# HIDDEN MARKOV MODELING OVER GRAPHS

*Mert Kayaalp*[⋆], *Virginia Bordignon*[⋆], *Stefan Vlaski*[†], *and Ali H. Sayed*[⋆]

[⋆]École Polytechnique Fédérale de Lausanne (EPFL)
[†] Imperial College London

## ABSTRACT

This work proposes a multi-agent filtering algorithm over graphs for finite-state hidden Markov models (HMMs), which can be used for sequential state estimation or for tracking opinion formation over dynamic social networks. We show that the difference from the optimal centralized Bayesian solution is asymptotically bounded for geometrically ergodic transition models. Experiments illustrate the theoretical findings and in particular, demonstrate the superior performance of the proposed algorithm compared to a state-of-the-art social learning algorithm.

***Index Terms*—** Hidden Markov models, distributed hypothesis testing, social learning, sequential state estimation

## 1. INTRODUCTION AND RELATED WORK

We consider a network of agents observing data that are emitted by some (hidden or) latent state of a dynamic Markov system. The goal is to cooperatively infer and track the time-varying state. The state can be any quantity of interest, e.g., location of a moving object, concentration of air pollutants, or whether it is a sunny or rainy day. The model is general enough and can be used for many applications such as source localization, environmental monitoring, target tracking, navigation, and analyzing opinion formation in social networks.

There exist several works addressing the decentralized multi-agent state estimation problem. One body of work considers distributed Kalman filters [1, 2, 3, 4]. These papers assume linear state dynamics and observations, whereas this work does not make these assumptions. For non-linear system models, some works use a Bayesian framework [5, 6, 7, 8], which we also adopt. It is based on calculating *beliefs*, which are distributions over the set of states. Typically, this involves temporal recursions that consist of computing the new beliefs from the previous beliefs by taking the observed data into account. The Bayesian framework is advantageous in the sense that it uses the complete information of a distribution rather than only some statistics of it such as the mean. Moreover, it enjoys optimality in the minimum mean-square-error and the maximum a-posteriori sense for the single agent mode of operation [9].

Among works that examine distributed Bayesian state estimation, the articles [5, 6] require multiple rounds of communication between agents for every state change. This might be impractical especially when the state changes rapidly. We will show that it is sufficient to communicate once per iteration to attain a bounded difference from the optimal estimator by means of our proposed solution. The works [5, 7] also consider observation likelihood functions that belong to the exponential family of distributions in order to have numerically tractable solutions. In our work, we consider finite dimensional HMMs where the true state/hypothesis can take only a finite set of values. Therefore, we do not need to restrict the distributions to analytically well-behaved families of distributions. Likewise, sub-optimal filters like particle filters [10, 11] are not necessary in this case. On the other hand, the work [8] proposes a distributed Bayesian filtering algorithm that does not require multiple consensus steps at every iteration. However, the agents are assumed to use an *effective* likelihood in place of their own likelihood function. Calculating the effective likelihood includes combining the neighbors' effective likelihoods from the previous time instant with their likelihood from that time instant. As the assumptions in the convergence analysis suggest, this would be beneficial only if the likelihood functions are changing slowly over time. This limits the maximum rate of state transition the algorithm can track. In comparison, our algorithm combines the likelihoods with *time-adjusted* priors, in which the neighbors' beliefs at the previous time instant are combined to compute the priors. Consequently, we get bounds (Theorem 1) that hold for a large class of transition models, namely, geometrically ergodic models, which includes rapidly mixing Markov chains.

Combining the neighbors' beliefs is also useful for modeling opinion formation over social networks. In these models, agents' beliefs/opinions over the set of states/hypotheses are formed based on their local observations and their interactions with the other agents [12, 13, 14, 15]. A line of social learning algorithms comprises of two iterative steps. First, each agent revises its belief via a Bayesian update based on new private data. Second, agents aggregate their neighbors' information into their beliefs with a distributed learning algorithm like consensus [14, 16, 17, 18] or diffusion

This work was supported in part by SNSF grant 205121-184999. Emails: mert.kayaalp@epfl.ch, virginia.bordignon@epfl.ch , s.vlaski@imperial.ac.uk , ali.sayed@epfl.ch .

