n}{n!} \approx \kappa_1(\theta)t + \frac{\kappa_2(\theta)}{2}t^2 \tag{21}$$

with the cumulants $\kappa_1(\theta) = \mathsf{m}_{\mathsf{ave}}(\theta)$ and $\kappa_2(\theta) = \mathsf{c}_{\mathsf{ave}}(\theta)$. Since the parabolic approximation is actually a Gaussian approximation, (21) is accurate only if $\boldsymbol{x}_{\mathsf{ave},i}(\theta)$ obeys a Gaussian distribution (e.g., in the canonical *shift-in-mean Gaussian problems* [15]). For non-Gaussian cases, (21) will be valid only in the small SNR regime [15]. The exact definition of the small SNR regime depends on the specific learning task, but it generally includes the scenarios where the hypotheses are close to each other and which makes the learning task difficult. As indicated by [16], this regime is related to detecting weak signals in the framework of *locally optimum detection* [17, 18]. In the small SNR regime, we can derive an explicit approximation result for the adaptation time.

Theorem 2 (**Adaptation time for the small SNR regime**). *Consider the uniform initial belief condition and the small SNR regime, then the adaptation time* $\mathsf{T}_{\mathsf{adap}}$ *can be approximated as*

$$\mathsf{T}_{\mathsf{adap}} \approx \frac{\log\left(1 - \sqrt{1-\epsilon}\right)}{\log(1-\delta)} \tag{22}$$

for any combination policy.

Sketch of Proof. According to (6), we can write

$$\boldsymbol{\lambda}_{k,i}(\theta) \overset{\mathrm{d}}{=} \delta \sum_{m=0}^{i-1} \sum_{\ell=1}^{N} (1-\delta)^m [A^{m+1}]_{\ell k} \boldsymbol{x}_{\ell,m+1}(\theta) \triangleq \widetilde{\boldsymbol{\lambda}}_{k,i}(\theta), \tag{23}$$

where $\overset{\mathrm{d}}{=}$ denotes equality in distribution. The key step in the proof is to upper bound the instantaneous error probability by using Markov's inequality and approximating the LMGF of $\widetilde{\boldsymbol{\lambda}}_{k,i}(\theta)$ via $\Lambda_{\mathsf{ave}}(t;\theta)$. Applying Markov's inequality, we have

$$\begin{aligned}
\mathbb{P}\left[\boldsymbol{\lambda}_{k,i}(\theta) \leq 0\right] &\overset{(23)}{=} \mathbb{P}\left[\frac{t_\theta^\star}{\delta}\widetilde{\boldsymbol{\lambda}}_{k,i}(\theta) \geq 0\right] \\
&\leq \mathbb{E}\left[\exp\left(\frac{t_\theta^\star}{\delta}\widetilde{\boldsymbol{\lambda}}_{k,i}(\theta)\right)\right].
\end{aligned} \tag{24}$$

Denoting the LMGF of $\widetilde{\boldsymbol{\lambda}}_{k,i}(\theta)$ by $\Lambda_{k,i}(t;\theta)$and using Eqs. (85) and (86) from [16], we can derive

$$\Lambda_{k,i}\left(\frac{t_\theta^\star}{\delta};\theta\right) = \frac{1}{\delta}\left[\int_{(1-\delta)^i t_\theta^\star}^{t_\theta^\star} \frac{\Lambda_{\mathsf{ave}}(\tau;\theta)}{\tau} d\tau + \mathcal{O}(\delta)\right]. \tag{25}$$

Replacing the parabolic approximation (21) for the small SNR regime into (25), we are able to prove that

$$\Lambda_{k,i}\left(\frac{t_\theta^\star}{\delta};\theta\right) \leq -\frac{1}{\delta}\left[(1-\epsilon)\Phi(\theta) + \mathcal{O}(\delta)\right] \tag{26}$$

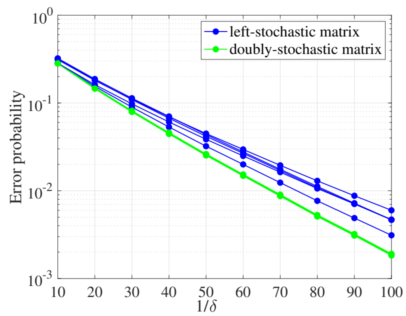

Fig. 1. Steady-state error probability.

