Paper ID: 1704

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

EPFL

Imperial College London

22-27 MAY 2022 ICASSP 2022 SINGAPORE

Abstract

This work studies the effect of combination policies on the performance of adaptive social learning. We proved that in the slow adaptation regime, combination policies with a uniform Perron eigenvector will provide the smallest steady-state error probability. 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.

Introduction

Adaptive social learning [1] is an inference process over multi-agent networks. The basic setup of a social learning network of N agents is:

- Local observations: $\boldsymbol{\xi}_{k,i} \in \mathcal{X}_k$ for each agent k at each time instant i
- Local likelihood models: $\{L_k(\cdot|\theta)\}$ parameterized by a hypothesis θ
- Finite hypothesis set: $\Theta = \{\theta_0, \dots, \theta_{H-1}\}$
- Global inference task: learning the true state θ^* in Θ that best explains the local observations.

To infer the true model using adaptive social learning, 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 adaptive social learning (ASL) algorithm [1] is described by:

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

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

The combination policy $A = [a_{\ell k}]$ satisfies

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

and $a_{\ell k} = 0$ for $\ell \notin \mathcal{N}_k$, where $\mathbb{1}$ denotes the N-dimensional vector of all ones. The Perron eigenvector π of matrix A has strictly positive entries,

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

Learning Performance of the ASL Algorithm

- **Steady-state learning accuracy:** In the slow adaptation (i.e., with small δ) regime, the steady-state error probability p_k of agent k decays exponentially with $1/\delta$:

$$p_k \simeq e^{-\Phi/\delta}, \quad (5)$$

where the notation $\simeq$ denotes equality to the leading order in the exponent as δ goes to zero. The decaying rate Φ is also called *error exponent.*

- **Adaptation ability:** The *adaptation time* $\mathsf{T}_{\mathsf{adap}}$ of the adaptive social learning is defined as the critical time instant i after which the instantaneous error probability $p_{k,i}$ of agent k decays with an error exponent $(1-\epsilon)\Phi$ for some small $\epsilon > 0$:

$$p_{k,i} \le e^{-\frac{1}{\delta}[(1-\epsilon)\Phi + \mathcal{O}(\delta)]}, \quad (6)$$

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

- **Error Exponent:** Let $\boldsymbol{x}_{k,i}(\theta) \triangleq \log \frac{L_k(\boldsymbol{\xi}_{k,i}|\theta_0)}{L_k(\boldsymbol{\xi}_{k,i}|\theta)}$ and $\Lambda_k(t;\theta) \triangleq \log \mathbb{E}[\exp(t\boldsymbol{x}_{k,i}(\theta))]$. The error exponent Φ is expressed as

$$\Phi = \min_{\theta \neq \theta_0} \left\{ -\inf_{t \in \mathbb{R}} \int_0^t \frac{\Lambda_{\mathsf{ave}}(\tau;\theta)}{\tau} d\tau \right\} \quad (7)$$

where $\Lambda_{\mathsf{ave}}(t;\theta) \triangleq \sum_{k=1}^N \Lambda_k(\pi_k t;\theta)$.

Role of the Combination Policy

Theoretical Result 1: Maximizing the Error Exponent

The maximum error exponent of the steady-state error probability is achieved when the Perron eigenvector is uniform, i.e.,

$$\frac{1}{N}\mathbb{1} \in \arg\max_{\pi} \Phi \quad \text{s.t.} \quad \mathbb{1}^\top \pi = 1 \text{ and } \pi_\ell > 0, \; \forall \ell = 1, 2, \dots, N. \quad (8)$$

Theoretical Result 2: Minimizing the Adaptation Time

Consider the uniform initial belief condition and the small signal-to-noise ratio (SNR) regime [2], then the adaptation time $\mathsf{T}_{\mathsf{adap}}$ can be approximated as

$$\mathsf{T}_{\mathsf{adap}} \approx \frac{\log(1 - \sqrt{1-\epsilon})}{\log(1-\delta)} \quad (9)$$

for any combination policy.

Simulation Results

Consider a 10-agent network. The agents will perform a social learning protocol with 3 hypotheses and Laplacian likelihood models. We consider 5 left-stochastic combination policies and 5 doubly-stochastic ones. The performance of adaptive social learning under different combination policies is investigated, both for the stationary and non-stationary environments.