[15, 19, 20]. A common assumption in these works is that the state of nature is fixed. In many practical applications, however, the state is time-varying. Linear transition models are considered in [21, 22, 23]. However, to tackle the possible drifts in the state of nature, an *adaptive social learning* (ASL) strategy was proposed in [24]. This algorithm enables the agents to respond to the changes faster. Nevertheless, it does not take the transition model into account. In many scenarios, the current true state will make some states more likely to occur than others in the future. By exploiting knowledge about the system dynamics, the true state can be tracked better. Indeed, we confirm this statement experimentally by comparing ASL against our proposed algorithm in Section 4.

**Contributions**. We propose a distributed Bayesian HMM filtering algorithm for detecting a time-varying state in Section 2.2. The algorithm requires only one round of communication per state change and takes advantage of prior information about state dynamics, which allow tracking rapidly changing states. We examine how close this distributed strategy gets to the optimal centralized solution in Section 3. More specifically, an asymptotic bound on the expected Kullback-Leibler (KL) divergence between the centralized belief and the agent-specific beliefs is established in Theorem 1 for geometrically ergodic transition models. Corollary 1 relates this bound to belief values evaluated at the true hypothesis. Simulation results in Section 4 support the theoretical results and compare the proposed algorithm with a state-of-the-art algorithm, ASL [24].

## 2. ALGORITHM DESCRIPTION

A network of $K$ communicating agents are exchanging beliefs with each other in order to keep track of the underlying state of nature, which is allowed to evolve over time. The belief of agent $k \in \mathcal{N}$, at time $i$, is denoted by $\mu_{k,i}$ and it is a probability simplex. The value of $\mu_{k,i}(\theta)$ represents the probability that agent $k$ believes the hypothesis $\theta \in \Theta$ is the true hypothesis at time $i$. The true hypothesis at that same time instant is denoted by $\boldsymbol{\theta}_i^\circ \in \Theta$, and it is assumed to belong to a finite set of $H$ hypotheses, $\Theta = \{0, 1, .., H-1\}$. Note that we are using boldface letters to refer to random variables. The transition model $T$, which is assumed to be known to the agents, is a Markov chain and we use the notation:

$$T(\theta_i|\theta_{i-1}) \triangleq \mathbb{P}(\boldsymbol{\theta}_i^\circ = \theta_i|\boldsymbol{\theta}_{i-1}^\circ = \theta_{i-1}) \tag{1}$$

At instant $i$, each agent $k$ observes a private observation $\boldsymbol{\xi}_{k,i}$ distributed according to the agent-specific likelihood function $L_k(\boldsymbol{\xi}_{k,i}|\boldsymbol{\theta}_i^\circ)$, which is known to agent $k$. The likelihoods can be probability mass or density functions depending on whether observations are discrete or continuous.

**Assumption 1.** *[5, 23] The observations are assumed to be independent across agents given the state. Denoting the joint observations by $\boldsymbol{\xi}_i \triangleq \{\boldsymbol{\xi}_{k,i}\}_{k=1}^K$ and its distribution by $L^\star(\boldsymbol{\xi}_i|\boldsymbol{\theta}_i^\circ)$, we have for all $\xi_i \triangleq \{\xi_{k,i}\}_{k=1}^K$ and $\theta_i \in \Theta$:*

$$L^\star(\xi_i|\theta_i) = \prod_{k=1}^K L_k(\xi_{k,i}|\theta_i) \tag{2}$$

□

Agents can communicate with each other once per iteration. We have the following assumption on the communication topology.

**Assumption 2.** *The communication topology underlying the network is a strongly-connected graph [25]. This means that the combination matrix $A \triangleq [a_{\ell k}]$ is a primitive matrix. The coefficient $a_{\ell k}$ weights the information sent by agent $\ell$ to $k$ and is nonzero if, and only if, $\ell \in \mathcal{N}_k$ (i.e., for every agent $\ell$ in the neighborhood of $k$). The matrix $A$ is doubly-stochastic and symmetric, i.e., it satisfies:*

$$A\mathbb{1}_K = \mathbb{1}_K, \quad A = A^\mathsf{T} \tag{3}$$

□

We also assume a regularity condition on the likelihood functions.