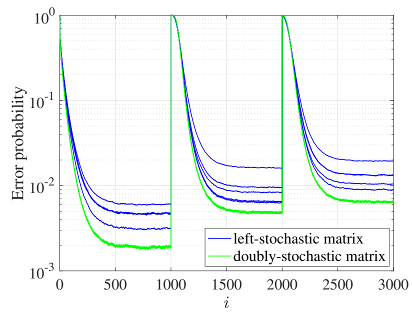

Fig. 2. Instantaneous error probability.

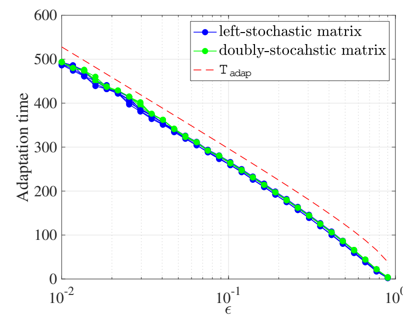

Fig. 3. Adaptation time

for $i \geq \mathsf{T}_{\mathsf{adap}}$. Using Boole's inequality, the instantaneous error probability is upper bounded by

$$p_{k,i} \overset{(8)}{=} \mathbb{P}\left[\bigcup_{\theta\neq\theta_0}\left\{\widetilde{\boldsymbol{\lambda}}_{k,i}(\theta)\leq 0\right\}\right] \leq \sum_{\theta\neq\theta_0}\mathbb{P}\left[\widetilde{\boldsymbol{\lambda}}_{k,i}(\theta)\leq 0\right]$$

$$\overset{(26)}{\leq} \sum_{\theta\neq\theta_0} e^{-\frac{1}{\delta}\left[(1-\epsilon)\Phi(\theta)+\mathcal{O}(\delta)\right]} \simeq e^{-\frac{1}{\delta}\left[(1-\epsilon)\Phi+\mathcal{O}(\delta)\right]}$$

for $i \geq \mathsf{T}_{\mathsf{adap}}$, which shows that $\mathsf{T}_{\mathsf{adap}}$ provides a reasonable estimate for the adaptation time in the small SNR regime. ■

Theorem 2 indicates that when the hypotheses are hard to distinguish, the influence of different combination policies on the adaptation time is insignificant. Instead, it is the step-size δ that plays the dominant role in the adaptation time [10]. This fact ensures that choosing a doubly-stochastic combination policy, as suggested by optimizing steady-state performance in Theorem 1, does not negatively impact the transient learning performance in the small SNR regime.

5. NUMERICAL SIMULATIONS

In this section, we present simulation results on an Erdös-Rényi random graph with connection probability 0.5. We also assume that each agent has a self-loop and will perform a social learning protocol with three hypotheses $\{\theta_0, \theta_1, \theta_2\}$. The learning rule for each agent k is presented in Algorithm 1. We consider a family of Laplace likelihood functions with scale parameter one, i.e., $f_n(\xi) \triangleq L(\xi|\theta_n) = \frac{1}{2}\exp\{-|\xi - 0.1n|\}$ for $n \in \{0, 1, 2\}$. The local likelihoods of the three hypotheses $\theta_0, \theta_1, \theta_2$ are respectively f_0, f_0, f_1 for agents 1–3, and f_0, f_2, f_2 for agents 4–7, and f_0, f_2, f_0 for the rest of agents.

Algorithm 1 ASL rule for each agent k

Initialization: $\boldsymbol{\mu}_{k,0}(\theta) = \frac{1}{3}, \forall\theta \in \Theta$.
for $i = 1, 2, \ldots$ **do**
  receives a new observation $\boldsymbol{\xi}_{k,i}$;
  updates the intermediate belief $\boldsymbol{\psi}_{k,i}$ according to (2a);
  obtains $\boldsymbol{\psi}_{\ell,i}$ from its neighbors $\ell \in \mathcal{N}_k$;
  updates $\boldsymbol{\mu}_{k,i}$ according to (2b).
end for

First, we study the effect of the combination policies on the error exponent. A stationary environment where the true hypothesis is selected as θ_0 is considered. For comparison, we employ 5 left-stochastic and 5 doubly-stochastic combination matrices with positive Perron eigenvectors. In Fig. 1, the average steady-state error probabilities under 10 combination matrices and different step-sizes are presented. For each step-size, we select the terminal time as 10^3 and run 10^6 Monte Carlo simulations to obtain the average results. It can be observed that all doubly-stochastic combination matrices lead to a lower error probability than the left-stochastic ones.

