Optimal Combination Policies for Adaptive Social Learning

Ping Hu, Virginia Bordignon, Stefan Vlaski, Ali.H Sayed

School of Engineering, École Polytechnique Fédérale de Lausanne
Department of Electrical and Electronic Engineering, Imperial College London

EPFL

Imperial College London

22-27 MAY 2022
ICASSP 2022
SINGAPORE

Outline

Introduction

Social learning problem

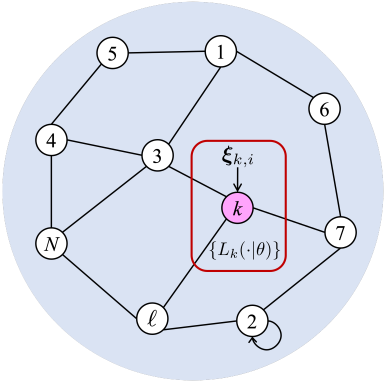

- Local observations: At each time instant i, each agent k receives a private signal $\boldsymbol{\xi}_{k,i}$ generated from an unknown probability distribution.
- Local likelihood models: Each agent k has a family of H likelihood models $\{L_k(\cdot|\theta)\}$ parameterized by a hypothesis belonging to a finite set

$$\Theta = \{\theta_0, \theta_1, \dots, \theta_{H-1}\}.$$

- Global learning task: Inferring the true state $\theta^* \in \Theta$ that best explains their observations.

Introduction

Social learning problem

- Strongly-connected network
- Left-stochastic combination policy $A = [a_{\ell k}]$:

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

and $a_{\ell k} = 0$ for $\ell \notin \mathcal{N}_k$.

- *Perron eigenvector* π:

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

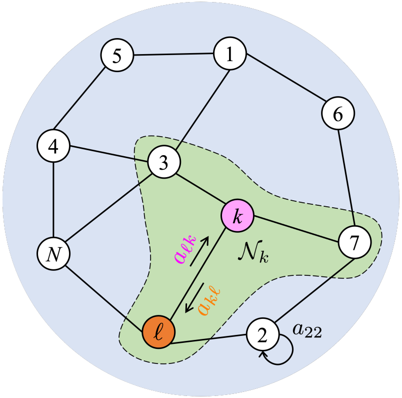

Introduction

Adaptive social learning (ASL) algorithm [2]

Adaptive Learning

1. Adaptation step:

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

Information Diffusion

2. Combination step:

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

- δ is the step size parameter
- $\boldsymbol{\mu}_{k,i}(\theta)$ is the belief of agent k on hypothesis θ at time instant i.

Introduction

Learning performance of the ASL algorithm

In the *slow adaptation regime* (i.e., with small δ), the ASL algorithm enables each agent in the network to

- learn the true hypothesis with a small error probability in the steady state (the steady-state error probability decays with $\frac{1}{\delta}$).
- track the changes of the true state faster than the non-adaptive social learning algorithms (adaptation time $\approx \mathcal{O}(\frac{1}{\delta})$).

➢ Can we improve the learning performance by optimizing the combination policy?

Introduction

Definitions and notations

Variables	Agent k	Network
Log-likelihood ratio	$\boldsymbol{x}_{k,i}(\theta) = \log \frac{L_k(\boldsymbol{\xi}_{k,i}\|\theta_0)}{L_k(\boldsymbol{\xi}_{k,i}\|\theta)}$	$\boldsymbol{x}_{\text{ave},i}(\theta) = \sum_{\ell=1}^{N} \pi_\ell \boldsymbol{x}_{\ell,i}(\theta)$
Logarithmic Moment Generating Function	$\Lambda_k(t;\theta) = \log \mathbb{E}\left[e^{t\boldsymbol{x}_{k,i}(\theta)}\right]$	$\Lambda_{\text{ave}}(t;\theta) = \sum_{\ell=1}^{N} \Lambda_\ell(\pi_\ell t;\theta)$
Instantaneous error probability	$p_{k,i} = \mathbb{P}\left[\arg\max_{\theta\in\Theta} \boldsymbol{\mu}_{k,i}(\theta) \neq \theta_0\right]$	$p_i^{\text{net}} = \frac{1}{N}\sum_{\ell=1}^{N} p_{\ell,i}$
Steady-state error probability	$p_k = \lim_{i\to\infty} p_{k,i}$	$p^{\text{net}} = \frac{1}{N}\sum_{\ell=1}^{N} p_\ell$

Theoretical Results

1. Maximizing the error exponent

Lemma 1[2]: Under the assumption of $\Lambda_{\text{ave}}(t;\theta) < \infty, \forall t$, the steady-state error probability p_k obeys a Large Deviation Principle (LDP) in the small-δ regime:

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

where the notation $\simeq$ denotes equality to the leading order in the exponent as δ goes to zero. The error exponent $\Phi = \min_{\theta \neq \theta_0} \Phi(\theta)$ is determined by the rate function given by

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

and $\Phi(\theta) = -\inf_{t \in \mathbb{R}} \phi(t;\theta)$ is the Legendre transform of $\phi(t;\theta)$ at point 0.

Theoretical Results

Question: What is the best Perron eigenvector π^* that leads to the largest error exponent?