**Assumption 3.** *For each agent $k$, log-likelihoods are bounded in absolute value, namely:*

$$|\log L_k(\cdot|\cdot)| \leq \alpha \tag{4}$$

□

For instance, this assumption is satisfied for truncated Gaussian likelihoods.

### 2.1. Optimal Centralized Belief Recursion

Given the observation history of all agents $\boldsymbol{\mathcal{F}}_i \triangleq \{\boldsymbol{\xi}_j\}_{j=1}^i$, the posterior distribution of the true hypothesis at time $i$ is denoted by:

$$\boldsymbol{\mu}_i^\star(\theta_i) \triangleq \mathbb{P}(\boldsymbol{\theta}_i^\circ = \theta_i|\boldsymbol{\mathcal{F}}_i) \tag{5}$$

This posterior satisfies the optimal Bayesian filtering recursion [9]:

$$\boldsymbol{\mu}_i^\star(\theta_i) = \frac{L^\star(\boldsymbol{\xi}_i|\theta_i)\boldsymbol{\eta}_i^\star(\theta_i)}{\sum_{\theta_i'} L^\star(\boldsymbol{\xi}_i|\theta_i')\boldsymbol{\eta}_i^\star(\theta_i')} \tag{6}$$

where $\boldsymbol{\eta}_i^\star(\theta_i)$ is the time-adjusted prior at time instant $i$ and given by:

$$\boldsymbol{\eta}_i^\star(\theta_i) \triangleq \sum_{\theta_{i-1}} T(\theta_i|\theta_{i-1})\boldsymbol{\mu}_{i-1}^\star(\theta_{i-1}) \tag{7}$$

In the sequel, we study how close the beliefs generated by the proposed distributed algorithm get to the above centralized posterior (6).

### 2.2. Diffusion HMM Strategy

Agents across the network need to update their beliefs based on their local streaming observations, as well as exchange their beliefs with each other in order to track the true state in the face of stochastic and dynamic conditions. To do so, we propose a *social* HMM filtering algorithm that is based on the diffusion strategy for cooperation over networks [25]. At every time instant $i$, agents first revise their belief at $i-1$ via the Chapman-Kolmogorov equation [9]:

$$\boldsymbol{\eta}_{k,i}(\theta_i) = \sum_{\theta_{i-1}} T(\theta_i|\theta_{i-1})\boldsymbol{\mu}_{k,i-1}(\theta_{i-1}) \qquad \text{(Evolve)} \quad (8)$$

Then, each agent $k$ forms an intermediate belief *locally* by a $\gamma$-*scaled* Bayesian update based on the received data:

$$\boldsymbol{\psi}_{k,i}(\theta_i) = \frac{(L_k(\boldsymbol{\xi}_{k,i}|\theta_i))^{\gamma}\boldsymbol{\eta}_{k,i}(\theta_i)}{\sum_{\theta_i'}(L_k(\boldsymbol{\xi}_{k,i}|\theta_i'))^{\gamma}\boldsymbol{\eta}_{k,i}(\theta_i')} \qquad \text{(Adapt)} \quad (9)$$

where $\gamma > 0$ is a step-size that scales the newly arrived data against prior information. Finally, agents combine the intermediate beliefs of their neighbors into their updated belief:

$$\boldsymbol{\mu}_{k,i}(\theta_i) = \frac{\exp\{\sum_{\ell\in\mathcal{N}_k} a_{\ell k}\log\boldsymbol{\psi}_{\ell,i}(\theta_i)\}}{\sum_{\theta_i'}\exp\{\sum_{\ell\in\mathcal{N}_k} a_{\ell k}\log\boldsymbol{\psi}_{\ell,i}(\theta_i')\}} \qquad \text{(Combine)} \quad (10)$$

Repeatedly exchanging and fusing beliefs will allow the local information to diffuse throughout the network.

Note that the proposed algorithm recovers standard log-linear social learning algorithms [16, 17, 19, 20] when agents perform local Bayesian updates with $\gamma = 1$ and the true hypothesis is *fixed*:

$$T(\theta_i|\theta_{i-1}) = \begin{cases} 1, & \theta_i = \theta_{i-1} \\ 0, & \theta_i \neq \theta_{i-1} \end{cases}$$

Moreover, the agents' beliefs will match the optimal centralized belief if $a_{\ell k} = 1/K\ \forall \ell, k \in \mathcal{N}$ (i.e., when the network is fully-connected), all priors are equal ($\mu_0^\star = \mu_{k,0}\ \forall k \in \mathcal{N}$) and $\gamma = K$. These two special cases motivate us to use a general step-size $\gamma > 0$.

