

Fundamental Social Learning Scaling Law for Tracking Hidden Markov Models

Malek Khammassi
EPFL
Lausanne, Switzerland
malek.khammassi@epfl.ch

Virginia Bordignon[†]
Logmind
Lausanne, Switzerland
virginia.bordignon@alumni.epfl.ch

Vincenzo Matta
University of Salerno
Fisciano, Italy
vmatta@unisa.it

Ali H. Sayed
EPFL
Lausanne, Switzerland
ali.sayed@epfl.ch

Abstract—This paper studies the problem of interconnected agents collaborating to track a dynamic state from partially informative observations, where the dynamic state evolves according to a slowly varying finite-state Markov chain. Although the centralized version of this problem has been extensively studied in the literature, the decentralized setting, particularly in the context of social learning, remains largely underexplored. The main result of this work is to establish that adaptive social learning (ASL), a recent social learning strategy suited for non-stationary environments, achieves the same error probability scaling law as the centralized solution in the rare transitions regime. Theoretical findings are supported by simulations, offering valuable insights into social learning under Markovian state transitions.

Index Terms—Adaptive social learning, large deviations, Markov chain, hidden Markov model, opinion formation.

I. INTRODUCTION AND BACKGROUND

Tracking a hidden dynamic state based on partially informative observations is a fundamental problem in numerous fields, including sensor networks, robotics, and control systems. This task becomes even more challenging when performed in a decentralized manner, where observations are scattered among multiple agents rather than being aggregated by a single processor. In such scenarios, the agents must work together to estimate and track the hidden state that explains their partial and often noisy observations.

Various approaches can be employed to model the underlying dynamic state, with one of the most common being the Markov model. In this context, agents within a network receive local observations that depend on a Markov process, which is not directly observable (hidden); essentially, a hidden Markov model (HMM) [1]. In this work, we focus on a family of Markov chains parameterized by a small positive constant ε , which dictates the speed of the drifts between states. These models are also referred to as *two-time-scale* Markov chains in the literature [2], [3]. Specifically, we consider a discrete-time homogeneous Markov chain θ_i^* (we use the bold notation for random quantities) taking values in the state space Θ with the following transition probabilities:

$$\mathbb{P}[\theta_i^* = \theta' | \theta_{i-1}^* = \theta] = \begin{cases} 1 - \varepsilon q_{\theta\theta'} & \text{if } \theta = \theta' \\ \varepsilon q_{\theta\theta'} & \text{if } \theta \neq \theta' \end{cases}, \quad \theta, \theta' \in \Theta, \quad (1)$$

where $0 < \varepsilon < 1$ represents the drift parameter, $q_{\theta\theta'} \geq 0$ for $\theta' \neq \theta$, and $\sum_{\theta' \neq \theta} q_{\theta\theta'} = q_{\theta\theta}$ for each $\theta \in \Theta$. We assume that the Markov chain is irreducible and aperiodic, thereby guaranteeing the existence of a stationary distribution, which we denote by $p_s = [p_s(1), \dots, p_s(H)]^T$, where $H = |\Theta|$; the cardinality of the set Θ .

As ε decreases, the transition matrix of the Markov chain approaches an identity matrix, indicating piecewise constant behavior

[†]This work was done while Virginia Bordignon was a post-doctoral scholar at the Adaptive Systems Laboratory, EPFL.

with infrequent changes. We refer to these chains as slowly varying or slow Markov chains, and we refer to the regime where ε approaches zero as the rare transitions regime. Such models are relevant in various applications, including linear economic models [4], queuing networks, computer systems, telecommunications, control, and optimization [2], [3], [5].

The problem of tracking a slow Markov chain from partially informative observations using a centralized approach has been extensively studied in the literature [6]–[9]. However, the decentralized setting, particularly within the context of social learning, remains relatively underexplored. One notable exception is [10], which devises the diffusion HMM strategy but requires knowledge of the full state transition matrix. Now, an important result from [6] provides a closed-form characterization of the steady-state error probability of the optimal centralized filtering solution. This characterization illustrates the laws governing the steady-state error probability in the rare transitions regime. Specifically, the steady-state error probability is shown to decay asymptotically as $\varepsilon \log(1/\varepsilon)$ when ε approaches zero.