- Solve the optimization problem: $\max_{\pi} \Phi \quad \text{s.t.} \ \ \pi > 0, \ \ \mathbb{1}^\top \pi = 1$

Optimal Perron eigenvector

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

$$\frac{1}{N}\mathbb{1} \in \arg\max_{\pi} \Phi \quad \text{s.t.} \ \pi > 0, \ \mathbb{1}^\top \pi = 1.$$

- Doubly-stochastic policies are preferred for improving the learning accuracy!

Theoretical Results

2. Minimizing the adaptation time

Definition 1[2]: The adaptation time, i.e., T_{adap}, 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 $\epsilon > 0$:

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

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

Theoretical Results

Question: What is the best combination policy that leads to the smallest adaptation time?

- Approximate the adaptation time for the small signal-to-noise ratio (SNR) regime [3]: $\Lambda_{\text{ave}}(t;\theta) \approx \kappa_1(\theta)t + \frac{\kappa_2(\theta)}{2}t^2$

Adaptation time for the small SNR regime

Theorem 2: Consider the uniform initial belief condition and the small SNR regime, then the adaptation time T_{adap} can be approximated as

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

for any combination policy.

Simulation Results

Simulation setup

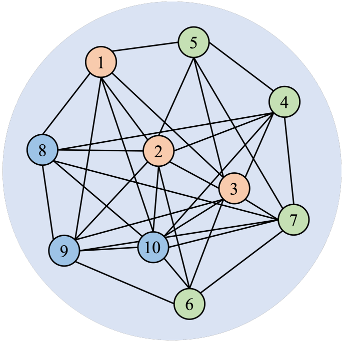

- Three hypotheses: $\Theta = \{\theta_0, \theta_1, \theta_2\}$
- Laplacian likelihood functions: $f_n(\xi) = \frac{1}{2}\exp\{-|\xi - 0.1n|\}, \ n \in \{0, 1, 2\}$
- Strongly-connected graph

Agent	Hypothesis		
	θ_0	θ_1	θ_2
1-3	$F_0(\xi)$	$F_0(\xi)$	$F_1(\xi)$
4-7	$F_0(\xi)$	$F_2(\xi)$	$F_2(\xi)$
8-10	$F_0(\xi)$	$F_2(\xi)$	$F_0(\xi)$

Simulation Results

Steady-state performance

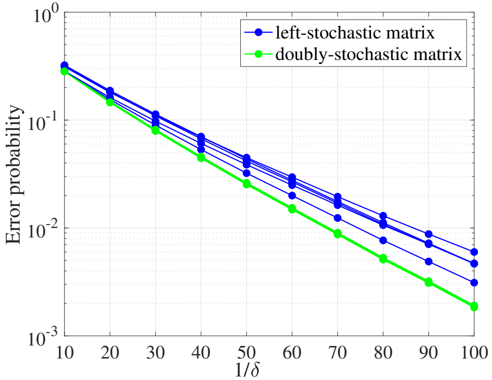

- 5 left-stochastic combination policies and 5 doubly-stochastic combination policies
- Stationary environment, where θ_0 is the underlying truth

Simulation Results

Transient behavior

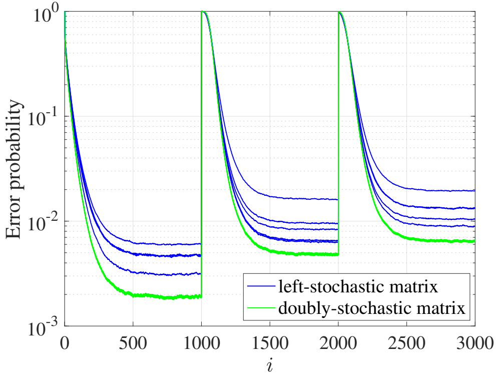

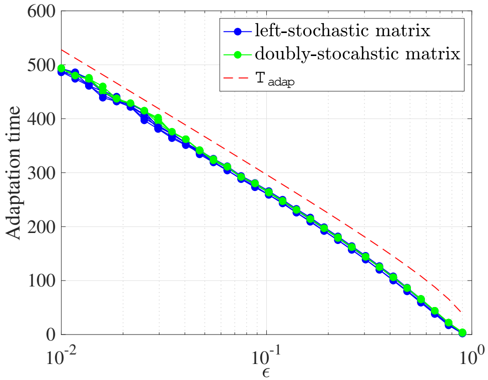

Concluding Remarks

The effect of combination policies on two key performance metrics of adaptive social learning:

1. Error exponent: The best error exponent can be achieved by doubly-stochastic combination policies.
2. Adaptation time: The difference of the adaptation time among different combination policies is negligible if the SNR between hypotheses is small.

References

References

Ping Hu, Virginia Bordignon, Stefan Vlaski, and Ali H. Sayed. Optimal aggregation strategies for social learning over graphs. *arXiv preprint arXiv:2203.07065*, 2022.

Virginia Bordignon, Vincenzo Matta, and Ali H. Sayed. Adaptive social learning. *IEEE Transactions on Information Theory*, 67(9):6053–6081, 2021.

Vincenzo Matta, Paolo Braca, Stefano Marano, and Ali 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.

Thank you!