## 3. MAIN RESULTS

To avoid discarding any hypothesis in the beginning, we have an assumption on the initial values of the beliefs.

**Assumption 4.** *[8, 16] We assume that all initial beliefs are strictly positive at every hypothesis, i.e., for each hypothesis $\theta \in \Theta$ and for each agent $k$, $\mu_{k,0}(\theta) > 0, \mu_0^\star(\theta) > 0$.* □

We can assess the performance of the algorithm (8)–(10) relative to the optimal centralized solution (6) by considering time-varying risks of the form:

$$J_i(\boldsymbol{\mu}_{k,i}) \triangleq \mathbb{E}_{\mathcal{F}_i} D_{\mathrm{KL}}(\boldsymbol{\mu}_i^\star||\boldsymbol{\mu}_{k,i}) \quad (11)$$

and

$$\widetilde{J}_i(\boldsymbol{\eta}_{k,i}) \triangleq \mathbb{E}_{\mathcal{F}_{i-1}} D_{\mathrm{KL}}(\boldsymbol{\eta}_i^\star||\boldsymbol{\eta}_{k,i}) \quad (12)$$

for each agent $k$, where $D_{\mathrm{KL}}(\cdot||\cdot)$ represents the KL divergence, and $\mathbb{E}_{\mathcal{F}_i}$ denotes the expectation over $\mathcal{F}_i$ with respect to $\mathbb{P}$. The risk in (11) measures the disagreement between agent $k$ and the centralized belief for the true state at time instant $i$ after the observations are emitted from that state. In comparison, the risk in (12) is the disagreement before the observations are emitted. In other words, $J_i$ is the *posterior* divergence while $\widetilde{J}_i$ is the divergence of time-adjusted *priors*.

The effect of the transition model will arise via the Dobrushin coefficient $\kappa(T) \in [0,1]$ defined as follows [26, 27, 9]:

$$\kappa(T) = \sup_{\theta,\theta'\in\Theta} D_{\mathrm{TV}}\Big(T(\cdot|\theta), T(\cdot|\theta')\Big) \quad (13)$$

where $D_{\mathrm{TV}}(\cdot,\cdot)$ represents the total variation distance. For instance, for a binary symmetric channel:

$$T(\theta_i|\theta_{i-1}) = \begin{cases} 1-\delta, & \theta_i = \theta_{i-1} \\ \delta, & \theta_i \neq \theta_{i-1} \end{cases}$$

we have $\kappa(T) = |1-2\delta|$. In general, the closer the Dobrushin coefficient is to zero, the faster the forgetting of the initial conditions will be.

For ease of notation, we define the column vector consisting of marginal likelihoods over hypotheses and agents of a joint observation $\xi_j$ as:

$$\mathcal{L}_{\xi_j} \triangleq \mathrm{col}\Big\{\mathrm{col}\Big\{\log L_\ell(\xi_{\ell,j}|\theta_j)\Big\}_{\theta_j=1}^{H}\Big\}_{\ell=1}^{K} \quad (14)$$

We first establish that the risk functions are asymptotically bounded for each agent $k$.

**Theorem 1.** *Under Assumptions 1, 2, 3, 4 and geometrically ergodic state transition models [9], i.e., $\kappa(T) < 1$, for each agent $k$ the risks are asymptotically bounded, namely:*

$$\limsup_{i\to\infty} J_i(\boldsymbol{\mu}_{k,i}) \leq \frac{2K\gamma\lambda}{1-\kappa(T)} \sup_{t>0} \mathbb{E}_{\xi_t}\|\boldsymbol{\mathcal{L}}_{\xi_t}\|_\infty \quad (15)$$