The objective of this work is to establish that it is possible, through social (i.e., decentralized) learning, to achieve the same steady-state error probability decay law in ε as the optimal centralized filtering solution. To this end, we focus on studying adaptive social learning (ASL), a decentralized decision-making strategy that has been recently proposed in [11] to learn the truth in non-stationary environments. Our previous work [12] established the *consistency* of ASL in tracking slowly varying Markov chains, showing that its steady-state error probability vanishes as ε approaches 0. In this paper, we go further by characterizing the decay law at which this convergence occurs.

ASL is part of a broader family of social learning strategies, which employ Bayesian and graph theories to address decentralized inference problems like hypothesis testing, decision making, classification, and opinion formation in social networks [13]–[20]. While most social learning research focuses on stationary environments, the work in [11] introduced an adaptive strategy that handles non-stationary environments where the true state drifts over time.

We aim to employ ASL for tracking the *state process* θ_i^* introduced in (1) from partially informative observations. To this end, we consider a set of K agents interconnected by a graph. At each time instant $i \geq 1$, each agent $k \in \{1, \dots, K\}$ receives an observation $\mathbf{x}_{k,i}$ belonging to some space \mathcal{X}_k . Agent k adopts a collection of likelihood models, denoted by $\{L_k(\cdot|\theta)\}_{\theta \in \Theta}$ known only to that agent. These models reflect the likelihood assumed by agent k that the observation has been generated by any of the possible states θ . However, the actual likelihood model that governs the distribution of $\mathbf{x}_{k,i}$ is determined by the state process θ_i^* . In other words, conditioned on $\theta = \theta_i^*$, each agent k receives an observation $\mathbf{x}_{k,i}$

distributed according to the likelihood model $L_k(\cdot|\theta_i^*)$.

To track θ_i^* using ASL, each agent k begins with an initial belief vector $\mu_{k,0}$ over the set of plausible hypotheses Θ . At each time step $i > 0$, the agent performs two recursive steps: first, a *self-learning* step to update the previous belief $\mu_{k,i-1}$ with the new observation $\mathbf{x}_{k,i}$, resulting in an *intermediate* belief vector $\psi_{k,i}$; and second, a *combination* step to integrate intermediate beliefs from its neighbors into an updated private belief vector $\mu_{k,i}$. Formally, these steps are written as:

$$\psi_{k,i}(\theta) \propto L_k(\mathbf{x}_{k,i}|\theta)\mu_{k,i-1}^{1-\delta}(\theta) \quad (\text{self-learning}) \quad (2)$$

$$\mu_{k,i}(\theta) \propto \prod_{\ell=1}^K [\psi_{\ell,i}(\theta)]^{a_{\ell k}} \quad (\text{combination}) \quad (3)$$

where the proportionality symbol \propto indicates that the entries of $\mu_{k,i}$ and $\psi_{k,i}$ are normalized to add up to 1. In (2), the positive scalar $\delta \in (0, 1)$ is an adaptation parameter that tunes the network's responsiveness to changes. It controls the weight of the prior belief $\mu_{k,i-1}$ in the update rule (2). A smaller δ gives less weight to prior information, allowing the algorithm to adapt more effectively to changes. In (3), the quantity $a_{\ell k}$ is a nonnegative weight assigned by agent k to the information received from neighbor ℓ , satisfying the following conditions:

$$0 \leq a_{\ell k} \leq 1, \quad \sum_{\ell=1}^K a_{\ell k} = 1. \quad (4)$$

It is shown in [11] that the adaptation time is inversely proportional to δ , while the steady-state error probability approaches zero as δ decreases. This illustrates a trade-off between adaptation and learning: a larger δ allows for quicker adaptation but results in a higher steady-state error probability. These results were derived without considering the nature of the underlying process driving the drifts, assuming instead that the true state remains constant for long enough intervals for agents to learn it before it changes. In this work, we provide a significant advance by establishing that when the underlying state evolves according to a Markov chain, ASL attains the same scaling law (in terms of steady-state error probability as $\varepsilon \rightarrow 0$) as the centralized solution. Notably, to reach this goal, ASL does not need detailed knowledge of the Markov chain dynamics. It only requires rough information regarding the average time between two drifts in the Markov chain to set the adaptation parameter δ , as we will show in the sequel.

Before proceeding with our result, we introduce a set of useful assumptions in the following section.

II. ASSUMPTIONS

We list in the following some assumptions that are traditionally employed in the analysis of social learning systems.