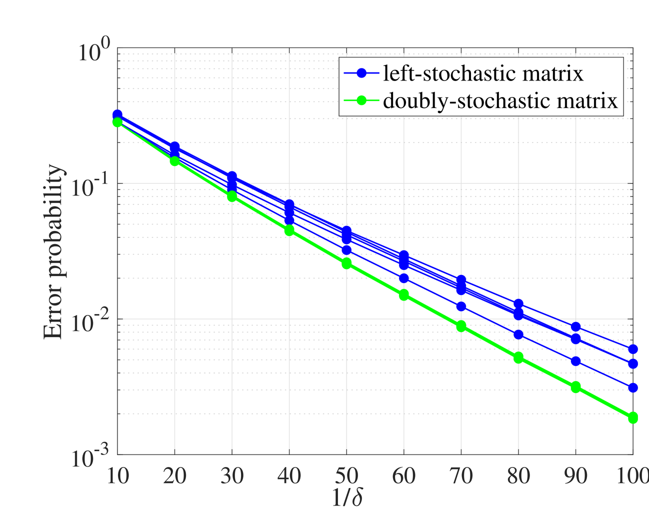

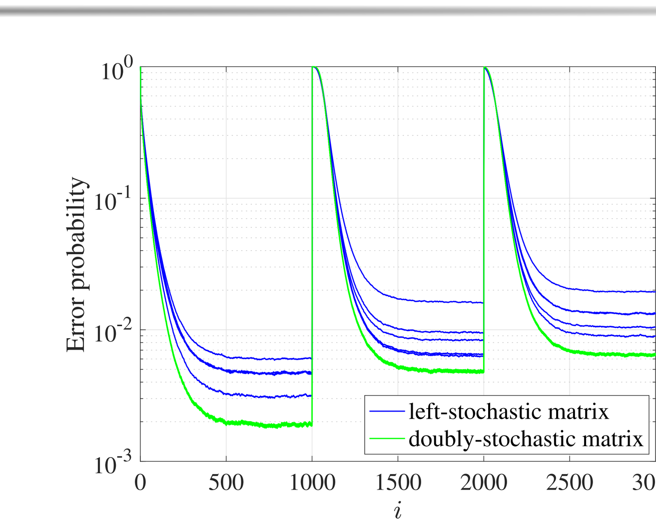

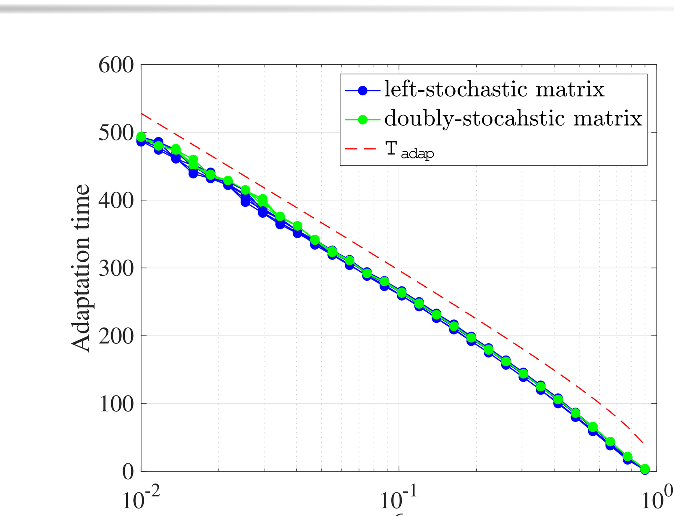

Concluding Remarks

- The largest error exponent can be achieved when the Perron eigenvector is uniform. Therefore, a doubly-stochastic combination policy will be beneficial for improving the steady-state learning accuracy.
- The combination policy plays a minor role in the adaptation time of the adaptive social learning when the SNR between hypotheses is small. It is then reliable to employ a combination policy with better steady-state learning accuracy in the small SNR regime.
- The results obtained in this work are based on the assumption that the likelihood model of the true state is accurate. The optimal combination policies for the generalized likelihood models have been established in our extended version [3].

References

[1] V. Bordignon, V. Matta, and A. H. Sayed. Adaptive social learning. *IEEE Transactions on Information Theory*, 67(9):6053–6081, 2021.

[2] 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*, 2(4):442–460, 2016.

[3] P. Hu, V. Bordignon, S. Vlaski, and A. H. Sayed. Optimal aggregation strategies for social learning over graphs. *arXiv preprint arXiv:2203.07065*, 2022.