*and*

$$\limsup_{i\to\infty} \widetilde{J}_i(\boldsymbol{\eta}_{k,i}) \leq \frac{2\kappa(T)K\gamma\lambda}{1-\kappa(T)} \sup_{t>0} \mathbb{E}_{\xi_t}\|\boldsymbol{\mathcal{L}}_{\xi_t}\|_\infty \quad (16)$$

*where $\lambda \triangleq \max\{|1-\frac{K}{\gamma}|, \rho_2\}$, and $\rho_2$ is the absolute value of the second largest magnitude eigenvalue of $A$.*

*Proof.* Omitted due to space limitations. □

Note that if the transition Markov chain is stationary, we have, for any $t' > 0$:

$$\sup_{t>0} \mathbb{E}_{\xi_t}\|\boldsymbol{\mathcal{L}}_{\xi_t}\|_\infty = \mathbb{E}_{\xi_{t'}}\|\boldsymbol{\mathcal{L}}_{\xi_{t'}}\|_\infty \quad (17)$$

The bounds are tight in the sense that they are equal to zero as expected when the centralized solution is matched with $a_{\ell k} = 1/K \ \forall \ell, k \in \mathcal{N}$ ($\rho_2 = 0$), equal priors ($\mu_0^\star = \mu_{k,0}$ $\forall k \in \mathcal{N}$), and $\gamma = K$ . In fact, the bounds (15) and (16) do not depend on the initial beliefs as long as they satisfy Assumption 4. Geometric ergodicity is sufficient to forget the initial conditions. Specifically, for $\rho_2 = 0$ and $\gamma = K$, the asymptotic risk is zero which means that the filter is *stable*.

Furthermore, for $\gamma \to K$, the risks are proportional to the mixing rate of the graph, $\rho_2$. This factor underlines the benefit of cooperation. More connected graphs, with smaller $\rho_2$, will track the centralized solution better whereas sparse networks or non-cooperative agents will have higher deviation from the optimal. The bounds are also proportional to the network size. The disagreement between the agents and the optimal centralized solution, which has access to all data of agents, increases with the number of agents.

Observe that if the transition Markov chain is mixing very fast such that $\kappa(T) = 0$, the bound in (16) goes to zero. This is expected because the transition model, in equations (7) and (8), will output the same distributions and the divergence between them will vanish. Therefore, (16) captures the effect of ergodicity accurately. However, the bounds are still not tight enough in the sense that the effect of observations is not sufficiently reflected. For instance, if the state of nature is fixed, we have $\kappa(T) = 1$, which is not a geometrically ergodic model. However, from the standard social learning literature [14, 15, 16, 17], we know that true state can be learned. Bounds that address both ergodicity and informativeness of the measurements is an interesting future work.

Now assume that the risks are asymptotically bounded with $\limsup_{i\to\infty} J_i(\boldsymbol{\mu}_{k,i}) \leq B$ , as suggested by Theorem 1. Next, we relate this upper bound to belief values at the true hypothesis.

**Corollary 1.** *Define the variance:*

$$\begin{aligned} &Var\Big(\log \frac{\boldsymbol{\mu}_i^\star(\boldsymbol{\theta}_i^\circ)}{\boldsymbol{\mu}_{k,i}(\boldsymbol{\theta}_i^\circ)}\Big) \\ &\triangleq \mathbb{E}_{\mathcal{F}_i,\theta_i^\circ}\Big|\log \frac{\boldsymbol{\mu}_i^\star(\boldsymbol{\theta}_i^\circ)}{\boldsymbol{\mu}_{k,i}(\boldsymbol{\theta}_i^\circ)} - \mathbb{E}_{\mathcal{F}_i,\theta_i^\circ}\Big[\log \frac{\boldsymbol{\mu}_i^\star(\boldsymbol{\theta}_i^\circ)}{\boldsymbol{\mu}_{k,i}(\boldsymbol{\theta}_i^\circ)}\Big]\Big|^2 \end{aligned} \tag{18}$$

*Then, with probability at least:*

$$p \triangleq 1 - \frac{Var\Big(\log \frac{\boldsymbol{\mu}_i^\star(\boldsymbol{\theta}_i^\circ)}{\boldsymbol{\mu}_{k,i}(\boldsymbol{\theta}_i^\circ)}\Big)}{\epsilon^2} \tag{19}$$

*where $\epsilon$ is an arbitrary positive constant such that $p \in (0, 1]$, the agent and centralized beliefs evaluated at the true hypothesis $\theta_i^\circ$ satisfy the following relation as $i \to \infty$:*

$$\mu_{k,i}(\theta_i^\circ) \geq \mu_i^\star(\theta_i^\circ) \exp\Big\{-\epsilon - B\Big\} \tag{20}$$

*Proof.* Omitted due to space limitations. □

This is a guarantee on the beliefs evaluated at the true hypothesis. For small risks (11) and variances (18), agents will assign a belief value which is close to the value assigned by the optimal belief.