Assumption 1 (Bounded log-likelihood ratios) *There exists a positive constant B such that*

$$\max_{k \in \{1, \dots, K\}} \max_{\theta, \theta' \in \Theta} \sup_{x \in \mathcal{X}_k} \left| \log \frac{L_k(x|\theta)}{L_k(x|\theta')} \right| \leq B. \quad (5)$$

Assumption 1 is automatically satisfied when the observations are discrete random variables with the same finite support.

Assumption 2 (Statistical model) *Let $\mathbf{x}_i \triangleq \{\mathbf{x}_{k,i}\}_{k=1}^K$ collect all observations from across the agents at time i . The joint likelihood at time i satisfies*

$$L(\mathbf{x}_i|\theta_i^*, \dots, \theta_1^*, \mathbf{x}_{i-1}, \dots, \mathbf{x}_1) = L(\mathbf{x}_i|\theta_i^*) \quad (6a)$$

$$= \prod_{k=1}^K L_k(\mathbf{x}_{k,i}|\theta_i^*). \quad (6b)$$

Assumption 3 (Global Identifiability) *For each pair $\theta, \theta' \in \Theta$ such that $\theta \neq \theta'$, there exists at least one agent k such that*

$$D_k(\theta, \theta') \triangleq \mathbb{E}_{\theta} \left[\log \frac{L_k(\mathbf{x}|\theta)}{L_k(\mathbf{x}|\theta')} \right] > 0, \quad (7)$$

where the subscript on the expectation operator means that the random variable \mathbf{x} follows the distribution $L_k(\cdot|\theta)$.

Assumption 4 (Positive initial beliefs) *The initial beliefs of all agents are positive, i.e., $\mu_{k,0}(\theta) > 0$ for all $k \in \{1, \dots, K\}$ and all $\theta \in \Theta$.*

To introduce the main assumption on the combination weights, it is convenient to define the *combination matrix* A , where each entry (ℓ, k) corresponds to the weight $a_{\ell k}$.

Assumption 5 (Primitive matrix) *The combination matrix A is assumed to be primitive [21].*

A sufficient condition for a primitive combination matrix is strong connectivity of the network, meaning that paths with nonzero weights exist in both directions between any two distinct nodes and that the network includes at least one self-loop, which means that some agent k has $a_{kk} > 0$.

Under Assumption 5, the combination matrix A is irreducible [21]. By the Perron-Frobenius theorem [21], [22], A has a spectral radius equal to 1 and a single eigenvalue at 1, associated with the Perron vector π , which is scaled to have all positive entries summing to 1, namely,

$$A\pi = \pi, \quad \sum_{k=1}^K \pi_k = 1, \quad \pi_k > 0 \text{ for } k = 1, 2, \dots, K. \quad (8)$$

III. ASL AND MARKOV CHAINS

We first introduce the steady-state error probability of the optimal centralized filtering solution [6, Theorem 1]. In this setting, we consider a centralized processor that receives all the streaming observations instead of a network of agents. At each time instant $i \geq 1$, the centralized processor receives the observation vector $\mathbf{x}_i = \{\mathbf{x}_{k,i}\}_{k=1}^K$, distributed according to the joint likelihood function $L(\cdot|\theta_i^*) = \prod_{k=1}^K L_k(\cdot|\theta_i^*)$. The processor then uses the set of observations up to that time instant, i.e., $\mathbf{x}_1, \dots, \mathbf{x}_i$, to construct an estimator $\hat{\theta}_i^*$ for θ_i^* , optimal in the sense of minimizing the error probability $\mathbb{P}[\hat{\theta}_i^* \neq \theta_i^*]$. Under Assumptions 1-3, Theorem 1 in [6] states that the steady-state error probability satisfies

$$\lim_{i \rightarrow \infty} \mathbb{P}[\hat{\theta}_i^* \neq \theta_i^*] = \left(\sum_{\theta \in \Theta} p_{\text{stat}}(\theta) \sum_{\theta' \neq \theta} \frac{q_{\theta\theta'}}{D_c(\theta', \theta)} \right) \varepsilon \log \frac{1}{\varepsilon} + o\left(\varepsilon \log \frac{1}{\varepsilon}\right), \quad (9)$$

where $D_c(\theta', \theta) \triangleq \mathbb{E}_{\theta'} \left[\log \frac{L(\mathbf{x}|\theta')}{L(\mathbf{x}|\theta)} \right]$ represents the KL divergence between the *centralized* likelihood models $L(\cdot|\theta)$ and $L(\cdot|\theta')$ for $\theta \neq \theta'$, computed under $L(\cdot|\theta')$. Furthermore, the notation $f_\varepsilon = o(g_\varepsilon)$ means that $\lim_{\varepsilon \rightarrow 0} f_\varepsilon/g_\varepsilon = 0$.