Next, we investigate the effect of combination policies on the adaptation time. Here, we consider a non-stationary environment where the true state changes from θ_0 to θ_1 at $i = 1000$ and from θ_1 to θ_2 at $i = 2000$. Under a small step-size $\delta = 0.01$, the transient dynamics of average error probability over $i \in [0, 3000]$ is depicted in Fig. 2. It is observed that even in the non-stationary environment, the adaptation time related to different combination matrices is very close to each other. To derive a quantitative comparison, we calculate the simulated adaptation time when the true state is θ_0 (i.e., $i \in [0, 1000]$). By the definition of adaptation time, we record the time instant i_0 after which the error probability satisfies $\log p_{k,i} \leq (1-\epsilon)\log p_k, i \geq i_0$. The simulated adaptation time under different values of ϵ is presented in Fig. 3. It is clear that the difference in adaptation time for all considered combination matrices is almost negligible irrespective of ϵ.

6. CONCLUSION

In this work, we discussed the effect of combination policies on two key performance metrics of adaptive social learning: the error exponent (i.e., steady-state learning ability) and the adaptation time (i.e., transient behavior). Our results show that in the small-δ regime, the best error exponent is achieved by doubly-stochastic combination policies. Moreover, the difference of the adaptation time among different combination policies is almost negligible if the SNR between hypotheses is small. Importantly, these results are in contrast to analogous results in the context of distributed optimization [12, 13, 19, 20] when agents have access to data of varying quality.

7. REFERENCES

[1] 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.

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

[3] 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, May 2017.

[4] A. Nedic, 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, Oct. 2017.

[5] 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, Aug. 2018.

[6] 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.

[7] 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, Oct. 2016.

[8] A. Mitra, J. A. Richards, and S. Sundaram, "A new approach to distributed hypothesis testing and non-bayesian learning: Improved learning rate and byzantine-resilience," *IEEE Transactions on Automatic Control*, vol. 66, no. 9, pp. 4084–4100, Oct. 2021.

[9] C.-X. Shi and G.-H. Yang, "Distributed learning over networks: Effect of using historical observations," *IEEE Transactions on Automatic Control*, vol. 65, no. 12, pp. 5503–5509, Feb. 2020.

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

[11] A. Jadbabaie, P. Molavi, and A. Tahbaz-Salehi, "Information heterogeneity and the speed of learning in social networks," *Columbia Business School Research Paper*, no. 13-28, May 2013.

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

[13] A. H. Sayed, "Adaptive networks," *Proceedings of the IEEE*, vol. 102, no. 4, pp. 460–497, Apr. 2014.

[14] F. Den Hollander, *Large Deviations*. American Mathematical Society, 2008.

[15] V. Matta, P. Braca, S. Marano, and A. H. Sayed, "Diffusion-based adaptive distributed detection: Steady-state performance in the slow adaptation regime," *IEEE Transactions on Information Theory*, vol. 62, no. 8, pp. 4710–4732, Jul. 2016.

[16] V. Matta, P. Braca, S. Marano, and A. H. Sayed, "Distributed detection over adaptive networks: Refined asymptotics and the role of connectivity," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 2, no. 4, pp. 442–460, Oct. 2016.

[17] P. Braca, S. Marano, V. Matta, and P. Willett, "Asymptotic optimality of running consensus in testing binary hypotheses," *IEEE Transactions on Signal Processing*, vol. 58, no. 2, pp. 814–825, Feb. 2010.

[18] S. A. Kassam, *Signal Detection in Non-Gaussian Noise*. Springer-Verlag New York, 1987.

[19] S. Vlaski and A. H. Sayed, "Distributed learning in non-convex environments—Part I: Agreement at a linear rate," *IEEE Transactions on Signal Processing*, vol. 69, pp. 1242–1256, 2021.

[20] S. Vlaski and A. H. Sayed, "Distributed learning in non-convex environments—Part II: Polynomial escape from saddle-points," *IEEE Transactions on Signal Processing*, vol. 69, pp. 1257–1270, 2021.