## 4. SIMULATION RESULTS

We consider first a $10-$agent network whose topology is displayed in Fig. 1. The combination matrix is given by the Metropolis rule [28, 25], resulting in a doubly-stochastic and symmetric matrix with $\rho_2 = 0.86$.

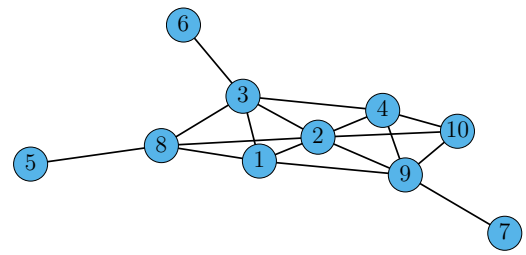


**Fig. 1**. Network diagram.

The network would like to track the true state of nature from a set of two hypotheses, $\Theta = \{0, 1\}$. We assume that all agents possess the same family of truncated Gaussian likelihoods :

$$L_k(\xi|\theta) = \begin{cases} \frac{1}{Z_\theta}\frac{1}{\sqrt{2\pi}} \exp\Big\{-\frac{(\xi-(\theta+1))^2}{2}\Big\}, & -1 \leq \xi \leq 2 \\ 0, & \text{otherwise} \end{cases}$$

for all $k = 1, 2, \ldots, K$, where $Z_\theta$ is the normalization constant:

$$Z_\theta \triangleq \int_{-1}^{2} \frac{1}{\sqrt{2\pi}} \exp\Big\{-\frac{(\xi-(\theta+1))^2}{2}\Big\} d\xi \tag{21}$$

Notice that, in this case, Assumption 3 is satisfied. The hidden state is assumed to be a Markovian random variable, whose transition matrix is given by:

$$T(\theta_i|\theta_{i-1}) = \begin{cases} 0.9, & \theta_i = \theta_{i-1} \\ 0.1, & \theta_i \neq \theta_{i-1} \end{cases}$$

which corresponds to the Dobrushin coefficient $\kappa(T) = 0.8$.

We first compare the performance of the proposed diffusion HMM filter (dHMM) with the centralized HMM filter (cHMM) for a particular realization of hidden states $\boldsymbol{\theta}_i^\circ$. This comparison is seen in the middle panel of Fig. 2. Notice that both dHMM, with choice $\gamma = K$, and cHMM behave similarly, which supports Corollary 1. They show a remarkable capacity of tracking the abrupt changes in the hidden state. In the same figure, we have also included the behavior of the