The result in (9) provides a closed-form expression for the leading-order term of the steady-state error probability of the centralized filtering solution, which is optimal in the minimum error probability sense.

Notably, the steady-state error probability diminishes proportionally to $\varepsilon \log \frac{1}{\varepsilon}$, ensuring its convergence to zero as ε vanishes.

Consider now the ASL decentralized setting. Agent k makes a mistake whenever its belief is not maximized at the true hypothesis. Thus, the instantaneous error probability of agent k at time instant i can be expressed as

$$p_{k,i} \triangleq \mathbb{P} \left[\theta_i^* \neq \arg \max_{\theta \in \Theta} \mu_{k,i}(\theta) \right]. \quad (10)$$

Having observed how the centralized steady-state error probability decays in terms of ε , we aim to investigate a choice for the adaptation parameter δ that leads to a similar decay law for the steady-state error probability of the decentralized solution (ASL). Before presenting our main result, we introduce some key quantities. Let $\Phi(\theta', \theta)$ and $t_{\theta', \theta}^*$ correspond to $\Phi(\theta) > 0$ and $t_{\theta}^* < 0$ from [11, Appendix F, Lemma 2], where θ' and θ in our notation represent θ and θ_0 in [11]. While the quantity $t_{\theta', \theta}^*$ is technical in nature, $\Phi(\theta', \theta)$ represents the error exponent, which rules the decay to zero of the steady-state error probability as $\delta \rightarrow 0$, and increases as the hypotheses θ and θ' become more distinguishable. Additionally, we define:

$$\Phi_{\min} \triangleq \min_{\substack{\theta, \theta' \in \Theta \\ \theta' \neq \theta}} \Phi(\theta', \theta), \quad (11)$$

$$\mathcal{D}(\theta', \theta) \triangleq \sum_{k=1}^K \pi_k D_k(\theta', \theta), \quad (12)$$

$$\Delta_{\max} \triangleq \max_{\substack{\theta, \theta' \in \Theta \\ \theta' \neq \theta}} |t_{\theta', \theta}^*| (\mathcal{D}(\theta', \theta) + B). \quad (13)$$

Theorem 1 (Scaling law) *Under Assumptions 1-4, and for any $0 < \nu < 1$, let*

$$\delta = \frac{\Phi_{\min}(1 - \nu)}{\log \frac{1}{\varepsilon}}. \quad (14)$$

Then, we have the following bound on the steady-state error probability:

$$\limsup_{i \rightarrow \infty} p_{k,i} \leq \kappa \varepsilon \log \frac{1}{\varepsilon} + o\left(\varepsilon \log \frac{1}{\varepsilon}\right), \quad (15)$$

where

$$\kappa \triangleq \frac{1}{(1 - \nu)\Phi_{\min}} \log \left(\frac{\Delta_{\max}}{\nu \Phi_{\min}} \right) \sum_{\theta \in \Theta} p_{\text{stat}}(\theta) q_{\theta\theta}. \quad (16)$$

Proof: Due to space limitations, we provide a sketch of the proof. In what follows, the notation $f_{\varepsilon} = O(g_{\varepsilon})$ means that the limit $\lim_{\varepsilon \rightarrow 0} f_{\varepsilon}/g_{\varepsilon}$ is finite.