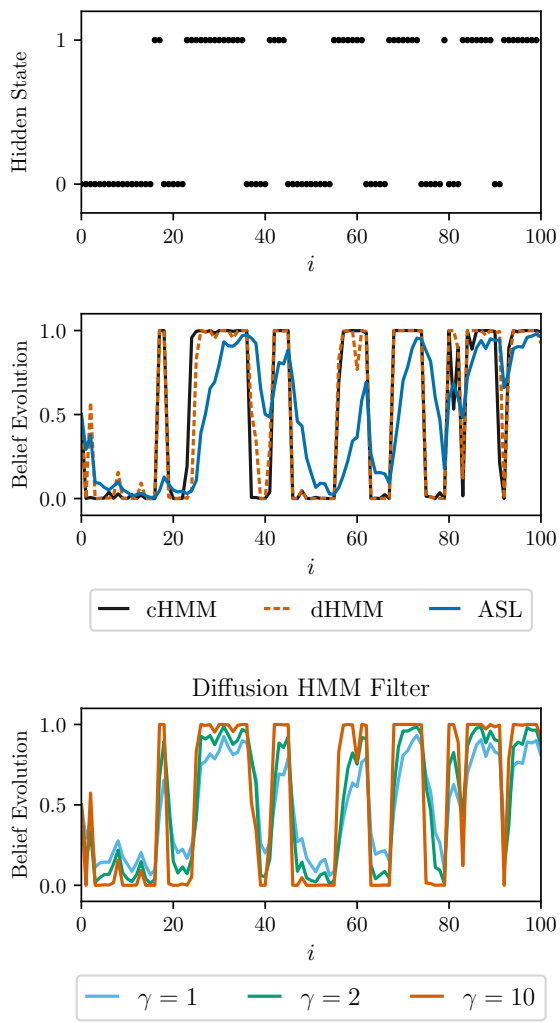


**Fig. 2**. *Top panel*: Evolution of the hidden state. *Middle panel*: Evolution of beliefs over time for different methods (centralized HMM, diffusion HMM, and ASL). *Bottom panel*: Evolution of beliefs over time for the diffusion HMM filter.

ASL algorithm with a choice of step size $\delta = 0.1$ . ASL does not utilize the transition model knowledge. Hence, it is slower to respond to changes compared to cHMM and dHMM.

In the bottom panel of Fig. 2, we explore different choices of $\gamma$ and observe the evolution of beliefs for the distributed HMM filter. We can see that as $\gamma$ approaches $K = 10$, the tracking performance of the algorithm increases, approaching the centralized performance.

In Fig. 3, we observe the evolution of the risk $J_i(\boldsymbol{\mu}_{k,i})$ over time for three different graph topologies: $i)$ a very sparse topology with $\rho_2 = 0.97$, $ii)$ the topology seen in Fig. 1 with $\rho_2 = 0.86$, and $iii)$ a fully connected network with $\rho_2 = 0$. The risk was approximated by averaging 1000 Monte Carlo simulations with the choice of $\gamma = K$. Fig. 3 shows that as the second larger eigenvalue $\rho_2$ decreases, the risk approaches

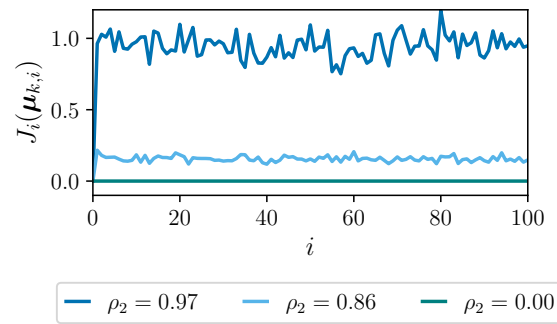


**Fig. 3**. Evolution of the risk over time, computed with 1000 Monte Carlo experiments for three different network topologies.

**Table 1**. Number of agents and average risks over networks

| $K$ | $\frac{1}{K}\sum_{k=1}^{K} J_\infty(\boldsymbol{\mu}_{k,\infty})$ | $\rho_2$ |
|---|---|---|
| 10 | 0.53 | 0.86 |
| 20 | 0.83 | 0.83 |
| 30 | 1.17 | 0.81 |
| 40 | 1.77 | 0.80 |
| 70 | 2.69 | 0.77 |

zero. This observation is expected in view of Theorem 1.

Having observed that the filter performance increases with step-size chosen around the number of agents, i.e., $\gamma \to K$, and more graph connectivity, i.e. $\rho_2 \to 0$, we finally compare the effect of the network size on the risk function. In Table 1, the average risk function over the network with respect to different numbers of agents can be found. For all cases, we set $\gamma = K$. It is hard to get network configurations with exactly the same $\rho_2$'s for different sizes. Therefore, for a fair comparison, we increase the connectivity of the larger graphs, i.e., higher $K$'s are associated with smaller $\rho_2$'s. Despite this fact, the average risk increases with increasing number of agents, which supports Theorem 1. This means that the deviation from the optimal centralized algorithm is higher for larger networks, for fixed graph mixing rates.

## 5. REFERENCES


[1] R. Olfati-Saber, "Distributed Kalman filtering for sensor networks," in *Proc. IEEE Conference on Decision and Control*, New Orleans, LA, USA, 2007, pp. 5492–5498.