Let $T_{\varepsilon} > 0$ be an integer function of ε . We can write the instantaneous error probability as follows:

$$\begin{aligned} p_{k,i} &= \mathbb{P} \left[\theta_i^* \neq \arg \max_{\theta \in \Theta} \mu_{k,i}(\theta) \right] \\ &= \mathbb{P} \left[\theta_i^* \neq \arg \max_{\theta \in \Theta} \mu_{k,i}(\theta), \text{ no jumps in } [i - T_{\varepsilon}, i] \right] \\ &+ \mathbb{P} \left[\theta_i^* \neq \arg \max_{\theta \in \Theta} \mu_{k,i}(\theta), \text{ one jump in } [i - T_{\varepsilon}, i] \right] \\ &+ \mathbb{P} \left[\theta_i^* \neq \arg \max_{\theta \in \Theta} \mu_{k,i}(\theta), \text{ at least two jumps in } [i - T_{\varepsilon}, i] \right], \end{aligned} \quad (17)$$

where a jump in the Markov chain θ_i^* is observed at instant j if $\theta_j^* \neq \theta_{j-1}^*$. We see from (17) that $p_{k,i}$ is the sum of three probability

TABLE I: Identifiability setup of the agents.

Agent k	Likelihood model: $L_k(\cdot \theta)$	
	$\theta = 0.1$	$\theta = 0.4$
1, 2, 4, 5, 7, 8	$f_{0.1}$	$f_{0.4}$
3, 9	$f_{0.1}$	$f_{0.1}$
6, 10	$f_{0.4}$	$f_{0.4}$

terms that depend on T_{ε} . We can show that if we set δ as in (14) and

$$T_{\varepsilon} = \left\lceil \rho \log \frac{1}{\varepsilon} \right\rceil, \quad \rho > 0, \quad (18)$$

we can upper bound each of the terms in (17) in the steady-state (i.e. as $i \rightarrow \infty$) to write:

$$\limsup_{i \rightarrow \infty} p_{k,i} \leq o\left(\varepsilon \log \frac{1}{\varepsilon}\right) + \varepsilon T_{\varepsilon} \sum_{\theta \in \Theta} p_{\text{stat}}(\theta) q_{\theta\theta} + O\left((\varepsilon T_{\varepsilon})^2\right), \quad (19)$$

for

$$\rho = \frac{1}{(1 - \nu)\Phi_{\min}} \log \frac{\Delta_{\max}}{\nu \Phi_{\min}}. \quad (20)$$

Then, by using (18) and (20) in (19), we obtain (15). \blacksquare

According to the theory of Bayesian filtering and smoothing [6] [7], the consistency of a learning strategy is examined as the drift parameter $\varepsilon \rightarrow 0$. On the other hand, according to the theory of adaptation [23], consistent learning for ASL was established in [11] as the adaptation parameter $\delta \rightarrow 0$ under stationary conditions. In this work, we establish an interesting link between the drifts in the underlying dynamical model and the adaptation parameter of the social learning strategy. That is, our study examines jointly the role of ε and δ . It can be shown that the former parameter is proportional to the average time between drifts in the Markov chain, whereas the latter is inversely proportional to the adaptation time of ASL. Therefore, our study relates the two fundamental time scales in our system, one dictating the speed of transitions in the dynamical model and the other the inherent adaptation capability of ASL. We find that choosing δ as in (14) yields a closed-form upper bound on ASL's error probability that vanishes as $\varepsilon \rightarrow 0$, ensuring consistency.

The consistency of ASL under slowly varying Markov chain was first established in [12]. Theorem 1 provides a much stronger result: it establishes the asymptotic scaling law for the error probability, showing that ASL scales with $\varepsilon \log(1/\varepsilon)$, which is the same scaling law attained by the optimal centralized scheme. The constant κ for ASL differs from the centralized case, and finding the optimal constant for the decentralized case is an open problem.

The choice of δ in (14) depends on the Markov chain parameter ε and decreases as ε approaches 0. This implies that δ becomes smaller as the environment becomes more stationary, consistent with findings in [11]. Remarkably, the choice of the adaptation parameter does not require any knowledge about the detailed structure of the Markov chain, that is, about the transition weights $q_{\theta, \theta'}$ in (1). It only requires knowledge about ε , which in practice means a rough knowledge about the average time between the drifts [12]. Additionally, δ depends on Φ_{\min} , which represents the worst-case error exponent in ASL [11]. As the classification problem becomes easier, Φ_{\min} increases, allowing for larger δ values to achieve the scaling law in (15). In such cases, ASL can adapt quickly while ensuring a good learning performance. Conversely, when states are harder to distinguish, a smaller δ is necessary to prevent misinterpretation of the true state.




Fig. 1: Strongly connected network with $K = 10$ agents.




Fig. 2: Comparison of steady-state error probabilities between optimal centralized filtering and ASL over ε .

IV. SIMULATION EXAMPLE

We consider the network topology illustrated in Fig. 1. Each agent is assumed to have a self-loop (not shown in the figure). It can be verified that the network is strongly connected, satisfying Assumption 5. Additionally, we design the combination matrix A to be doubly stochastic according to the metropolis policy [23]. We assume that the Markov chain θ_i^* takes on values in the state space $\Theta = \{0.1, 0.4\}$ with the following transition probabilities


$$\mathbb{P}[\theta_i^* = \theta | \theta_{i-1}^* = \theta'] = \begin{cases} 1 - \varepsilon & \text{if } \theta = \theta' \\ \frac{\varepsilon}{H-1} & \text{if } \theta \neq \theta' \end{cases}, \quad \theta, \theta' \in \Theta. \quad (21)$$

The likelihood models of the agents are chosen from the following family of binomial distributions with number of trials $n = 5$ and probability of success $\theta \in \Theta$.


$$f_\theta(x) = \binom{5}{x} \theta^x (1 - \theta)^{5-x}, \quad x \in \{1, \dots, 5\}, \quad (22)$$

where $\binom{5}{x} = \frac{5!}{x!(5-x)!}$ and $f_\theta(x)$ represents the probability of x successes in n trials. Thus, in this setting, the observations $\mathbf{x}_{k,i}$ are drawn from a binomial distribution with a success probability θ_i^* . It is easily seen that the choice of likelihood models in (22) satisfies Assumption 1. Furthermore, we assume that the observations $\mathbf{x}_{k,i}$ are independent across the agents (Assumption 2). We also assume that some agents cannot locally differentiate all pairs of hypotheses, as it can be verified from the identifiability setup in Table I, which satisfies Assumption 3.


We analyze the ASL performance with the adaptation parameter designed using (14). Specifically, we select $\nu = 0.1$ and 0.3 , and compute numerically $\Phi_{\min} \approx 3.80$. For 50 values of ε in the range $[0.005, 0.05]$, we run both the centralized filtering solution and the ASL algorithm for each δ , estimating error probabilities via Monte Carlo simulations with 80000 draws. The steady-state error probability as a function of ε is shown in Fig. 2. We see that



(a) Time evolution of the state process.



(b) Time evolution of the beliefs of centralized filtering.



(c) Time evolution of the beliefs of ASL.

Fig. 3: Belief evolution of the centralized filtering solution and ASL for $\nu = 0.3$ and $\varepsilon = 0.005$.

the steady-state error probability vanishes as ε decreases. Notably, the error probability scales with ε similarly for both ASL and the centralized filtering solution, differing only by the multiplicative constant κ . These simulations confirm the asymptotic scaling law $\varepsilon \log(1/\varepsilon)$, predicted by Theorem 1.

Next, we examine the belief evolution of ASL over time, specifically for $\nu = 0.3$. For comparison purposes, we also examine the belief evolution of the optimal centralized filtering solution, where $\mathbb{P}[\hat{\theta}_i^* = \theta | \mathbf{x}_1, \dots, \mathbf{x}_i]$ is denoted as $\mu_i(\theta)$. Simulations are performed with $\varepsilon = 0.005$ over 1000 time samples, and the beliefs are averaged over 1000 Monte Carlo draws. The resulting belief evolution is visualized in Fig. 3. We see in Fig. 3b that the centralized filtering solution accurately tracks the state process changes. On the other hand, the belief evolution of ASL, depicted in Fig. 3, indicates that ASL effectively tracks the state process evolution. We can also notice that the belief curves for ASL are slightly noisier than the curves pertaining to the centralized solution, which matches well the error probability behavior observed in Fig. 2.

V. CONCLUDING REMARKS

This paper employs social learning to address the decentralized tracking of a dynamic state evolving according to a slow Markov chain. Our key contribution is demonstrating that ASL, a (decentralized) social learning strategy, can achieve the same steady-state error probability decay law in the rare transitions regime as the optimal centralized filtering solution, without requiring full knowledge of the Markov chain's transition probabilities. We design the ASL adaptation parameter δ in terms of ε , the drift parameter of the Markov chain, which can be estimated from the average time the Markov chain remains between two drifts. This approach allows us to derive a closed-form characterization of ASL's steady-state error probability and to identify the laws governing the rare transitions regime.