[2] U. A. Khan and J. M. F. Moura, "Distributing the Kalman filter for large-scale systems," *IEEE Transactions on Signal Processing*, vol. 56, no. 10, pp. 4919–4935, 2008.

[3] F. S. Cattivelli and A. H. Sayed, "Diffusion strategies

for distributed Kalman filtering and smoothing," *IEEE Transactions on Automatic Control*, vol. 55, no. 9, pp. 2069–2084, 2010.

[4] S. P. Talebi, S. Werner, V. Gupta, and Y.-F. Huang, "On stability and convergence of distributed filters," *IEEE Signal Processing Letters*, vol. 28, pp. 494–498, 2021.

[5] O. Hlinka, O. Slučiak, F. Hlawatsch, P. M. Djurić, and M. Rupp, "Likelihood consensus and its application to distributed particle filtering," *IEEE Transactions on Signal Processing*, vol. 60, no. 8, pp. 4334–4349, 2012.

[6] G. Battistelli and L. Chisci, "Kullback–Leibler average, consensus on probability densities, and distributed state estimation with guaranteed stability," *Automatica*, vol. 50, no. 3, pp. 707–718, 2014.

[7] K. Dedecius and P. M. Djurić, "Sequential estimation and diffusion of information over networks: A Bayesian approach with exponential family of distributions," *IEEE Transactions on Signal Processing*, vol. 65, no. 7, pp. 1795–1809, 2017.

[8] S. Bandyopadhyay and S. Chung, "Distributed Bayesian filtering using logarithmic opinion pool for dynamic sensor networks," *Automatica*, vol. 97, pp. 7–17, 2018.

[9] V. Krishnamurthy, *Partially Observed Markov Decision Processes: From Filtering to Controlled Sensing*, Cambridge University Press, 2016.

[10] A. H. Sayed, P. M. Djurić, and F. Hlawatsch, "Distributed Kalman and particle filtering," in *Cooperative and Graph Signal Processing*, P. M. Djurić and C. Richard, Eds., pp. 169–207. Academic Press, 2018.

[11] G. Papa, R. Repp, F. Meyer, P. Braca, and F. Hlawatsch, "Distributed Bernoulli filtering using likelihood consensus," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 5, no. 2, pp. 218–233, 2019.

[12] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar, "Bayesian learning in social networks," *The Review of Economic Studies*, vol. 78, no. 4, pp. 1201–1236, 2011.

[13] 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 Processing Magazine*, vol. 30, no. 3, pp. 43–57, 2013.

[14] 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, 2012.

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

[16] A. Nedić, 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, 2017.

[17] 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, 2018.

[18] D. Luengo, L. Martino, V. Elvira, and M. Bugallo, "Efficient linear fusion of partial estimators," *Digital Signal Processing*, vol. 78, pp. 265–283, July 2018.

[19] 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.

[20] K. Ntemos, V. Bordignon, S. Vlaski, and A. H. Sayed, "Deception in social learning," *arXiv preprint arXiv:2103.14729*, 2021.

[21] D. Acemoglu, A. Nedic, and A. Ozdaglar, "Convergence of rule-of-thumb learning rules in social networks," in *IEEE Conference on Decision and Control*, 2008, pp. 1714–1720.

[22] R. M. Frongillo, G. Schoenebeck, and O. Tamuz, "Social learning in a changing world," in *Internet and Network Economics*. 2011, pp. 146–157, Springer Berlin Heidelberg.

[23] S. Shahrampour, S. Rakhlin, and A. Jadbabaie, "Online learning of dynamic parameters in social networks," in *Advances in Neural Information Processing Systems*, 2013, vol. 26.

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

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

[26] R. L. Dobrushin, "Central limit theorem for nonstationary Markov chains. I," *Theory of Probability & Its Applications*, vol. 1, no. 1, pp. 65–80, 1956.

[27] Y. Polyanskiy and Y. Wu, "Strong data-processing inequalities for channels and Bayesian networks," in *Convexity and Concentration*, New York, NY, 2017, pp. 211–249, Springer New York.

[28] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, "Equation of state calculations by fast computing machines," *The Journal of Chemical Physics*, vol. 21, no. 6, pp. 1087–1092, Jun 1953.