ACKNOWLEDGMENT

The authors wish to thank Ph.D. student Mert Kayaalp for useful discussions on the topic of this paper.

REFERENCES

- [1] Y. Ephraim and N. Merhav, "Hidden Markov processes," *IEEE Trans. Inf. Theory*, vol. 48, no. 6, pp. 1518–1569, Jun. 2002.
- [2] G. Yin and Q. Zhang, "Singularly perturbed discrete-time markov chains," *SIAM J. Appl. Math.*, vol. 61, no. 3, pp. 834–854, 2000.
- [3] —, *Continuous-Time Markov Chains and Applications: A Singular Perturbation Approach*, 2nd ed. Springer, 2013.
- [4] H. A. Simon and A. Ando, "Aggregation of variables in dynamic systems," *Econometrica*, vol. 29, no. 2, pp. 111–138, 1961.
- [5] P. J. Courtois, *Decomposability: Queueing and Computer System Applications*. Academic Press, 1977.
- [6] R. Khasminskii and O. Zeitouni, "Asymptotic filtering for finite state Markov chains," *Stochastic Processes and their Applications*, vol. 63, pp. 1–10, 1996.
- [7] L. Shue, B. Anderson, and F. De Bruyne, "Asymptotic smoothing errors for hidden Markov models," *IEEE Trans. Signal Process.*, vol. 48, no. 12, pp. 3289–3302, Dec. 2000.
- [8] G. Yin and V. Krishnamurthy, "LMS algorithms for tracking slow Markov chains with applications to hidden Markov estimation and adaptive multiuser detection," *IEEE Trans. Inf. Theory*, vol. 51, no. 7, pp. 2475–2490, 2005.
- [9] G. G. Yin and V. Krishnamurthy, "Least mean square algorithms with Markov regime-switching limit," *IEEE Trans. Autom. Control*, vol. 50, no. 5, pp. 577–593, May 2005.
- [10] M. Kayaalp, V. Bordignon, S. Vlaski, and A. H. Sayed, "Hidden Markov modeling over graphs," in *Proc. IEEE Data Science and Learning Workshop*, Singapore, May 2022, pp. 1–5, see also arXiv:2212.02565, Dec. 2022.
- [11] V. Bordignon, V. Matta, and A. H. Sayed, "Adaptive social learning," *IEEE Trans. Inf. Theory*, vol. 67, no. 9, pp. 6053–6081, Jul. 2021.
- [12] M. Khammassi, V. Bordignon, V. Matta, and A. H. Sayed, "Adaptive social learning for tracking rare transition markov chains," in *Proc. Eur. Signal Process. Conf. (EUSIPCO)*, 2024, pp. 1032–1036.
- [13] A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi, "Non-Bayesian social learning," *Games Econ. Behav.*, vol. 76, no. 1, pp. 210–225, Sep 2012.
- [14] X. Zhao and A. H. Sayed, "Learning over social networks via diffusion adaptation," in *Proc. Asilomar Conf. Signals, Syst. and Comput.*, 2012, pp. 709–713.
- [15] 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, Mar. 2017.
- [16] M. Pooya, T.-S. Alireza, and J. Ali, "A theory of non-bayesian social learning," *Econometrica*, vol. 86, no. 2, pp. 445–490, Mar. 2018.
- [17] A. Lalitha, T. Javidi, and A. D. Sarwate, "Social learning and distributed hypothesis testing," *IEEE Trans. Inf. Theory*, vol. 64, no. 9, pp. 6161–6179, May 2018.
- [18] H. Salami, B. Ying, and A. H. Sayed, "Social learning over weakly connected graphs," *IEEE Trans. Signal and Inf. Process. over Networks*, vol. 3, no. 2, pp. 222–238, Feb. 2017.
- [19] P. Molavi, K. R. Rad, A. Tahbaz-Salehi, and A. Jadbabaie, "On consensus and exponentially fast social learning," in *Amer. Control Conf. (ACC)*, 2012, pp. 2165–2170.
- [20] 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.
- [21] A. H. Sayed, *Inference and Learning from Data*. Cambridge University Press, 2022.
- [22] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge Univ. Press, 1985.
- [23] A. H. Sayed, "Adaptation, learning, and optimization over networks," *Found. Trends Mach. Learn.*, vol. 7, no. 4–5, p. 311–801, jul 2014.